GHOST DAY

Applied Machine Learning Conference

Conference Booklet

Full Content List

Poznan University of Technology

09-10 May 2025

ORGANIZED BY

About the event

GHOST DAY
Applied Machine Learning Conference

GHOST Day: AMLC aims at creating a friendly and vivid space for the machine learning enthusiasts to exchange experiences and broaden their knowledge of the rapidly changing discipline of data analysis.

Our speakers include recognized representatives of the scientific community publishing at top-tier global conferences such as NeurIPS or ICML and many experts from leading companies building machine learning-based products.

What is our goal?

The conference is not focused on any specific programming language or technology, but on the techniques, methods, and algorithms. It is also a gateway for talented engineers wanting to enter the fascinating world of artificial intelligence. We give our guests the opportunity to meet companies operating in the industry, participate in presentations at an intermediate level, and present their own projects during the poster session.

What is GHOST?

Group of Horribly Optimistic STatisticians (GHOST) is a student organization with roots at Poznań University of Technology focused on conducting research and studying machine learning. GHOST is divided into study groups focused on fields like TinyML, Probabilistic modeling, or NLP. You may have also heard about us at various hackathons (we are good at them :)), Kaggle competitions, and other conferences.

Location

The conference takes place at Poznań University of Technology – leader of EUNICE, the European University and one of the two Polish institutions ranked in prestigious Shanghai Ranking 2019 in Computer Science Engineering (on positions 401-500, ex aequo with University of Warsaw) and the first Polish higher education institution offering bachelor degree studies on Artificial Intelligence.

Organizing Committee

GHOST DAY
Applied Machine Learning Conference

Oskar Kałuziak
Project Leader

Dmytro Romaniv
Assistant Project Leader

Łukasz Sztukiewicz
Board Member

Michał Wiliński
SpecOps Coordinator

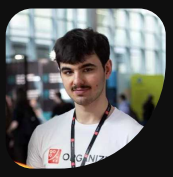

Jakub Kubiak
Finance Coordinator

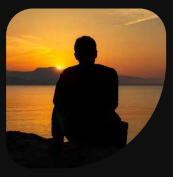

Wiktor Kamzela
Speakers Coordinator

Rafał Malcewicz
Sponsors Coordinator

Jagoda Grzybowska
Marketing Coordinator

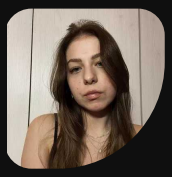

Agata Chruśniak
Marketing Coordinator

Bartek Szaliński
Main Designer

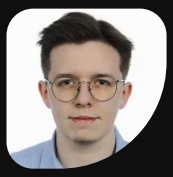

Alex Terentowicz
SpecOps Coordinator

Kostek Kordas
Finance Team

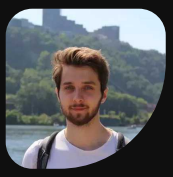

Ignacy Stępka
Speakers Team

Joanna Cicha
Finance Team

Piotr Wyrwiński
Finance Team

Konrad Szewczyk
Speakers Team

Jakub Bilski
Marketing Team

Julia Lorenz
Finance Team

Zuzanna Tabisz
Finance Team

Karol Cyganik
Sponsors Team

Organizing Committee

GHOST DAY
Applied Machine Learning Conference

Maciej Kurzawa
Sponsors Team

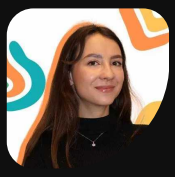

Aleksandra Benowska
Marketing Team

Natalia Król
Marketing Team

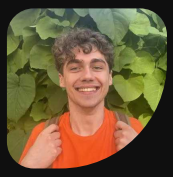

Michał Skrzypek
Marketing Team

Jagoda Janowska
Finance Team

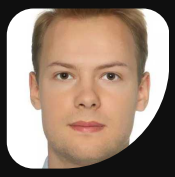

Tomasz Kazulak
Sponsors Team

Marcin Sitarek
Finance Team

Jakub Krym
SpecOps Team

Adam Dobosz
Speakers Team

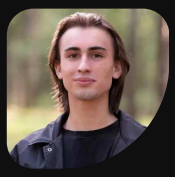

Kanstantsin Laurysiuk
Speakers Team

Viacheslav Krilyshyn
Finance Team

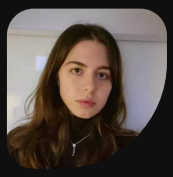

Julia Korzekwa
Finance Team

Jakub Madejczyk
Speakers Team

Kamil Pluciński
Marketing Team

Grzegorz Możdzyński
Finance Team

Patryk Maciejewski
SpecOps Team

Advisory Board

GHOST DAY
Applied Machine Learning Conference

Alicja Kwasniewska, Ph.D.
Amazon Web Services

Artur Dubrawski, Prof.
Carnegie Mellon University

Halima Bouzidi, Ph.D.
University of California, Irvine

Hanna Klimczak-Plucińska
Google DeepMind

Jerzy Stefanowski, Prof.
Poznan University of Technology

Krzysztof Krawiec, Prof.
Poznan University of Technology

Krzysztof Walas, Ph.D.
IDEAS NCBR, PUT

Mateusz Lango, Ph.D.
Charles University

Michal Valko, Ph.D.
Stealth Startup, Inria, ENS

Wojciech Kotłowski, Prof.
Poznan University of Technology

Partners & Sponsors

GHOST DAY
Applied Machine Learning Conference

Honorary Patronages

Minister of Science and Higher Education
Republic of Poland

Minister of Digital Affairs
Republic of Poland

Ministry of Economic Development and Technology
Republic of Poland

Polish Artificial Intelligence Society

POZnań*
Honorary Patronage
of the Mayor of Poznań

IRIM

WOMEN IN AI

PAN
POLISH ACADEMY OF SCIENCES

Platinum Sponsor

allegro

Partners & Sponsors

GHOST DAY
Applied Machine Learning Conference

Gold Sponsors

AMD
data rabbit
Pearson
RESEARCH
OLX
Google
Google DeepMind

Silver Sponsors

Sii
CAICS

Media Partners

Group of Horribly Optimistic STatisticians
ML WORKOUT
ML in PL ASSOCIATION
FUNDACJA KANAŁ STUDENCKI
prg.ai
AI LAB
Machine Learning Prague
PyData
Ogólnopolska Konferencja Studentów Matematyki
θβlicZε
WARSAW.AI
Polish ML Community
QUANTUM AI
mlgdańsk
GOLEM
AI EMBASSY
KNSM
ATLAS
BLISS
DATA SCIENCE
GAUSS
GRADIENT

Agenda

GHOST DAY
Applied Machine Learning Conference

DAY 1: 09-05-2025

8:00 - 8:45	-1 FLOOR	**Registration**
8:45 - 9:00	AULA	**Conference Opening**
9:00 - 10:00	AULA	**Keynote Lecture 1: Graph Neural Networks Use Graphs When They Shouldn't** (Maya Bechler-Speicher, Meta)
10:00 - 10:30		**Coffee Break with** OLX
10:30 - 12:00	CW 1	**Session 1: Natural Language Processing** 10:30 - 11:00 → Babak Ehteshami Bejnordi 11:00 - 11:30 → Roman Grebennikov 11:30 - 12:00 → Grzegorz Bilewski PEARSON
	CW 2	**Session 2: AI for Science** 10:30 - 11:00 → Adam Kosiorek 11:00 - 11:30 → Jakub Adamczyk 11:30 - 12:00 → Piotr Ludynia
	CW 3	**Session 3: Computer Vision** 10:30 - 11:00 → Chih-Chen Kao AMD 11:00 - 11:30 → Krzysztof Krawiec 11:30 - 12:00 → Paweł Ekk-Cierniakowski & Kajetan Wencierski
12:15 - 13:15	AULA	**Keynote Lecture 2: LLMs for Code** (Baptiste Roziere, Mistral AI)
13:15 - 14:15		**Lunch Break**
14:15 - 15:30	AULA	**Panel Discussion: Research in Academia vs. Industry** (Michał Nowicki, Baptiste Rozière, Stanisław Jastrzębski, Faustyna Krawiec)
15:30 - 17:30		**Poster Session & Coffee Break**
16:30 - 17:30	AULA	**Student Session** (Dawid Siera, Michał Stefanik, Jakub Drzymała & Mateusz Konat, Kacper Cybiński, Łukasz Sztukiewicz, Kacper Wachnik)
19:00 - 1:00		**After Party** Blue Note Jazz Club (Imperial Castle, Kościuszki 79, Poznań)

Agenda

GHOST DAY
Applied Machine Learning Conference

DAY 2: 10-05-2025

9:00 - 10:30	AULA	**Keynote Lecture 3: AI Hardware and Real-world AI** (Andrew Fitzgibbon, Graphcore)
10:30 - 11:00		**Coffee Break**
11:00 - 11:30	AULA	**Accurate Structure Prediction of Biomolecular Interactions with AlphaFold 3** (Augustin Zidek, Google Deepmind)
11:40 - 12:40	CW 1	**Session 4: Optimization Techniques** 11:40 - 12:10 → Gergely Neu 12:10 - 12:40 → Patryk Rygiel
	CW 2	**Session 5: AI for Science** 11:40 - 12:10 → Fatima Sanchez-Cab 12:10 - 12:40 → Anastasiia Ponkratov
	CW 3	**Session 6: Computer Vision** 11:40 - 12:10 → Pavlo Melnyk 12:10 - 12:40 → Jędrzej Kopiszka OLX
12:40 - 13:40		**Lunch Break**
13:40 - 14:40	AULA	**Keynote Lecture 4: Oscillators and Traveling Waves in Machine Learning** (Max Welling, CuspAI)
14:40 - 15:10		Networking Break
15:10 - 16:40	CW 1	**Session 7: Natural Language Processing** 15:10 - 15:40 → Grigory Sapunov 15:40 - 16:10 → Martin Genzel 16:10 - 16:40 → Patrícia Schmidtová
	CW 2	**Session 8: Explainable AI** 15:10 - 15:40 → Stefan Haufe 15:40 - 16:10 → Jacek Karolczak 16:10 - 16:40 → Bartłomiej Sobieski
	CW 3	**Session 9: Applied Machine Learning** 15:10 - 15:40 → Michał Mikołajczak DATARABBIT 15:40 - 16:10 → Maciej Piernik 16:10 - 16:40 → Riccardo Belluzzo ALLEGRO
16:45 - 17:00	AULA	**Closing Remarks**

