

NORA's annual conference 2025

Complete list of abstracts for 12 oral presentations and 28 poster presentations

Oral presentations

1. Keynote: AI Ethics Across Borders: Values, Data, and Cultural Context

Giada Pistilli

Principal Ethicist at Hugging Face

Ethics in AI isn't just a checkbox: it's about making deliberate choices that shape how technology impacts our lives and societies. As Large Language Models (LLMs) become integrated into our daily activities, we face important questions about their development and deployment — those conversational AI systems are rapidly transforming how we work and communicate. But whose values do they reflect, and who controls the data they're built on? Each language carries unique concepts, values, and worldviews that reflect distinct cultural value systems and ways of understanding reality. When AI training prioritizes dominant languages and perspectives, we risk amplifying existing inequalities and losing the rich diversity that multiple linguistic traditions offer to our collective understanding of reality. Drawing from recent research initiatives like CIVICS, which examines how AI systems respond to culturally sensitive topics, and data governance frameworks developed through the multilingual BigScience workshop, we'll explore how values and cultural contexts shape these technologies. By connecting ethical principles to practical implementation strategies, we can build AI systems that genuinely respect human diversity while delivering meaningful benefits across communities and contexts.

2. Contributed talk: How good is your LLM? Systematic evaluation of Norwegian skills and knowledge for Large Language Models

Lilja Øvrelid

Leader of the Language Technology Group, University of Oslo

In recent years, a number of Large Language Models (LLMs) have been trained specifically for Norwegian. The motivation behind these efforts have been to make available models that generate higher quality Norwegian text and that are better attuned to Norwegian language and values. A critical question moving forward is therefore how we can assess the quality of LLM-generated Norwegian text. In this talk, I will briefly outline the current LLM landscape of Norway and Scandinavia. I will then present NorEval: a new and comprehensive evaluation suite for large-scale standardized benchmarking of Norwegian generative language models. NorEval covers a broad spectrum of task categories targeting Norwegian language understanding and

generation, establishes human baselines, and focuses on both of the official written standards of the Norwegian language: Bokmål and Nynorsk. I will present this new resource and the results of benchmarking 19 open-source pretrained and instruction-tuned LLMs for Norwegian. While NorEval constitutes an important first step toward a more informed use of LLMs in Norway, I will conclude with thoughts on the missing pieces and important next steps for LLM research in Norway.

3. Contributed talk: Iris.ai - reflections on a 10-year journey of AI for Research

Anita Schjøll Abildgaard

CEO Iris.ai

Iris.ai is Norway's first AI startup, and we've been on the forefront of AI for Deep Knowledge and Research for almost a decade. In this talk, we will walk you through our journey - both from a startup, technology and research perspective. The ups and downs, the past and current challenges and the opportunities that lie ahead. The moments where we were about to give up, and the moments where everything fell into place. Of course we will dig into the research and development of our tech stack, presenting some things old and some brand new - like our newly published context sensitivity metric and our critical analysis of reliability in LLM evaluation metrics. And finally we will cover how we navigate our product offerings in a highly shifting AI market.

4. Contributed talk: ML for real-time Audio Enhancement

Erlend Aune

CTO at HANCE

In AI there is an increasing number of papers achieving state-of-the-art performance without considering the computational cost or how to achieve real-time inference. These models are - for the most part - unusable for real-time AI inference on the edge. In this talk I will present how HANCE works with smaller quasi-causal models, capable of running efficiently on mobile devices in real-time. I will touch upon the design choices we meet when making new models, and how we address training of these models. Throughout the presentation, I will give examples of models, how they work, and which application they are intended for, while connecting these applications to the design choices we face when making the models.

5. Contributed talk: Norwegian language models for Norwegian conditions – dream or reality?

Svein Arne Bygfeld

Head of the AI Lab, National Library of Norway

The race among tech giants to develop the best language models is intense. At the same time, we see that these same giants are adapting their services to the political frameworks within which they operate, and clear political biases can be observed.

Increasingly, people are using these services in more and more areas, and perhaps Norway as a nation is being influenced linguistically and culturally.

As a measure to ensure access to language models that better reflect the Norwegian language and Norwegian society, the Norwegian Parliament (Stortinget) has allocated 40 million Norwegian kroner in the 2025 national budget for the National Library of Norway (NB) to deliver language models for free use by all parts of society. The models will support variants of Nynorsk and Bokmål, as well as Sami languages. They will be built on and reflect the daily reality of Norwegian life. Computational resources will be provided by Sigma2.

The National Library's digital collection, combined with data from other sources, is tailor-made for training language models. However, the copyright restrictions that apply to the collection otherwise remain in place. Therefore, the government has asked the National Library to negotiate with rights holders for permission to use copyrighted textual content in the training of language models.

If successful in both areas, both the public sector and the business community will gain access to Norwegian language models that can be used in service development. The licenses for the models will provide good opportunities for widespread use, including for commercial product development.