Building Plan

GHOST DAY
Applied Machine Learning Conference

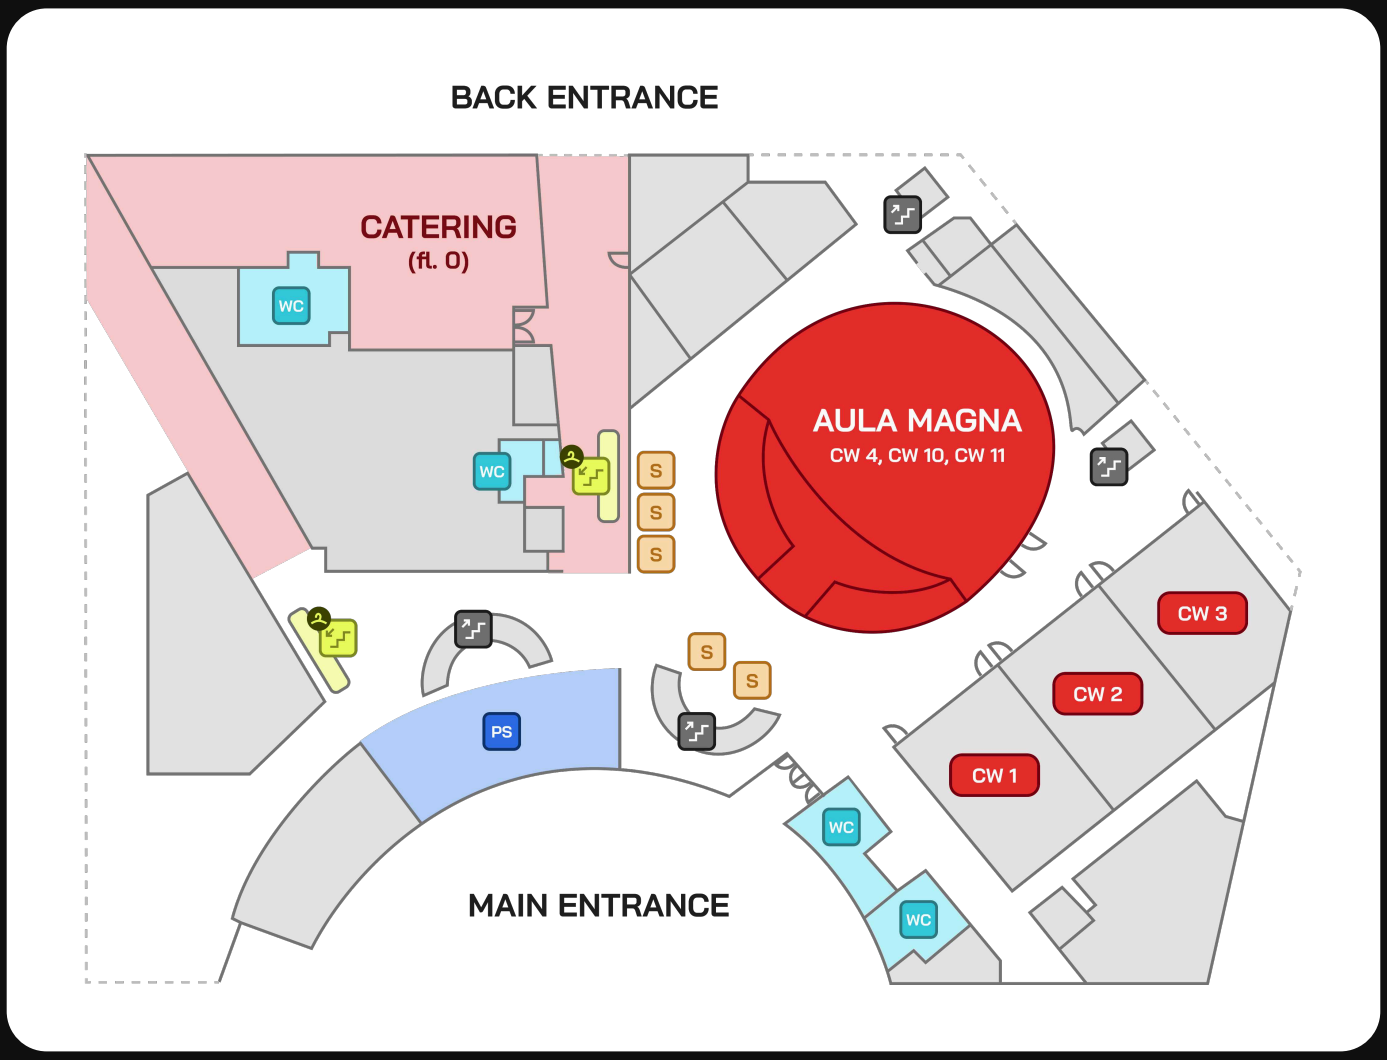

LEGEND (floor no.)

- Cloakroom & Registration (fl. -1)
- WC — Bathrooms
- PS — Poster Session (fl. 0)
- S — Sponsors Booths (fl. 0)
- CW 1 — Session Halls (fl. 0)

Maya Bechler-Speicher

Research Scientist | Meta

Keynote Lecture
Graph Neural Networks Use Graphs When They Shouldn’t

Biography

Maya is a Research Scientist at Meta and a last-year Ph.D. candidate in Computer Science at Tel Aviv University, where she is also a Lecturer, teaching "Machine Learning with Graphs." She is a domain expert in Graph Machine Learning, with research spanning both theoretical and applied aspects, including its integration with LLMs and Interpretability. Her work focuses on Graph Neural Networks , exploring their foundations and real-world applications.

Abstract

Learning over graph-structured data is central to a wide range of domains, including social networks, biological systems, and medicine. Graph Neural Networks (GNNs) have emerged as the leading methodology for such tasks, predicated on the assumption that the input graph encodes information relevant to the predictive objective.
Nevertheless, there exist scenarios in which improved generalization can be obtained by modifying the input graph—or disregarding it entirely. While GNNs are theoretically capable of attenuating the influence of uninformative edges, empirical evidence suggests they frequently overfit to the provided structure, even when it is misaligned with the target signal.
In this talk, I will present findings demonstrating that GNNs tend to overfit the input graph even in regimes where optimal performance is achieved by ignoring it, or by replacing it with special regular graphs devoid of any task-relevant information. I will further discuss the implications of this behavior from the perspective of implicit bias, and analyze how the alignment between the graph structure and the learning objective critically affects both the success and failure modes of GNNs. Finally, I will provide practical advice on how to identify and mitigate harmful structure in graph data.

Baptiste Rozière

Researcher and Code Generation Team Leader | Mistral AI

Keynote Lecture
LLMs for Code

Biography

Baptiste is leading the code generation team at Mistral AI. Previously, he was a research scientist in the codegen team at Meta AI in Paris working. He contributed to Llama and led Code Llama. During his PhD at Meta AI and Université Paris Dauphine, Baptiste conducted research on unsupervised translation of programming languages and model pre-training for code. His work was featured in dozens of news articles in more than ten languages. Prior to that, Baptiste worked as an applied scientist in the dynamic advertising team at Amazon.

Abstract

This talk explores how Large Language Models (LLMs) can revolutionize software development by enhancing programmer productivity. We will discuss key applications such as in-IDE code completion and code-assistant chatbots, demonstrating how LLMs can automate tasks and allow developers to focus on high-level design.

Andrew Fitzgibbon

Engineering Fellow | Graphcore

Keynote Lecture
AI Hardware and Real-world AI

Biography

Andrew Fitzgibbon is an Engineering Fellow at Graphcore, working on the future of computing hardware and programming for AI and numerical computing. He is best known for his work on computer vision: he was a core contributor to the Emmy-award-winning 3D camera tracker "boujou", having co-founded the company "2d3", with Andrew Zisserman, Julian Morris, and Nick Bolton; at Microsoft, he introduced massive synthetic training data for Kinect for Xbox 360; and was science lead on the real-time hand tracking in Microsoft's HoloLens. His research interests are broad, spanning computer vision, graphics, machine learning, neuroscience, and most recently programming languages. He has published numerous highly-cited papers, and received many awards for his work, including ten "best paper" prizes at various venues, the Silver medal of the Royal Academy of Engineering, and the BCS Roger Needham award. He is a fellow of the Royal Academy of Engineering, the Royal Society, the British Computer Society, and is very proud to be a Distinguished Fellow of the British Machine Vision Association. Before joining Graphcore in 2022, he spent 15 years at Microsoft, and before then, he was a Royal Society University Research Fellow at Oxford University, having previously studied at Edinburgh University, Heriot-Watt University, and University College, Cork.

Abstract

AI is fast becoming a significant consumer of the world's computational power, so it is crucial to use that power wisely and efficiently. Our approaches to doing so must span all levels of the research stack: from fundamental theoretical understanding of the loss surfaces and regularization properties of machine learning models, to efficient layout at the transistor level of floating-point multipliers and RAM. I will talk about projects, such as real-time computer vision on the Microsoft HoloLens HPU (about 3.5 GFLOPS), which required extreme efficiency in both objective and gradient computations, and how this relates to the training of massive AI models on Graphcore's IPU (about 350 TFLOPS). Key to this work is how we empower programmers to communicate effectively with such hardware, and how we design frameworks and languages to ensure we can put theory into practice. So this talk contains aspects of: mathematical optimization, automatic differentiation, programming languages, and silicon design. Despite this range of topics, the plan is for it to be accessible and useful to anyone who loves computers.

Max Welling

CTO | CuspAI & Professor | University of Amsterdam

Keynote Lecture
Oscillators and Traveling Waves in Machine Learning

Biography

Prof. Dr. Max Welling is a full professor and research chair in machine learning at the University of Amsterdam and a Merkin distinguished visiting professor at Caltech. He is co-founder and CTO of the startup CuspAI in Materials Design. He is a fellow at the Canadian Institute for Advanced Research (CIFAR) and the European Lab for Learning and Intelligent Systems (ELLIS) where he served on the founding board. His previous appointments include Partner and VP at Microsoft Research, VP at Qualcomm Technologies, professor at UC Irvine. He finished his PhD in theoretical high energy physics under supervision of Nobel laureate prof. Gerard 't Hooft. He then switched fields to focus on machine learning, first as a postdoc at Caltech under supervision of prof. Pietro Perona and then as postdoc under supervision of Nobel laureate prof. Geoffrey Hinton at UCL & U. Toronto. Max Welling has served as associate editor in chief of IEEE TPAMI from 2011-2015, he serves on the advisory board of the Neurips foundation since 2015, he is co-founder of the European Lab for Learning and Intelligence Systems (ELLIS) and served on its board until 2021, he has been program chair and general chair of Neurips in 2013 and 2014 respectively. He was also program chair of AISTATS in 2009 and ECCV in 2016 and general chair and co-founder of MIDL 2018. Max Welling is recipient of the ECCV Koenderink Prize in 2010, and the 10 year Test of Time awards at ICML in 2021 and ICLR in 2024.

Abstract

Traveling waves of neural activity are increasingly observed in biological brains. In this talk I will discuss a number of models that support traveling waves and explore how they can be used to build better ML models. In particular, the AKOrN model is able to synchronise collections of neurons and bind them together to represent larger, coarse grained, more abstract concepts. Under certain circumstances, the dynamic activities of the oscillatory neurons form traveling waves. We show that these models can help with memory tasks, can reason, are robust to adversarial attacks and can do unsupervised image segmentation. Their dynamical and biologically plausible nature makes them interesting candidates for a new generation of neural architectures.

Speakers

GHOST DAY
Applied Machine Learning Conference

Panel
Research in Academia vs. Industry

PANELIST

Michał Nowicki

Research Assistant Professor | Poznan University of Technology, Staff Engineer | AeroVect

PANELIST

Baptiste Rozière

Researcher and Code Generation Team Leader | Mistral AI

PANELIST

Stanisław Jastrzębski

CTO MoleculeOne

PANELIST

Faustyna Krawiec

Doctoral Student University of Cambridge

Introducing the panel discussion for
GHOST Day: Applied Machine Learning Conference 2025 🚀

Freedom vs. Impact. Why do some choose industry for better resources and career prospects, while others embrace academia despite its constraints? Let's delve into these and more questions from both academic and industry perspectives, exploring the challenges, motivations, and synergies behind choosing a research path in academia vs. industry.

Joining the discussion are some of the most respected voices in the field: Michal Valko, Andrew Fitzgibbon, Stanisław Jastrzębski, and Faustyna Krawiec, with Michał Nowicki moderating the conversation.

MODERATOR

Michał Wiliński

Carnegie Mellon Univerisity

Augustin Žídek

Research Engineer | Google DeepMind

Talk

Accurate structure prediction of biomolecular interactions with AlphaFold 3

Biography

Augustin Žídek works as a Research Engineer at Google DeepMind and has been a member of the AlphaFold team since 2017. He studied Computer Science at the University of Cambridge. He enjoys working at the boundary of research and engineering, hiking, playing musical instruments and fixing things.