6. Contributed talk: From Research and Development – to Society and Industry, in the Public Interest

Gunnar Bøe

Managing Director at Sigma2

To stay competitive globally, to safeguard the environment and resources, and to deliver the best services to customers and citizens, it is essential to adopt the latest advancements in AI. This requires access to vast computational resources and expertise. Equally important is ensuring a seamless and efficient transfer of knowledge and development between research, society, and industry.

AI computational resources demand continuous large-scale investments, deep expertise, and careful management of energy and environmental challenges. Additionally, computational capacity must be utilized as efficiently as possible, striking a balance between national control and governance and relying solely on commercial providers.

For over a decade, Sigma2 and NRIS have provided national high-performance computing (HPC) and data services, primarily for research. Moving forward, there will be a stronger focus on integrating these efforts with industry and public administration. In addition to existing AI resources, Sigma2 is making significant investments to meet growing demands and support this increasingly important collaboration.

In this presentation, we will provide insights into Norway's computational resources—both existing and upcoming—as well as the expertise and support services available, including the forthcoming LUMI AI Factory.

7. Contributed talk: Beyond Text: Alignment Challenges in Multimodal AI Systems

Michael Riegler

Head of AI at Simula

As AI systems evolve to process and generate content across multiple modalities - including text, images, audio, video, and physical interactions - the challenge of alignment becomes increasingly complex. This talk examines how the multimodal nature of modern AI systems introduces novel considerations for alignment research. We explore how different modalities might interact in ways that complicate traditional alignment approaches, from potential divergences in concept grounding across modalities to emergent behaviors that arise from modal interactions. Special attention will be given to three key challenges: ensuring consistent ethical behavior across modalities, developing effective oversight mechanisms for multimodal systems, and addressing potential failure modes unique to cross-modal learning. Drawing from recent research in multimodal AI and alignment theory, we propose a framework for thinking about alignment that explicitly accounts for the unique challenges and opportunities presented by multimodal systems.

8. Keynote: Aspects of XAI in neural multimodal learning (MML)

Robert Jenssen

Director of Visual Intelligence, UiT

The first part of this talk will outline some basic principles for neural multimodal learning (MML). Focus then shifts to aspects of eXplainable AI (XAI) in the context of MML. Leveraging the so-called CLIP model, the second part of the talk presents an approach to have vision networks “explain” how they function in terms of capturing concepts (such as colors, patterns, objects, etc) with particular focus on “transformer” networks. This is important for the general understanding of such AI systems and to shed light on pretraining vs finetuning. The third part of the talk presents ideas for developing self-explaining MML by learned masks for classification purposes. The aim is to identify the most informative features within each modality that influence predictions. This is potentially important for users, for reliability and trustworthiness.

9. Host Keynote: AI and the Unsettling Shift in Education: Are We Approaching a Tipping Point or a Transformational Change?

Irina Engeness

Professor at the Dept. of Languages, Literature and Culture, HiØ

Artificial Intelligence (AI) is affecting education, bringing both opportunities and challenges. As AI-driven technologies become more integrated into learning environments, fundamental questions arise: How will AI shape the roles of teachers and

students? Will it enhance learning, teaching, and assessment, or does it risk undermining traditional educational practices? Can AI foster creativity and critical thinking, or will it lead to intellectual stagnation and passive knowledge consumption?

This presentation will explore these critical issues, examining the benefits and dilemmas of AI in education. It will consider AI's role in assessment and feedback, its impact on student engagement, and the growing reliance on generative AI for content creation. Additionally, it will address concerns about the devaluation of human expertise, the risks of bias in AI-driven decision-making, and the broader implications for academic integrity and independent learning.

This session aims to provide a nuanced perspective on the future of AI in education: Are we approaching a tipping point of another wave of technological disruption, or will AI cause fundamental transformational changes in the educational landscape?

10. Contributed talk: Datafication and the Midas touch: ethics in AI-enabled healthcare and beyond

Anna Smajdor

Professor at the Practical Philosophy Group, UiO

The rapid deployment of AI in medicine has made datafication an increasingly integral feature of our health services. Patient data is systematically collected and analyzed; complex human experiences become mere numbers. Datafication transforms everything into digital artifacts, just as Midas turned all he touched to gold. Indeed, these digital artifacts are immensely valuable. The data we can yield is hungrily sought by those who develop AI systems. The ethical response to this has so far focussed on the need to bolster privacy, anonymity and confidentiality, stripping sensitive data of its connections to the individuals from whom it originates. However, this detachment makes it easier for data to be commodified and exploited in ways that may not necessarily align with the original donors' interests. As an alternative, I propose a new approach to privacy and confidentiality. Instead of obliterating the links between people and their data, the significance of these connections should be recognized and preserved. In turn, this opens the possibility of offering individuals greater involvement and control over the AI technologies developed from their data. Ultimately, anonymisation is not ethically neutral: it feeds into a false perception of data as unclaimed property up for grabs.