Abstract

The introduction of AlphaFold 2 has spurred a revolution in modelling the structure of proteins and their interactions, enabling a huge range of applications in protein modelling and design. In this talk, I will describe our new AlphaFold 3 model that is capable of predicting the joint structure of complexes including proteins, nucleic acids, small molecules, ions and modified residues. I will also discuss the development of AlphaFold 1, 2, and 3, as well as the impact that AlphaFold has had so far.

Babak Ehteshami Bejnordi

Research Scientist | Qualcomm AI Research

Session 1: Natural Language Processing

Efficient Deployment of Large Language Models on Edge Devices

Biography

Babak Ehteshami Bejnordi is a Research Scientist at Qualcomm AI Research in the Netherlands, leading a research group focusing on conditional computation for efficient deep learning. His primary research focus lies in the realm of efficient Deep Learning for Large Language Models (LLMs) and Computer Vision. His recent research works have been in the areas of Efficient Autoregressive decoding in LLMs, Mixture of Experts, Multi-Task Learning, and Continual Learning. Babak obtained his Ph.D. in machine learning for breast cancer diagnosis from Radboud University in the Netherlands. During his Ph.D., he organized the CAMELYON16 challenge on breast cancer metastases detection which demonstrated one of the first medical diagnostic tasks in which AI algorithms outperform expert pathologists. Before joining Qualcomm he was a visiting researcher at Harvard University, BeckLab, and a member of the Broad Institute of MIT and Harvard. He has been the organizer of the Qualcomm Innovation Fellowship Program in Europe since 2019.

Abstract

Deploying large language models (LLMs) on edge devices has emerged as a compelling solution to achieve reduced costs, lower latency, enhanced privacy, and personalized user experiences. However, on-device deployment introduces unique challenges due to the resource constraints of edge devices, which differ significantly from high-throughput inference environments. Key obstacles include sequential token generation with batch size one and the inability of these devices to accommodate the massive size of LLMs within their limited DRAM. Additionally, directly loading model weights from flash storage can severely impact latency.

In this talk, we delve into the intricacies of on-device LLM inference during both prompt encoding and token generation stages. We then explore innovative approaches for developing efficient on-device language models. Specifically, we introduce a novel caching strategy that enables state-of-the-art Mixture of Expert models, that exceed the available DRAM on mobile devices, to be deployed effectively. Experimental results demonstrate up to 2X speedups on mobile hardware.

Furthermore, we present a hybrid methodology that combines language models of varying sizes to optimize autoregressive decoding while preserving high predictive performance. By conditioning a small language model (SLM) on representations generated by an LLM, we achieve substantial token generation speedups of up to 4X with minimal performance trade-offs (1–2% degradation) for tasks like translation and summarization.

Speakers

GHOST DAY
Applied Machine Learning Conference

Roman Grebennikov

Principal Engineer | Delivery Hero

Session 1: Natural Language Processing
LLMs for Machine Translation are here - but not quite yet

Biography

A principal ML engineer and an ex startup CTO working on modern search and recommendations problems. A pragmatic fan of open-source software, functional programming, LLMs and performance engineering.

Abstract

• How are traditional machine translation models trained and evaluated? A look at Google Translate, NLLB, and the evolution from n-gram to neural-based evaluation metrics.
• Wait, is seq2seq dead in 2025? How well do open-source LLMs like Llama and Qwen perform out of the box on MT tasks? And what about GPT-4o?
• Sure, modern LLMs aren't trained specifically for MT -- but what if we try? Enter X-ALMA and friends. Big and noisy vs. small and clean datasets for MT fine-tuning, and the curse of multilinguality.
• A few DIY experiments with Gemma 2/3 fine-tuning -- and how we accidentally beat NLLB after just 10 minutes of training.

SPONSOR SPEAKER - PEARSON

Grzegorz Bilewski

Software Engineering Manager | Pearson

Session 1: Natural Language Processing
LLM's in Education: Smart Lesson Generator case study

Biography

Grzegorz is an experienced Software Engineering Manager with over a decade of expertise in building and leading cross-functional teams, driving digital product development, and scaling technical organisations. At Pearson, he oversees a portfolio of global EdTech products, combining strategic leadership with hands-on technical direction.

Currently, Grzegorz leads the development of Smart Lesson Generator—an AI-powered teaching assistant designed to revolutionise language instruction. Built on AWS and integrated with large language models, the tool generates curriculum-aligned, level-appropriate classroom activities in seconds. It supports both Pearson and third-party content, ensuring flexibility for educators while aligning every activity with the Global Scale of English (GSE) for measurable learning outcomes. This initiative is a cornerstone of Pearson's broader AI strategy, focused on personalised education at scale, and already serves thousands of teachers worldwide.

Abstract

In this presentation, we will explore how Large Language Models (LLMs) are transforming education through the Smart Lesson Generator (SLG) case study. We will discuss how Pearson uses courseware data and the Global Scale of English within LLMs to create personalized learning experiences. Additionally, we will highlight new services driven by business needs and examine the challenges faced during SLG's development and implementation, and how we overcame them to ensure its effectiveness and reliability.

Adam Kosiorek

Senior Research Scientist | Google Deepmind

Session 2: AI for Science
Learning the Language of Life: AI in Genomics & Cell Biology

Biography

Adam R. Kosiorek is a Senior Research Scientist at Google DeepMind in London. His research interests include unsupervised representation learning, particularly learning object-centric representations from images and videos without supervision, and generative modelling. He is also interested in inference methods for probabilistic generative models. Previously, Adam interned at Google Brain in Toronto with Geoff Hinton, where he developed Stacked Capsule Autoencoders. He received his Ph.D. from the University of Oxford, MSc from the Technical University of Munich and BSc from the Warsaw University of Technology. In his free time, Adam likes reading books and training calisthenics.

Abstract

Large–language models can speed up paperwork and even suggest a diagnosis, yet they stop short when no test exists or no cure is known. Those breakthroughs lie deeper—in the tangled code of our genomes and in the shifting states of living cells. This talk traces the journey from messy biological measurements to the algorithms that turn them into insight and, ultimately, new therapies. We'll unpack real case studies that show how ML advances translate to spotting disease–causing DNA changes, predict how a drug will steer a cell's fate, and even design new genetic "switches" from scratch. Along the way, you'll see the key technical hurdles, the skills that matter most, and why progress here opens doors that language models alone cannot. You'll leave with a clear map of the questions worth tackling and a realistic view of how AI in genomics and cell biology could reshape medicine in the next decade.

Jakub Adamczyk

PhD candidate; AGH, Data Science Engineer | Placewise

Session 2: AI for Science
ML in agrochemistry and bee pesticide toxicity prediction

Biography

He is a PhD candidate in Computer Science at AGH University of Krakow and a member of the Graph ML and Chemoinformatics Group at the Faculty of Computer Science. His research focuses on fair evaluation, graph representation learning, graph classification, chemoinformatics, and molecular property prediction. He is also interested in time series, natural language processing (NLP), and MLOps, and teaches these subjects at AGH. He works at Placewise as a Data Science Engineer, where he addresses various machine learning challenges in tabular data, computer vision, and NLP, along with their end-to-end MLOps implementations. Outside of his professional work, he trains in Historical European Martial Arts (HEMA), specializing in messer and longsword, and enjoys reading and tabletop role-playing games.

Abstract

ML is widely applied in medicinal chemistry and pharmaceutical industry. Chemoinformatics and molecular ML have been used for decades for safer, faster drug design. However, the important area of agrochemistry has been relatively neglected. It has incredible potential with both predictive and generative models, e.g. to predict pesticide toxicity, or generate new, safer agrochemicals. New regulations and focus on ecological agriculture necessitate the need for fast development of alternative pesticides, and ML can help.

In this talk, I will describe how and why we can apply ML in agrochemistry. In particular, I will present ApisTox, a novel dataset about pesticide bee toxicity, how we can construct such dataset from publicly available data sources, and what challenges we meet. Lastly, results for predictive models will be presented, including e.g. molecular fingerprints, graph kernels, and GNNs. We will also discuss potential future work in this interesting area.

Piotr Ludynia

Researcher | AGH University of Krakow

Session 2: AI for Science

Surpassing GNNs and Transformers with Simple Feature Extraction for Peptide Function Prediction

Biography

Piotr is a master's student in Machine Learning and Data Science at AGH University of Kraków and a member of AGH ML and Chemoinformatics Group. He is one of the main maintainers of scikit-fingerprints, a molecular fingerprint Python library for efficient molecular vectorization. In addition to his academic work, he is currently employed at Intel Technology Poland, working on deep learning research and neural network optimization.

Abstract

Peptides, small proteins, are gaining attention for their therapeutic potential in combating cancer, viruses, and antibiotic-resistant bacteria. Antimicrobial peptides, in particular, offer a promising alternative to traditional antibiotics in addressing the growing resistance crisis. Accurately predicting peptide properties is crucial for drug discovery, and recent research has focused on deep learning approaches like graph neural networks and protein language models. However, it turns out that a much simpler graph feature extraction can achieve even better results with a fraction of the computational cost. We show that molecular fingerprints - domain specific molecular graph vectorization algorithms - combined with tree-based ensemble models, achieve state-of-the-art performance on multiple benchmarks and datasets. This method is not only more accurate but also much faster and scalable, opening new possibilities for graph machine learning in chemoinformatics and peptide-based drug design.

SPONSOR SPEAKER - AMD

Chih-Chen Kao

Member of Technical Staff | AMD

Session 3: Computer Vision

Advancing Ray Tracing with Neural Networks: Neural Intersection Function

Biography

Chih-Chen Kao is a member of the technical staff and a software engineer at AMD in Munich, Germany. Prior to joining AMD, he worked on simulation and perception systems for autonomous driving research. He earned his Ph.D. in Computer Science from National Taiwan University, Taipei, Taiwan, in 2017. His research interests span real-time ray tracing, physically-based rendering, heterogeneous computing, and the application of deep learning in computer graphics. In 2011, Dr. Kao was invited to participate in the OPTIMI (Online Predictive Tools for Intervention in Mental Illness) program, a research initiative funded by the European Union's 7th Framework Programme for Personal Health Systems, hosted at ETH Zurich in Switzerland.

Abstract

Bounding Volume Hierarchies (BVHs) are widely used to accelerate ray-object intersection queries in Monte Carlo ray tracing, particularly for evaluating visibility. However, traditional BVH traversal exhibits highly irregular memory access patterns and branch divergence, limiting its efficiency on massively parallel GPUs. Recent advances in neural representations suggest an opportunity to replace or augment BVH structures to address these inefficiencies. In this work, we propose a unified framework that integrates a Neural Intersection Function into a conventional BVH-based ray tracing pipeline. Our approach utilizes a multilayer perceptron (MLP) with dense matrix multiplications and predictable memory access patterns, allowing us to bypass the most irregular and time-consuming parts of BVH traversal during secondary ray queries. Furthermore, we introduce Locally-Subdivided Neural Intersection Functions (LSNIFs), a refined method that replaces bottom-level BVHs with neural networks trained to predict visibility, hit-point locations, and material indices directly. Key innovations include a sparse hash grid encoding scheme combined with geometry voxelization, a scene-agnostic training data collection pipeline, and a custom loss function to ensure high fidelity in hit-point prediction. LSNIFs are trained offline per object, enabling flexible and efficient support for arbitrary ray types and viewpoints at inference time. By leveraging neural approximations to replace or enhance critical components of the rendering pipeline, our work paves the way toward more efficient, neural-augmented ray tracing systems.

Krzysztof Krawiec

Full Professor | Poznan University of Technology

Session 3: Computer Vision

Neurosymbolic Autoencoders for Image Interpretation: From Physics-based ML to Program Synthesis

Biography

Krzysztof Krawiec is a Professor of Computer Science at the Poznan University of Technology in Poland, where he currently serves as the head of the Neurosymbolic Systems Group. His primary research areas include program synthesis, neurosymbolic systems, evolutionary computation, and medical imaging. He has authored over 180 publications on these topics and has received the Fulbright Senior Advanced Research Award, two ACM SIGEVO Impact Awards, and was a visiting professor at the University of California and Massachusetts Institute of Technology. He also served as the general chair of GECCO'21, the largest scientific event in the field of evolutionary computation and as an advisor at the Confederation of Laboratories for Artificial Intelligence in Europe. Krzysztof is also a co-founder of the Center for Artificial Intelligence and Machine learning, part of the Horizon 2020 Foundations of Trustworthy AI project, and an associate editor of Genetic Programming and Evolvable Machines and ACM Transactions on Evolutionary Learning and Optimization. In addition to his academic contributions, Professor Krawiec serves as the Chief AI Officer at Optopol Technology and CTO of Hylomorph Solutions Ltd.

Abstract

In order to make reliable predictions and robust inferences, image interpretation models need to achieve as complete an understanding of scene content as possible. This goal cannot be achieved in supervised learning, which focuses on narrow, task-specific goals. In this talk, I will argue for neurosymbolic autoencoders, a class of unsupervised architectures that enable holistic scene modeling by forcing the formation of transparent latent representations expressed in tangible terms of objects, their appearance, and their spatial arrangement. I will also illustrate how the learned representations can help solve downstream tasks in medical imaging and remote sensing.

Paweł Ekk-Cierniakowski & Kajetan Wencierski

Senior Manager AI & ML; Data Scientist | SoftwareOne

Session 3: Computer Vision
Machine learning in video processing

Biography

For over 10 years Paweł has been professionally involved in data analytics as a data scientist, team leader and project manager. He has participated in many projects in the area of advanced data analytics, such as monitoring production lines, opinion analysis, fraud detection and price forecasting. His main experience and interests are in the pharmaceutical and healthcare industries, but he has been involved in projects in various areas such as finance, media, energy and agriculture. Currently, he is responsible for designing and implementing data solutions, mainly in the field of machine learning and artificial intelligence. He shares his knowledge as a data science trainer, lecturer and speaker at conferences. Co-author of scientific papers, mainly in the field of medicine and statistics, published e.g. in journals from the Master Journal List.
Kajetan Wencierski works as a Data Scientist at SoftwareOne, holding a Master's degree in Computer Science from Poznan University of Technology. While studying there, he was also a member of GHOST. He is deeply passionate about the diverse applications of artificial intelligence, with a particular focus on the intersection of machine learning and audio. Kajetan especially enjoys pushing the boundaries of what is achievable by applying state-of-the-art AI models to solve real-world, everyday problems, finding the process deeply rewarding.

Abstract

During our lecture, we will delve into the comprehensive 2-year journey of our project for the media house, which began with the transcription of videos and the detection of background sounds. This initiative aims to enhance content accessibility for individuals with disabilities. Our presentation will cover the entire process, from the initial Proof-of-Concept stage, through the development of a Minimum Viable Product (MVP), and finally to the deployment in a production environment.
We will discuss various strategies to improve the quality of transcriptions and increase the accuracy of background sound detection. Additionally, we will highlight the integration of advanced features based on Azure services. These features include the creation of concise video summaries, automatic translations, and dubbing. Our goal is to showcase how these enhancements contribute to making content more accessible and user-friendly.

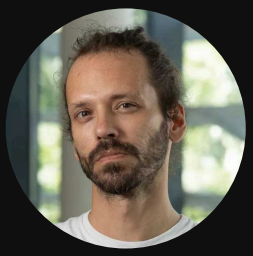

Gergely Neu

Professor | Pompeu Fabra University, Barcelona

Session 4: Optimization Techniques
Optimal transport distances for Markov chains

Biography

Gergely Neu is a research assistant professor at the Pompeu Fabra University, Barcelona, Spain. He has previously worked with the SequeL team of INRIA Lille, France and the RLAI group at the University of Alberta, Edmonton, Canada. He obtained his PhD degree in 2013 from the Budapest University of Technology and Economics, where his advisors were András György, Csaba Szepesvári and László Györfi. His main research interests are in machine learning theory, with a strong focus on sequential decision making problems. Dr. Neu was the recipient of a Google Faculty Research award in 2018, the Bosch Young AI Researcher Award in 2019, and an ERC Starting Grant in 2020.

Abstract

How can one define similarity metrics between stochastic processes? Understanding this question can help us design better representations for dynamical systems, study distances between structured objects, formally verify complex programs, and so on. In the past, the dominant framework for studying this question has been that of bisimulation metrics, a concept coming from theoretical computer science. My recent work has been exploring an alternative perspective based on the theory of optimal transport, which has led to surprising results, including a proof of the fact that bisimulation metrics are, in fact, optimal transport distances. This realization allowed us to import tools from optimal transport and develop computationally efficient methods for computing distances between Markov chains via the reduction of the problem to a finite-dimensional linear program. In this talk, I will introduce this framework and the most recent algorithmic developments, as well as discuss the potential for representation learning in more detail.

Patryk Rygiel

PhD candidate | University of Twente

Session 4: Optimization Techniques

Geometric deep learning in blood flow modelling

Biography

Patryk is a PhD candidate in the EU Horizon VASCUL-AID project at the University of Twente and a guest researcher at Amsterdam UMC. He is a former AI research scientist at med-tech company Hemolens Diagnostics and the best MSc graduate in AI at the Wrocław University of Science and Technology. His current research focuses on applications of geometric deep learning and implicit neural representations in medical imaging and blood flow modelling in the assessment of cardiovascular diseases. He is a co-author of international patents, a winner of the ICML Topological Deep Learning Challenge 2024 and a recipient of the MICCAI 2023 STAR Award. Next to his work, he is actively involved in AI community by being a member of the ML in PL association and a general chair of the Medical Imaging Symposium for PhDs and Postdocs (MISP^2) 2025.

Abstract

Computational fluid dynamics (CFD) is a commonly used method of in-silico blood flow modelling in the assessment of cardiovascular diseases (CVDs). CFD, a numerical approach to solving the Navier-Stokes equations that govern fluid dynamics, can provide hemodynamic markers that have been found to correlate with the development and progression of various CVDs. However accurate, the CFD simulations are known to be very computationally demanding in terms of both time and resources required. To remedy that, in recent years, geometric deep learning methods, operating directly on 3D shapes, have been proposed as compelling surrogates of CFD, providing accurate estimates in just a few seconds.

In this talk, we will show how utilizing geometric algebra transformers allows for the creation of accurate, generalizable and rapid CFD surrogates. We will showcase how their performance can be further enhanced by being coupled with physics-informed neural networks as well as deep vectorized operators, that can effectively map between function spaces. Additionally, we will cover the problem of training data acquisition in the real-life scenarios by delving into active learning methods that allow for creating data-efficient and robust surrogates.

Fatima Sanchez-Cabo

IP Computational Systems Biomedicine Lab | CNIC

Session 5: AI for Healthcare
AI enhances cardiovascular research

Biography

Dr. Fátima Sánchez Cabo graduated in Mathematics from the Complutense University of Madrid in 2000. From there she moved to the University of Manchester where she obtained a grant from the BBSRC to develop her doctoral work on statistical analysis and mathematical data modeling of microarrays. In 2005 she joined the Institute of Genomics and Bioinformatics of the Polytechnic University of Graz, Austria, where she developed her work first as a postdoctoral researcher and later as an associate professor. Since 2008 she has been working at the CNIC, leading the Bioinformatics Unit since 2017 and the Computational Systems Biomedicine Lab since 2025. Dr. Sánchez-Cabo has published more than 100 articles in peer-reviewed journals and is especially interested in the use of AI algorithms to enhance the understanding of biological systems in the context of cardiovascular disease and aging. Likewise, she has a strong commitment to the training of researchers in areas related to bioinformatics. Since 2021 is associate professor of the Autonoma University in Madrid. She is also the vicepresident of the Spanish Society for Computational Biology and Bioinformatics and member of the Advisory Board of ELIXIR.

Abstract

AI is transforming research information changing our life and the way we work in science, particularly in cardiovascular research. In my talk, I will be presenting several AI-based methods that might help improving the diagnostic, treatment and study of cardiovascular research. In particular, I will present EN-PESA, a Machine Learning algorithm that predicts the extent of subclinical atherosclerosis, improving the current definition of CV risk based on traditional scores. Also, I will be introducing CARMINA (CARdiovascular research MINer Assistant) a RAG implemented and locally running at CNIC, to ease the mining of bibliographical information and knowledge about CV research.

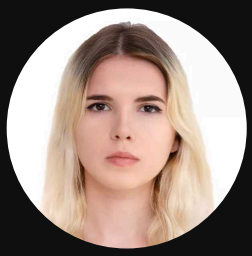

Anastasiia Ponkratova

Student | Polish-Japanese Academy of Information Technology

Session 5: AI for Healthcare

Automatic diagnosis of systemic and ophthalmic diseases using deep learning techniques in analyzing fundus images

Biography

Anastasiia Ponkratova is a recent graduate with a Bachelor's degree in Computer Science, possessing a strong foundation in programming, data analysis, and artificial intelligence. Passionate about machine and deep learning, Anastasiia is focused on applying AI in medicine, including image analysis, natural language processing, and predictive modeling. She actively expands her knowledge through online courses, conferences, and participation in open-source projects. She is particularly interested in how AI can enhance diagnostics, personalize treatment, and improve the quality of healthcare. Anastasia already has experience collaborating with professionals and medical research institutions and aims to further develop these collaborations to contribute to innovations at the intersection of technology and healthcare.

Abstract

I invite you to join me in exploring the cutting-edge applications of deep learning in medical diagnostics! The presentation will delve into the innovative use of the RETFound model, a state-of-the-art AI framework, for the automated diagnosis of systemic and ophthalmological diseases through color retinal image analysis. Using a unique private dataset of retinal images provided by a leading medical university, we demonstrate how this interdisciplinary approach enhances the accuracy of detecting conditions like hypertension, diabetes, glaucoma, and diabetic retinopathy. Discover the model's implementation, its adaptation for real-world clinical use, and the promising results achieved. Whether you are a researcher, clinician, or AI enthusiast, this talk will provide valuable insights into the future of AI-driven medical diagnostics and its potential to revolutionize healthcare. Don't miss this opportunity to learn about groundbreaking advancements that could transform patient care! 🙂

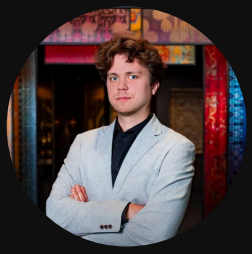

Pavlo Melnyk

Data Engineer | Skandinaviska Enskilda Banken (SEB)

Session 6: Business Data Science

The Backbone of AI: How Data Engineering Powers Machine Learning

Biography

Pavlo is an experienced Data Engineer with over 5 years of experience building scalable data solutions on AWS and Google Cloud Platform (GCP). With a strong passion for big data and data pipelining, he specializes in leveraging Python and Scala to design and optimize data workflows. Pavlo spent almost 2 years working within one of the largest Scandinavian banks, where he contributed to cutting-edge data infrastructure projects. A graduate of the Big Data analysis major at Warsaw School of Economics and of the applied mathematics at University of Warsaw, Pavlo combines technical expertise with a deep understanding of business and analytics. Outside of work, he is an avid learner with interests in Spanish and Mandarin, a passionate photographer, and an enthusiastic mountain hiker. He also enjoys exploring philosophy and applied neuroscience, blending his technical and creative sides to drive innovation in data engineering.

Abstract

In the world of machine learning and artificial intelligence, models are only as good as the data they are built on. While much attention is given to algorithms, model tuning and searching of the hyperparameters golden mean, the critical role of data engineering is often overlooked. This talk will dive deep into how data engineering forms the backbone of successful ML systems, enabling scalable, reliable, and efficient pipelines that transform raw data into actionable insights.

As a result of this talk, attendees will understand why data engineering is not just a supporting function but a critical enabler of machine learning success. Whether you're a data scientist, ML engineer, or data engineer, this session will provide actionable insights into building robust data systems that drive impactful ML solutions.

Speakers

GHOST DAY
Applied Machine Learning Conference

SPONSOR SPEAKER - OLX GROUP

Jędrzej Kopiszka

Junior Data Scientist | OLX Group

Session 6: Business Data Science
Tales of moderation. How ML helps us detect fraud at OLX

Biography

Jędrzej works as a Junior Data Scientist at OLX in the Moderation team, focusing on fraud detection through data science. He holds a Bachelor's degree in Artificial Intelligence from Poznań University of Technology and is currently pursuing a Master's in Artificial Intelligence at Johannes Kepler University in Linz, Austria. A former GHOST member and past GHOST Day coordinator, Jędrzej is a true believer in the mission of this conference's organizers. In his free time, he is an avid climber, cyclist, skier, and an enthusiast of urbanism.

Abstract

This talk will guide you through the moderation process at OLX – Poland's largest classifieds platform... and an app known to all of you! Specifically, I will focus on fraud detection in OLX chat, whether it involves images or text. I will showcase how our solutions have evolved over the past few years, the techniques we use, what we are working on, and why tackling fraud often feels like a cat-and-mouse game with fraudsters. Additionally, I will discuss the use of classic machine learning, our latest advancements in online learning, and content classification powered by our in-house generative AI platform.

Grigory Sapunov

CTO & Co-Founder | Intento

Session 7: Natural Language Processing
LLMs and Multilinguality

Biography

Grigory is the CTO and co-founder of Intento, a company dedicated to advancing machine learning and artificial intelligence technologies. Before founding Intento, he gained extensive experience in both industry and academia, working at Yandex and the Higher School of Economics. With a career spanning over 25 years in software engineering, Grigory has dedicated nearly 20 years to data analysis, artificial intelligence, and machine learning, building expertise in these cutting-edge fields. Since 2011, Grigory has been deeply involved in deep learning, contributing to the development and application of these transformative technologies. His work bridges theoretical advancements and practical implementations, making him a leader in the AI community. Grigory is also a Google Developer Expert in Machine Learning, a recognition of his deep technical knowledge and active contributions to the developer ecosystem.

Abstract

With over 7,000 languages spoken globally, the field of Natural Language Processing (NLP) has historically been dominated by English, with some attention to other high-resource languages like Chinese, Spanish, and French. This focus has led to significant advancements in language technologies for these languages, but it has also resulted in a disparity in NLP resources and tools available for less-represented languages. This talk will explore the current state of multilinguality in Large Language Models (LLMs), examining the challenges and opportunities in expanding language coverage beyond the most popular and digitized languages. We will discuss the limitations of existing models in handling diverse linguistic structures and the efforts being made to create more inclusive language technologies. By highlighting recent research and developments, this presentation aims to shed light on the path toward truly multilingual AI systems that can serve a broader spectrum of the world's linguistic communities.

Martin Genzel

Applied Machine Learning Researcher | Merantix Momentum

Session 7: Natural Language Processing
Can Compressing Foundation Models Be as Easy as Image Compression?

Biography

Martin Genzel is a Staff Machine Learning Researcher at Merantix Momentum, developing deep learning solutions for tabular and time-series data in real-world applications. As an applied mathematician by training, he got his Ph.D. from TU Berlin working on compressed sensing and high-dimensional signal processing. In his postdocs at Utrecht University and the Helmholtz Centre Berlin, he explored deep learning techniques for ill-posed inverse problems and computational imaging.

Abstract

The widespread adoption of Foundation Models, especially LLMs, is often hindered by their substantial size and computational demands, especially in resource-limited settings. While post-training compression offers a promising avenue to mitigate these challenges, the process can feel like a ""black box"" for the user, requiring significant expertise and trial-and-error to find the right balance between model size and performance.
This talk introduces Any Compression via Iterative Pruning (ACIP), a novel algorithmic approach designed with the user in mind. ACIP allows for intuitive and direct control over the compression-performance trade-off, akin to compressing an image. It leverages a single gradient descent run of iterative pruning to establish a global parameter ranking, from which models of any target size can be immediately materialized.
ACIP demonstrates strong predictive performance on downstream tasks without costly fine-tuning. Across various open-weight LLMs, it achieves state-of-the-art compression results compared to existing factorization-based methods. Moreover, it seamlessly complements common quantization techniques for even greater compression.

Patrícia Schmidtová

PhD student | Charles University

Session 7: Natural Language Processing
Evaluating LLM-generated text at scale

Biography

Patricia Schmidtova is a Ph.D. student in Natural Language Processing (NLP) at Charles University, focusing on semantic accuracy in natural language generation and its evaluation methodologies. Her work has earned her Best Paper Awards at EACL and INLG. With seven years of industry experience, she has specialized in implementing NLP components for task automation in the banking sector. In her talk, Patricia will share her experience designing an automatic evaluation protocol to assess the reliability of millions of summaries generated by large language models (LLMs).

Abstract

Large language models (LLMs) allow us to invent new ways of transforming data and text. In many cases, they can do it without a larger collection of training data – a couple of examples are often all that is required. However, this normally means there is also no curated evaluation dataset available. So what can we do to understand and track the quality of the generated texts?Benchmarks are popular on social networks, but they rarely simulate specific real-world applications. Metrics that are most frequently used in research require a larger set of human-written reference solutions that are frequently not available and cost a substantial amount to collect and curate. In my talk, I will share my experience with evaluating LLM-generated summaries of hotel descriptions. By documenting my journey with this task, I aspire to give you a recipe for building an evaluation pipeline that covers your needs. Spoiler alert, there is unfortunately no silver bullet.

Stefan Haufe

Head of the UNIML group & Professor | Technische Universität Berlin

Session 8: Explainable AI

How can explainable AI provide quality control for ML?

Biography

Stefan Haufe is a joint Associate Professor of Uncertainty, Inverse Modeling, and Machine Learning at Technische Universität Berlin and Physikalisch-Technische Bundesanstalt Berlin. He also serves as a Group Leader at Charité - Universitätsmedizin Berlin, where his research focuses on developing advanced computational methods. His group specializes in signal processing, inverse modeling, and machine learning techniques for analyzing neuroimaging and other medical data. This work aims to address challenges in understanding complex physiological processes and improving medical diagnostics and treatment strategies. Additionally, he has a keen interest in model interpretation and explainable artificial intelligence, striving to make machine learning applications more transparent and interpretable for practical use in medical and scientific contexts. Through collaborations with interdisciplinary teams, Stefan Haufe contributes to advancing both theoretical and applied aspects of computational neuroscience and biomedical data analysis.

Abstract

There is a desire for humans to ""understand"" decisions of machine learning (ML) models. Explainable AI (XAI) seemingly addresses this need by providing information beyond a model's output. XAI methods have been proposed to help identifying and correcting failure modes in models and data, to facilitate scientific discovery by identifying ""important"" features, and to perform counterfactual reasoning.
However, rarely are XAI methods designed to fulfill any of such purposes, which raises questions about the generalizability of reported empirical findings and the conditions under which certain interpretations of XAI outputs (e.g., feature attribution maps) are really valid.
Here, we present theoretical analyses of minimal problem settings as well as empirical results on ground-truth benchmarks showing that feature attribution methods cannot in general distinguish between features associated with the predicted variable or not. We discuss how this can lead to incorrect conclusions about models and data, and to suboptimal downstream decisions such as the rejection of an optimal model.
We present preliminary efforts to formalize distinct ""explanation problems"", emphasizing that useful explanations typically require understanding not only of the model function but also the distribution and causal structure of the data it was trained on.
Finally, we outline novel approaches addressing specific explanation problems.

Jacek Karolczak

PhD student | Poznan University of Technology

Session 8: Explainable AI
Explainable AI: Moving from Numbers to Meaningful Insights

Biography

Jacek Karolczak is a Ph.D. student at Poznan University of Technology. His research focuses on improving the interpretability of machine learning models, particularly in dynamic environments where data continuously evolves. He believes the world of explainable AI (xAI) is shifting beyond feature-importance explanations, embracing high-level concept-based interpretations that better align with the language and reasoning of end users.

Abstract

The increasing complexity of machine learning models has heightened the demand for their explainability. While most presentations at conferences like GHOST Day focus on feature importance, particularly Shapley values, these explanations are often criticized as incomprehensible, even to machine learning experts. In contrast, the XAI 2.0 manifesto advocates for concept-based explanations, such as ia prototypes - representative instances. Besides introducing the problem, the author will explore existing approaches and discuss recent contributions to concept-based explainability, including their own work on prototype-based concept drift detection, which ensures the approach's intrinsic interpretability. Additionally, the potential of prototypes to improve professional ML applications will be discussed.

Bartłomiej Sobieski

AI Researcher | University of Warsaw, MI2.ai

Session 8: Explainable AI
Rethinking Visual Counterfactual Explanations Through Region Constraint

Biography

He is a PhD student in Computer Science at the University of Warsaw, focusing on the development of novel algorithms that combine generative modeling with explainability in computer vision. His primary research interests involve exploring complex mathematical fields, such as the theory of differential equations and differential geometry, which serve as inspiration for innovative research ideas. His academic record includes publications in venues such as the International Conference on Artificial Intelligence in Medicine, Artificial Intelligence in Medicine, International Conference on Agents and Artificial Intelligence, Computer Vision and Pattern Recognition Conference Workshops, and the European Conference on Computer Vision.

Abstract

Visual counterfactual explanations (VCEs) have recently gained immense popularity as a tool for clarifying the decision-making process of image classifiers. This trend is largely motivated by what these explanations promise to deliver -- indicate semantically meaningful factors that change the classifier's decision. However, we argue that current state-of-the-art approaches lack a crucial component -- the region constraint -- whose absence prevents from drawing explicit conclusions, and may even lead to faulty reasoning due to phenomenons like confirmation bias. To address the issue of previous methods, which modify images in a very entangled and widely dispersed manner, we propose region-constrained VCEs (RVCEs), which assume that only a predefined image region can be modified to influence the model's prediction. To effectively sample from this subclass of VCEs, we propose Region-Constrained Counterfactual Schrödinger Bridges (RCSB), an adaptation of Schrödinger Bridges to the problem of conditional inpainting, where the conditioning signal originates from the classifier of interest. In addition to setting a new state-of-the-art by a large margin, we extend RCSB to allow for exact counterfactual reasoning, where the predefined region contains only the factor of interest, and incorporating the user to actively interact with the RVCE by predefining the regions manually.

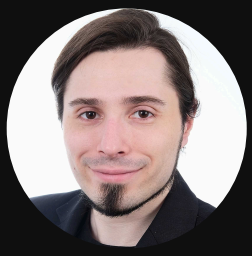

SPONSOR SPEAKER - DATARABBIT.AI

Michał Mikołajczak

Lead Solutions Architect | CEO | datarabbit.ai

Session 9: Applied Machine Learning
From Labs and PoCs to Production: Hard–Earned Lessons from 20+ Real–World AI Projects

Biography

Michal Mikolajczak is the founder and lead solutions architect at datarabbit.ai, a custom data/AI/cloud solutions provider that helps organizations leverage their data and gain a competitive advantage by designing, building, and delivering artificial intelligence and data-driven solutions for their businesses. Due to his work on various projects in different industries, he has a wide range of diverse ML experience, with an emphasis on its productization. He has particular expertise in the medical industry - having worked in medical imaging for several years, including as CTO of a startup that was successfully acquired (80 million USD transaction value) by a NASDAQ-listed company. Privately, a big fan of BI / visualization systems of all kinds that enable storytelling with data and Pratchett's work.

Abstract

There are many ideas for AI projects that make it and make a significant impact on business. There are many more, that seem great, but ultimately won't. Most of them, in fact. Why do so many promising ideas and pilots never make it to production? In this session, based on distilled knowledge from over 20 real-world projects (of various sizes and for various industries), we'll unpack some universal patterns behind why AI and data science projects stall, and what it takes to get them over the finish line. Through actual use cases and stories – some triumphant, others cautionary – we'll explore the key reasons AI initiatives fail to scale: misaligned expectations and business goals, processes and communication problems, data and infrastructure pitfalls, tech issues themselves. Then, we'll shift focus to practical strategies that work: aligning with business goals, starting small and iterating fast, building reliable MLOps foundations, and planning for real-world usage from day one. Whether you're a hands-on data scientist or a business leader navigating AI strategy, you'll leave with actionable insights and a fresh perspective on how to turn your next PoC into a production-grade success. If you're serious about delivering real value with AI, this is the talk you don't want to miss.

Maciej Piernik

Assistant Professor | Poznan University of Technology

Session 9: Applied Machine Learning
Retail Intelligence: From Noisy Receipts to Accurate Purchase Predictions

Biography

Dr. Maciej Piernik is a computer scientist specializing in artificial intelligence applications in healthcare. With a PhD in Computer Science from Poznan University of Technology, he combines academic research as an Assistant Professor with practical industry experience. His expertise lies in machine learning algorithms for complex data processing, particularly in biomedical applications. As a data scientist and consultant, Dr. Piernik collaborates with various businesses to build and deploy AI/analytical models, as well as medical professionals researching chronic diseases including multiple sclerosis, systemic lupus erythematosus, and inflammatory bowel diseases. With over 20 scientific publications to his credit, Dr. Piernik's research interests span artificial intelligence, cellular and molecular biology, and consciousness. He has developed several open-source tools for medical data analysis, including tools for gene expression data processing in IBD patients.

Abstract

This talk presents a unified approach addressing two critical challenges in retail analytics: product categorization from noisy receipt data and next-basket recommendation. We demonstrate how improved product labeling directly enhances recommendation accuracy in real-world retail environments. Our system processes OCR-extracted receipt data through a robust classification pipeline that achieves 80% mean accuracy despite inconsistent training labels, outperforming traditional rule-based methods. We compare LLM-based categorization against more efficient embedding-based neural networks, finding that lightweight models can match performance while enabling local deployment. Building on these enhanced product labels, our recommendation system achieves a 90% AUC for predicting future purchases, significantly outperforming recency and frequency baselines. We'll discuss our methodology for handling hierarchical categories, capturing seasonality patterns, and maintaining scalability. This work provides practical insights for implementing end-to-end retail intelligence systems that operate under real-world constraints while delivering meaningful improvements in recommendation quality.

SPONSOR SPEAKER - ALLEGRO

Riccardo Belluzzo

Research Engineer | Allegro

Session 9: Applied Machine Learning

From Regret to Retry: Training LLMs for Self-correcting SQL Generation

Biography

Riccardo is a Senior Research Engineer at Allegro, working in the Language Intelligence division. Over the past few years, he's been involved in a wide range of projects, from training multilingual language models to developing representation learning for e-commerce. He's also recently led several successful projects, bringing LLM-based solutions into production by applying the latest NLP breakthroughs. Riccardo is also a dedicated member of the local ML community, sharing his expertise at meetups and conferences. When he's not researching, he enjoys playing electric guitar, cooking delicious Italian meals, and exploring the world

Abstract

Recent developments in the field of "reasoning LLMs" have demonstrated that language models trained to follow chain-of-thoughts (CoT) excel at solving problems requiring logical thinking, like math or symbolical reasoning. A critical question arises: can an LLM recognize when its chain of thoughts contains a logical error? In this talk, I will explore this fascinating topic, showcasing recent discoveries in LLM theory through a practical use case: training an open-source coding LLM to solve the text-to-SQL task, which involves translating natural language questions into SQL queries. Throughout the talk, we will discover that LLMs not only exhibit patterns of regret when generating erroneous SQL reasoning steps, but can also be trained to recognize and self-correct those errors

STUDENT SPEAKER

Dawid Siera

Student | Poznan University of Technology

Student Session
SageML: Accelerating Auto-ML through neural network-driven model selection

Biography

He is a student at Poznań University of Technology, currently in the 6th semester of a bachelor's degree in Artificial Intelligence. He is the leader of Ghost's Quantum Computing group and is also interested in Explainable AI (XAI) and Automated Machine Learning (AutoML). He is currently working at PSNC as a quantum software developer, focusing on the development of quantum machine learning and quantum optimization software in Python. Beyond his academic and professional activities, he has initiated and organized several open-source projects, such as Torchboard and Decorify, which have been successfully published on PyPI. He is passionate about developing innovative AI tools, pushing the boundaries of quantum computing, and making machine learning more transparent and accessible.

Abstract

The selection of optimal machine learning models for specific datasets is a critical yet time-consuming task in the development of data-driven applications. Traditional AutoML systems automate this process, but often require extensive computational resources and time to explore a vast search space of potential models. We introduce SageML, a Python-based framework designed to accelerate the model selection process by leveraging combined topological dataset analysis and pre-trained neural networks. SageML predicts the performance of various machine learning algorithms, selecting the best candidate for a given dataset. By utilizing a pre-trained neural network to estimate model accuracies, SageML significantly reduces the search time while maintaining high predictive performance. The preliminary results show promising reductions in computational cost compared to traditional AutoML approaches, allowing users with limited domain knowledge or computational resources to achieve high-quality models efficiently.

Speakers

GHOST DAY
Applied Machine Learning Conference

STUDENT SPEAKER

Michał Stefanik

Student | AGH University of Krakow

Student Session
Molecular property prediction with Graph-of-Graphs approach

Biography

Michał is a Computer Science student at AGH University of Krakow. His current interests cover applying machine learning techniques to data of biological origin. Software engineer by profession.

Abstract

Accurately predicting the properties of chemical molecules remains a critical challenge in computational chemistry and drug discovery. These properties encompass a wide range of characteristics, including physical attributes such as solubility and melting point, as well as more complex biological interactions like the inhibition of viral growth or the ability to traverse cellular membranes. The latter are particularly pertinent in the context of drug development.

Currently, molecular fingerprints and Graph Neural Networks (GNNs) are among the most widely used approaches in the industry. Each offers unique strengths and weaknesses in modeling molecular properties. In this presentation, I will introduce a ""graph-of-graphs"" methodology that integrates the strengths of both techniques into a novel framework. Specifically, this approach leverages molecular fingerprints for vectorization, followed by node classification within a molecular similarity graph. I will discuss the key advantages of this method, including its interpretability and training stability, and how it can enhance the predictive accuracy and efficiency of molecular property estimation.

STUDENT SPEAKER

Jakub Drzymała & Mateusz Konat

Students | Poznan University of Technology

Student Session

Teaching Machines to See Forests: ML-Powered Tree Classification for QGIS

Biography

Jakub is an DevOps and Platform Engineer, working with Kubernetes, different Cloud Providers, programming scripts and way too much yaml. Born in the 90's in Berlin, he help companies build products from zero to one, deploying and maintaining them in production. His particular interest lies in making painful processes more pleasant, by using Infrastructure as a Code, Terraform, Ansible, automation, pipelines, Kubernetes, scripting and cloud. He is a technocrat building new tools every day. He spend a lot of time learning new skills and actively help other people learn software development and become tech-savvy through a variety of help groups and through writing own code. Besides programming, He is passionate about fighting sports, from Bruce Lee or old style Greco-Roman to modern MMA. He really like hiking in the mountains, sailing and motorcycles.

Mateusz is a first-year Artificial Intelligence student at Poznań University of Technology with a strong interest in the intersection of AI and biology. He is an active member of the GHOST machine learning science club, where he contributes to a Building an ML Model for Tree Recognition project in collaboration with the IRIM Institute. He has also worked on predicting ankylosing spondylitis risk based on HLA-B27 clades in European populations. As a GHOST members, they are participating in the IRIM project, developing a tree classification model for the QGIS Deepness Plugin, which they would like to show you during their short Lightning Talk during this year GHOST Day.

Abstract

The GHOST student research club, in collaboration with the Institute of Robotics and Machine Intelligence (IRIM) at Poznan University of Technology, is at the forefront of utilizing machine learning for tree classification through the QGIS Deepness Plugin. In this presentation, we will share our journey of developing machine learning models that automatically identify tree species from aerial and satellite imagery, thus enhancing environmental analysis and urban planning efforts.

We will guide you through the essential phases of our project, including data collection, preprocessing, and model training using readily available resources, along with our current results. Additionally, we will discuss the challenges we encountered, such as managing a diverse range of tree species and optimizing model performance.

STUDENT SPEAKER

Kacper Cybiński

Student | University of Warsaw

Student Session

Speak so a physicist can understand you! TetrisCNN for detecting phase transitions and order parameters

Biography

He is a master's student at the Faculty of Physics at the University of Warsaw (UW), majoring in Quantum Physics and Chemistry within the framework of the Individual Research Studies program. He holds a bachelor's degree from the College of Inter-Faculty Individual Studies in Mathematics and Natural Sciences (MISMaP) at the University of Warsaw, with a specialization in physics. Alongside academic development, he is actively engaged in social initiatives, organizing leading scientific conferences in Poland in the fields of quantum physics, optics, and photonics. During his studies, he gained valuable research experience through scientific stays at the Institut de Ciències Fotòniques (ICFO, Spain) and The Flatiron Institute (New York, USA), where he pursued scientific interests at the intersection of artificial intelligence and quantum physics. These activities were supported by the Excellence Initiative – Research University (IDUB) program. He actively participates in international conferences, presenting his research through talks and scientific posters. He has been involved in projects funded by the National Science Center, the European Research Council (ERC), and the Ministry of Science and Higher Education, including initiatives such as “Best of the Best 4.0.” In recognition of his scientific achievements, he was awarded a scholarship from the Minister of Science and Higher Education in 2024. In his free time, he enjoys sailing, hiking, climbing, and traveling in search of exceptional cuisine.

Abstract

Recently, neural networks (NNs) have become a powerful tool for detecting quantum phases of matter. Unfortunately, NNs are black boxes and only identify phases without elucidating their properties. Novel physics benefits most from insights about phases, traditionally extracted in spin systems using spin correlators. Here, we combine two approaches and design TetrisCNN, a convolutional NN with parallel branches using different kernels that detects the phases of spin systems and expresses their essential descriptors, called order parameters, in a symbolic form based on spin correlators. We demonstrate this on the example of snapshots of the one-dimensional transverse-field Ising model taken in various bases. We show also that TetrisCNN can detect more complex order parameters using the example of two-dimensional Ising gauge theory. This work can lead to the integration of NNs with quantum simulators to study new exotic phases of matter.

Speakers

GHOST DAY
Applied Machine Learning Conference

STUDENT SPEAKER

Łukasz Sztukiewicz

Student | Poznan University of Technology

Student Session
DetoxAI - Python Package for Debiasing Neural Networks

Biography

Łukasz Sztukiewicz holds a Bachelor of Science degree in Artificial Intelligence from Poznan University of Technology. He participated in the prestigious Robotics Institute Summer Scholar Program at Carnegie Mellon University and currently works as a machine learning engineer at molecule.one.

Abstract

The lack of specialized debiasing methods and appropriate software frameworks leaves a critical gap in addressing fairness issues for machine learning systems. To address these challenges, we present DetoxAI, a Python-based software framework for post-hoc debiasing of neural networks in image classification tasks. Designed with deep learning in mind, DetoxAI integrates state-of-the-art interventions, evaluation metrics, and visualization tools into a unified, production-ready ecosystem. Our approach applies post-training adaptation, allowing users to mitigate bias while maintaining model performance. By focusing on high-level semantic representations, DetoxAI addresses the unique challenges posed by vision data, where proprietary attributes such as race or gender are not explicitly encoded. This toolkit provides a modular interface for bias mitigation, making it accessible and adaptable for real-world applications. Through experimental studies, we quantitatively demonstrate that DetoxAI reliably improves upon baseline vanilla models on a fairness-performance tradeoff. Furthermore, we qualitatively show with attribution maps that DetoxAI's methods can shift the model's focus away from protected attributes.

STUDENT SPEAKER

BEST STUDENT TALK AWARD

Kacper Wachnik

Student | Poznan University of Technology

Student Session

MisterCar - AI agents for every virtual environment

Biography

He is a computer science engineer and a master's student in artificial intelligence, passionate about coding, automation, and AI agents. He is dedicated to knowledge sharing through talks, articles, and video courses. His journey in AI began with DeepMind's AlphaZero chess system, reflecting his love for the game. His interests later evolved to include computer vision and autonomous vehicles, and eventually expanded to large language models and AI agents following the release of ChatGPT. With professional experience in automating recruitment processes at a startup, he is currently developing his own agentic framework for researchers and developers. When not immersed in AI, he enjoys action-adventure books, watching and playing chess, and listening to epic soundtrack music from movies and games.

Abstract

MisterCar is a new framework built on the Sense-Plan-Act architecture that enables AI agents to perceive, understand, and interact with any virtual environment. Developed initially during autonomous driving research, MisterCar has evolved into a versatile toolkit with equal strength in two key domains: Video Game integration for researchers to transform games into experimental AI environments, and Computer Use automation that brings advanced capabilities to business workflows, GUI navigation, and process automation.

MisterCar provides unprecedented contextual richness through its specialized sensors — capturing screenshots, recording audio, monitoring input devices, interfacing with external hardware, and even inspecting memory and files. Combined with comprehensive deep learning support, you can train custom agents rather than relying solely on existing models.

DetoxAI - Python Package for Debiasing Neural Networks

Ignacy Stępka, Łukasz Sztukiewicz, Michał Wiliński, Jerzy Stefanowski

Abstract

The lack of specialized debiasing methods and appropriate software frameworks leaves a critical gap in addressing fairness issues for machine learning systems. To address these challenges, we present DetoxAI, a Python-based software framework for post-hoc debiasing of neural networks in image classification tasks. Designed with deep learning in mind, DetoxAI integrates state-of-the-art interventions, evaluation metrics, and visualization tools into a unified, production-ready ecosystem. Our approach applies post-training adaptation, allowing users to mitigate bias while maintaining model performance. By focusing on high-level semantic representations, DetoxAI addresses the unique challenges posed by vision data, where proprietary attributes such as race or gender are not explicitly encoded. This toolkit provides a modular interface for bias mitigation, making it accessible and adaptable for real-world applications. Through experimental studies, we quantitatively demonstrate that DetoxAI reliably improves upon baseline vanilla models on a fairness-performance tradeoff. Furthermore, we qualitatively show with attribution maps that DetoxAI's methods can shift the model's focus away from protected attributes.

How synthetic data can help with imbalance image classification

Karol Cyganik

Abstract

Imbalanced datasets pose a significant challenge in image classification, often leading to biased models that underperform on underrepresented classes. But what if we could generate realistic synthetic data to balance the scales? This research explores the potential of synthetic image generation using diffusion models to enhance training data diversity.

I investigate multiple strategies for incorporating synthetic images, including active learning, active generation, and similarity-based selection, to determine the most effective way to integrate artificial samples. By iteratively refining data selection and augmentation, I aim to improve model performance while maintaining realism and diversity in synthetic images.

Vocabulary-constrained story generation for language learners

Wiktor Kamzela, Mateusz Lango

Abstract

One of the most popular ways to learn a foreign language is to use a Spaced Repetition System (SRS) such as Anki, SuperMemo or HackChinese. SRS is essentially an application for learning new vocabulary by reviewing a deck of flashcards. Every day, SRS selects a set of flashcards containing words that have been learned in the past, but that the algorithm predicts the user may forget in the near future. In this way, the system optimises the user's ability to recall all the vocabulary they have learned, while minimising the number of repetitions. Although this strategy proved to be effective in teaching the user to recall new words, it is of limited value for real vocabulary acquisition. The words are reviewed out of context and learners often find it difficult to use them naturally in a sentence. Also, systematically reviewing flashcards with new vocabulary can become boring.

We propose a new generation of SRS systems that replace traditional flashcard review with the reading of a short story generated for a particular user with specific vocabulary estimated by a classical SRS. The reviewed words would be used in real sentences, providing the learner with the necessary context and examples of word usage. On the other hand, the words that would be new to the user would appear several times in the generated story, often in a context that allows the reader to guess the meaning of the word. This would make the learning process more enjoyable and effective, but also presents a new technological challenge to automatically generate such stories in different languages. In this work, we exploit the ability of large language models to generate high quality, fluent texts and propose several prompting strategies for SRS story generation. According to NLG metrics, our approach achieves good performance for different words, learner levels and languages.

Artificial Intelligence in Data Management: The Effectiveness Threshold of LLMs in Real-World Applications

inż. Mateusz Bugajak, dr inż. Anna Gorawska

Abstract

Large Language Models (LLMs) are increasingly used for automating various tasks, including database and data warehouse querying. However, their effectiveness in handling real-world, unstructured, and complex data environments remains an open question. Small and medium-sized enterprises (SMEs) often struggle with disorganized data structures, which hinder efficiency and scalability. Using real-world NGO workflows, this study evaluates the ability of state-of-the-art LLMs (e.g., GPT-4, Deepseek R1) to generate clear and optimized database queries from unstructured data representations. By analyzing gaps between automated outputs and a manually designed benchmark, the study aims to identify critical thresholds where AI intervention becomes viable versus scenarios requiring human expertise.

While final results remain pending, preliminary findings indicate that high data entropy significantly reduces LLM applicability, often requiring human intervention. Our findings aim to delineate the threshold where AI-driven query generation remains effective and where human expertise becomes indispensable. This research provides insights into the viability of LLMs in practical database management scenarios and their potential to streamline data querying for SMEs. Central to this work is the question: Where does AI's applicability begin in real-world data management?

BEST SCIENTIFIC QUALITY POSTER AWARD

Improving performance of distributed learning through density estimation

Kacper Trębacz, Jack Henry Good, Artur Dubrawski

Abstract

In domains like healthcare and military, training data is often distributed among independent clients who cannot share data due to privacy concerns or limited bandwidth. This scenario is typically addressed by federated learning, where a central server coordinates communication of models between clients. However, when a central server is absent, clients must train models in a distributed, peer-to-peer manner. Ongoing work by the Auton Lab is developing a novel approach based on Function Space Regularization (FSR) to train models in distributed fashion with low communication overhead. In their work, the authors have shown that for certain high dimensional data sets, the performance is far lower than that of a central model. It is partially due to the fact that function space regularization enforces agreement on the whole domain, which overly penalizes models for disagreeing on out of distribution data. To address this, issue, we propose an extension to the Distributed AI framework that leverages density estimates to appropriately weight the FSR by these estimates. We show this approach in a scenario using Decision Trees coupled with Gaussian Mixture Models for density estimation. Our benchmark results on popular machine learning datasets as well as synthetically created datasets demonstrate improvements of up to 25% in F1 score compared to unweighted FSR.

AUDIENCE CHOICE POSTER AWARD

BEST SCIENTIFIC QUALITY POSTER AWARD

Exploring the Stability Gap in Continual Learning: The Role of the Classification Head

Wojciech Łapacz, Daniel Marczak, Filip Szatkowski, Tomasz Trzciński

Abstract

Continual learning (CL) has emerged as a critical area in machine learning, enabling neural networks to learn from evolving data distributions while mitigating catastrophic forgetting. However, recent research has identified the stability gap - a phenomenon where models initially lose performance on previously learned tasks before partially recovering during training. Such learning dynamics are contradictory to the intuitive understanding of stability in continual learning where one would expect the performance to degrade gradually instead of rapidly decreasing and then partially recovering later. To better understand and alleviate the stability gap, we investigate it at different levels of the neural network architecture, particularly focusing on the role of the classification head. We introduce the nearest-mean classifier (NMC) as a tool to attribute the influence of the backbone and the classification head on the stability gap. Our experiments demonstrate that NMC not only improves final performance, but also significantly enhances training stability across various continual learning benchmarks, including CIFAR100, ImageNet100, CUB-200, and FGVC Aircrafts. Moreover, we find that NMC also reduces task-recency bias. Our analysis provides new insights into the stability gap and suggests that the primary contributor to this phenomenon is the linear head, rather than the insufficient representation learning.

Continual learning for satellite data analysis

Joanna Wiekiera

Abstract

Continual learning plays a crucial role in satellite data analysis, enabling models to adapt to new information without forgetting previously learned knowledge. This work reviews state-of-the-art approaches to continual learning applied to both telemetry and satellite imagery, focusing on strategies for mitigating catastrophic forgetting in neural networks. Key methods, including regularization-based, replay-based, and architectural techniques, are analyzed in the context of class-incremental learning. Additionally, preliminary insights from an experimental study are presented, evaluating the effectiveness of these approaches on benchmark datasets. The challenges of real-world deployment, such as data distribution shifts and computational constraints, are also discussed. This review provides a foundation for further research on enhancing the adaptability and robustness of AI models for space applications.

Disentangling Visual Priors: Unsupervised Learning of Scene Interpretations with Compositional Autoencoder

Antoni Nowinowski, Krzysztof Krawiec

Abstract

Contemporary deep learning architectures lack principled means for capturing and handling fundamental visual concepts, like objects, shapes, geometric transforms, and other higher-level structures. We propose a neurosymbolic architecture that uses a domain-specific language to capture selected priors of image formation, including object shape, appearance, categorization, and geometric transforms. We express template programs in that language and learn their parameterization with features extracted from the scene by a convolutional neural network. When executed, the parameterized program produces geometric primitives which are rendered and assessed for correspondence with the scene content and trained via auto-association with gradient. We confront our approach with a baseline method on a synthetic benchmark and demonstrate its capacity to disentangle selected aspects of the image formation process, learn from small data, correct inference in the presence of noise, and out-of-sample generalization.

Efficient Deep Learning Architecture for Image Watermarking with Adaptive Feature Enhancement

Paweł Duszejko, Mateusz Gabor, Zuzanna Gawrysiak, Tomasz Hawro, Tymoteusz Lindner

Abstract

In this work, we present a novel deep learning architecture for image watermarking that integrates key modules from recent advancements in the field. Our architecture incorporates a skip-layer excitation module, which enhances feature representation while maintaining computational efficiency. Additionally, we utilise conditional batch normalisation, allowing for adaptive adjustments to the watermarking process based on message variations. Furthermore, our approach leverages techniques for message spreading and mask generation, to effectively embed watermarks while minimising perceptual distortion. The resulting architecture is lightweight and optimised for speed, demonstrating robust performance against noise interference. We evaluate our model on standard benchmark datasets, achieving competitive results that underscore its effectiveness and resilience. This work contributes to the ongoing development of efficient watermarking solutions that balance fidelity, robustness, and computational efficiency, paving the way for practical applications in digital media protection.

Cleaning AFM Images with U-Net: Deep Learning for Nanoscale Artifact Removal

Jakub Klimkowski

Abstract

The Atomic Force Microscopy (AFM) image can be affected by artifacts like blobs, tip distortions, and scanning noise, all of which can obscure nanoscale details and necessitate extensive human intervention to rectify. Existing methods for denoising images either oversmooth the image or are too rigid in the parameter settings.

This research addresses these challenges by generating algorithmic synthetic data for training a U-Net-based model. Randomized noise injection simulates artifact by multi-scale height-map transformations and stochastic geometric deformations to create paired datasets of pristine/corrupted surfaces. The synthetic data approach eliminates tedious steps of manual annotation while simulating a range of defect profiles found in nanomaterial samples.

Through experiments performed on real AFM data, the model selects artifacts while maintaining details of the surface structure below the sub-nanometer. Integrating synthetic training data with real-time inference, the methodology simplifies the process of artifact elimination without needing extensive manual pre-processing.

Enhancing Crowd Management with Drones, Computer Vision and Deep Learning

Bartosz Ptak, Marek Kraft

Abstract

Effective crowd management at large events, protests, and marathons requires accurate population estimation, yet traditional counting methods often introduce significant errors. Leveraging recent advances in deep learning and aerial drone imaging, we present an end-to-end solution designed to improve the accuracy of crowd size estimation, achiving trajectory counting error of 15%. Our approach not only refines head detection and density estimation but also provides an explicit error analysis, ensuring transparent and reliable results. This system can support emergency services, city authorities, and event organizers in making informed decisions, ultimately enhancing public safety and resource allocation.

Explaining Predictive Models Through Invariants

Bartłomiej Sobieski, Łukasz Niedźwiedzki, Emilia Kaczmarczyk, MIriam Lipniacka, Antoni Janowski

Abstract

Our project introduces a novel method for explaining image classifiers using diffusion models by generating invariant images—synthetically created examples that produce identical activations in a specified network layer. Unlike traditional counterfactual explanations that modify inputs to alter classifier decisions, our approach focuses on generating images that elicit similar responses, offering insights into the most critical features influencing model predictions. We employ diffusion models, specifically DDPM and DDIM, to iteratively generate these images while guiding the process with classifier activations. Our experiments with ResNet-50 trained on ImageNet reveal that targeting different neurons or class probabilities leads to distinct generated outputs. Future work will refine guidance strategies that explore all the solutions in data manifold that give the expected neuron activation. We also plan to optimize loss functions, and explore advanced generative models to improve our method.

Application of vision transformers to protein-ligand affinity prediction

Jakub Poziemski, Paweł Siedlecki

Abstract

The transformer architecture has revolutionized many areas related to AI. It was originally adopted for natural language processing (NLP), but in recent years there has been rapid development of transformer architectures for computer vision (CV) data, the so-called Vision Transformers (ViT). ViT is achieving spectacular results in many CV areas, displacing architectures based on convolutional neural networks. (CNN). In this paper, we present a successful application of ViT to the problem of protein-ligand affinity prediction based on 3D crystallographic complexes. Despite the relatively small dataset and the very complex nature of the problem, ViT achieves results comparable to the best methods used for this problem. The paper also includes extensive model diagnostics that provide information on important aspects of the input data and its representation.

Fair Comparison Benchmark for Molecular Property Prediction

Mateusz Praski

Abstract

Navigating the rapidly expanding landscape of machine learning models for molecular property prediction requires robust evaluation frameworks. Our research presents a comprehensive benchmark of multiple pre-trained machine learning architectures across various property prediction tasks. We address both theoretical challenges in model evaluation, including metric selection, benchmarking methodology, and dataset selection, as well as practical challenges, such as automation processes. The goal of this benchmark is to evaluate novel approaches to molecular chemistry using pre-trained models, compared with previously popular approaches, such as molecular fingerprints. Join us to discuss how pre-trained architectures are utilized in molecular chemistry and explore our benchmarking process for transparent and reproducible comparisons across a diverse spectrum of models.

GHOST DAY
Applied Machine Learning Conference

Experimental Evaluation of Molecular Filters

Adam Staniszewski, Jakub Adamczyk

Abstract

We study the effectiveness of filters commonly used in chemoinformatics for molecular screening. Filters are rule-based tools, based on physicochemical descriptors or sustructural patterns, used to preprocess molecular datasets and remove unwanted compounds, e.g. potentially toxic. Despite their widespread use, there is a lack of quantitative evaluation of those methods, and even experimental procedures for such evaluation. To remediate this, we assess over thirty common filters across a number of molecular datasets.

We compare physicochemical and substructural filters on various tasks. Proposed validation procedure compares the efficiency in removing unwanted compounds and ability to keep bioactive compounds, where there is typically a tradeoff between the two. Results indicate redundancy of some of the filters and rather weak evidence for superiory of more complex designs. We conclude that while molecular filters offer a computationally inexpensive and interpretable tool for molecular dataset curation, there is still room for improvements in this space.

Implementation Experience of BER Degradation Prediction Using Random Forest in GANA KP for 5G & Beyond Optical Networks

Cristian Zumelzu Scheel

Abstract

Predicting Bit Error Rate (BER) degradation is crucial for ensuring reliable optical transmission in 5G and beyond networks. In this work, we present our implementation experience of BER degradation prediction using Random Forest models within the ETSI GANA Knowledge Plane (KP) for autonomic network management. Leveraging real-world optical network data, we applied a sliding window approach to forecast BER degradation, proactively mitigating service disruptions. Our model addresses data imbalance through augmentation techniques, improving prediction accuracy across diverse network conditions. Experimental results demonstrate the model's effectiveness in early-stage BER degradation detection, enabling proactive maintenance and optimizing Quality of Transmission (QoT). Attendees will gain insights into the model's design, key performance metrics, and potential for real-time deployment. If you're interested in machine learning for network reliability, join us to explore how predictive analytics can transform optical network management.

Quantum-Enhanced Signal Processing Meets Deep Learning and Reinforcement Learning: A Hybrid Framework for Robust Beam Position Monitoring

Quentin Bruant, Abdelaziz Guelfane, Ismail Cherkaoui, Yuxian ZUO, DALENA Barbara, Francesca Bugiotti

Abstract

High-energy particle accelerators rely on Beam Position Monitors (BPMs) for precise beam alignment and operational safety, yet their performance is often compromised by high-frequency noise, mechanical instabilities, and sporadic sensor failures. In this work, we propose a novel hybrid framework that integrates quantum-inspired signal processing, deep learning, and reinforcement learning to robustly denoise BPM signals, detect latent anomalies, and adaptively control beam operations.

Our approach begins with a multi-scale decomposition using Complete Ensemble Empirical Mode Decomposition with Adaptive Noise (CEEMDAN), followed by a quantum-inspired spectral filtering module that selectively enhances beam-relevant features while suppressing both stochastic and structured noise. The refined signals are then analyzed by a hybrid anomaly detection pipeline that synergizes self-supervised contrastive learning with transformer-based architectures. Complementing this, a reinforcement learning component is incorporated to dynamically adjust filtering parameters and guide diagnostic decision-making. By interacting with the accelerator control system, the RL agent learns optimal policies that maximize system stability and enhance early fault detection through reward-driven adaptation.

Evaluations on real-world BPM datasets from operating particle accelerators demonstrate that our method outperforms state-of-the-art baselines, achieving over a 30% improvement in failure classification precision while offering a robust adaptive response to emerging anomalies. This work establishes a new benchmark for predictive maintenance in accelerator environments, illustrating the synergistic potential of quantum-enhanced signal processing, advanced deep learning, and reinforcement learning in developing resilient, adaptive industrial AI solutions.

Risk Indicators and Their Prediction in Continuous Auditing: A Machine Learning Approach

Izabela Kartowicz-Stolarska, Jerzy Krawczuk, Marek Tabędzki, Piotr Hońko, Marek Krętowski

Abstract

Continuous auditing is a modern control framework that leverages real-time monitoring and evaluation of organizational processes. By integrating advanced IT systems and machine learning techniques, it enables proactive anomaly detection, allowing for rapid corrective actions. In banking projects, continuous auditing is critical due to evolving regulatory requirements and stringent security standards. Risk indicators play a central role in threat monitoring, and their predictive analysis is essential for early risk mitigation.

This poster presents a machine learning–based risk assessment tool developed by the Faculty of Computer Science of Białystok University of Technology in collaboration with a leading Polish bank as part of a project co-funded by the National Centre for Research and Development (NCBiR). We propose a hybrid methodology combining statistical models and artificial intelligence (AI) to forecast key risk indicators (KRIs). The study outlines the system's architecture, evaluates the predictive performance of machine learning algorithms, and presents empirical results demonstrating the model's effectiveness in dynamic auditing environments.

Machine Learning meets blockchain: addressing key challenges in anomaly detection

Julia Merta, Michał Gorawski, Dawid Janduda, Michał Kliemt, Kamil Kaczmarek, Wojciech Trojanowski, Anna Gorawska

Abstract

In recent years, blockchain technology has gained significant popularity. As usage of the Blockchain increases, its vulnerability to attacks also grows. Developing machine learning models for the identification of anomalous transactions poses a significant challenge. This study emphasizes the critical issues involved in this process. One of the main challenges in this process is the limited number of publicly available, properly labeled data sets. A new dataset was created for analysis using raw Bitcoin transactions from the Binance exchange, utilizing publicly available information to identify the exchange's addresses. Additional characteristics of each address were calculated using the transaction data, such as the number of incoming and outgoing transactions, the sum of incoming and outgoing transaction values, and the difference between the funds sent and received by the address. Statistical analysis revealed significant differences in the volume of transactions and the amounts transferred among network participants. Leveraging the available information and calculated characteristics, an attempt was made to classify addresses into four categories: hot wallets, cold wallets, deposits, and ordinary network users. In addition, the operation of mixing addresses and the challenges of detecting them were discussed, highlighting the difficulties in identifying this type of anomaly.

SageML: Accelerating Auto-ML Through Neural Network-Driven Model Selection

Anatol Kaczmarek, Dawid Siera, Adam Tomys

Abstract

The selection of optimal machine learning models for specific datasets is a critical yet time-consuming task in the development of data-driven applications. Traditional AutoML systems automate this process, but often require extensive computational resources and time to explore a vast search space of potential models. We introduce SageML, a Python-based framework designed to accelerate the model selection process by leveraging combined topological dataset analysis and pre-trained neural networks. SageML predicts the performance of various machine learning algorithms, selecting the best candidate for a given dataset. By utilizing a pre-trained neural network to estimate model accuracies, SageML significantly reduces the search time while maintaining high predictive performance. The preliminary results show promising reductions in computational cost compared to traditional AutoML approaches, allowing users with limited domain knowledge or computational resources to achieve high-quality models efficiently.

MisterCar - AI agents for every virtual environment

Kacper Wachnik

Abstract

MisterCar is a new framework built on the Sense-Plan-Act architecture that enables AI agents to perceive, understand, and interact with any virtual environment. Developed initially during autonomous driving research, MisterCar has evolved into a versatile toolkit with equal strength in two key domains: Video Game integration for researchers to transform games into experimental AI environments, and Computer Use automation that brings advanced capabilities to business workflows, GUI navigation, and process automation.

MisterCar provides unprecedented contextual richness through its specialized sensors — capturing screenshots, recording audio, monitoring input devices, interfacing with external hardware, and even inspecting memory and files. Combined with comprehensive deep learning support, you can train custom agents rather than relying solely on existing models.