11. Keynote: Title to be determined.

Samuel Kaski

Leader of FCAI and ELLIS Institute Finland

12. Keynote: Embodied AI: Learning Through Curiosity and Play

Amy Loutfi

Program Director of WASP

In this keynote, Professor Amy Loutfi explores how embodied AI—intelligent systems that learn through physical interaction—can be enhanced by integrating human-inspired traits like curiosity and play. Drawing from her research at Örebro University and Linköping University, she discusses how these traits enable robots to adapt to real-world environments, improving perception, reasoning, and human-robot collaboration. As a Wallenberg Scholar, Loutfi's work advances AI's capacity for autonomous learning and interaction.

Poster presentations

1. ENFIELD: European Lighthouse to Manifest Trustworthy and Green AI.

Pankaj Pandey¹

1. Department of Information Security and Communication Technology,
Norwegian University of Science and Technology, Gjøvik, Norway

ENFIELD: European Lighthouse to Manifest Trustworthy and Green AI, is a large European Commission-funded research and innovation project where 30 European partners from 18 countries from academia, SMEs, and industry are participating. ENFIELD received over EUR 11 Million in funding to create a unique European Centre of Excellence that excels in fundamental research in the scientific pillars of Adaptive, Green, Human-Centric, and Trustworthy AI that are new, strategic and of paramount importance to successful AI development, deployment, and acceptance in Europe and will further advance the research within verticals of healthcare, energy, manufacturing and space industry to reinforce a competitive EU position in AI and create significant socio-economic impact for the benefit of European citizens and businesses by attracting the best talents, technologies and resources from world-class research and industry players in Europe and by carrying out top-level research activities in synchronisation with industry challenges to reinforce a competitive EU position in AI and create significant socio-economic impact for the benefit of European citizens and businesses.

ENFIELD, through its results, aims to positively contribute to several Sustainable Development Goals (SDG), namely, SDG 3 - Good Health and Well-being; SDG 4 - Quality Education; SDG 5 - Gender Equality; SDG 7-Affordable and Clean Energy; SDG 8 - Decent Work and Economic Growth; SDG 9 - Industry, Innovation and Infrastructure; SDG 11 - Sustainable Cities and Communities; SDG 12 - Responsible Consumption and Production; SDG 13 - Climate Action; and SDG 17 - Partnerships for the Goals.

ENFIELD, through its third-party funding schemes, funds Third-Party Exchange Researchers (TES) for up to 6 months of researcher mobility and Third-Party Innovation Scheme (TIS) for innovation projects up to EUR 60,000.

2. AI4IS - Towards AI-based forecasting of Antarctic Ice shelf calving

Jacob Hay¹, Hamze Issa¹, Daniele Fantin¹, David Parkes², Amber Leeson², Malcom McMillan², Kate Briggs², Jan Wuite³, Thomas Nagler³

1. Science and Technology AS, Oslo, Norway

2. Lancaster Environment Centre, Lancaster University, Lancaster, UK

3. ENVEO IT GmbH, Innsbruck, Austria

AI4IS is an ESA-funded project developing AI-based forecasting for Antarctic ice shelf calving—critical to understanding ice shelf instability, one of the major uncertainties in sea level rise projections. Calving events, driven by complex, non-linear forces, are difficult to predict but vital for improving future sea level scenarios.

AI4IS explores using a multivariate data cube combining diverse Earth Observation (EO) and modelled datasets (e.g., wind speed, ice velocity, strain rate, surface elevation, supraglacial lakes) to train a convolutional neural network for calving prediction. Alongside existing EO products like ice velocity and surface melt, the project also generates new datasets on strain rate and ice shelf damage to enhance monitoring of key ice shelf parameters.

To harmonize data of varying spatial resolutions, a Gaussian random field method is used to interpolate all datasets onto a shared monthly grid. This approach enables effective fusion of heterogeneous data, infilling of missing values, and propagation of uncertainty estimates, which are included as input features in the AI model.

The AI model, based on a modified U-Net architecture within the SNTML framework, is trained on manually labeled calving events from Qi et al. (2021). It identifies the time and location of potential calving events, marking a crucial step toward data-driven prediction of ice shelf collapse.

This work has significant scientific and societal implications, contributing to more accurate forecasting of ice shelf instability. With a long-term vision to integrate this framework into a digital twin of Antarctica, AI4IS also lays the groundwork for future “what if?” scenario testing and enhanced coupling with other Earth system models, such as ocean circulation models.

3. Development of a machine learning-based ice shelf coastline (ISCL) delineation in Antarctica

Hamze Issa¹, Daniele Fantin¹, Jacob Hay¹ and the AIS CCI Team

1. Science and Technology AS, Oslo, Norway

Recent advances have shown that deep learning can effectively delineate ice sheet front locations, leading to the development of Icelines, a Sentinel-1 radar-based service.

However, relying solely on radar data can be limiting in noisy or ambiguous scenes. Integrating additional sources, such as altimetry, could improve accuracy. CryoSat-2 (CS2) swath altimetry has previously been used for front-line delineation via edge-detection on gridded elevation data, but this approach suffers from low temporal resolution—requiring months of data to produce a DEM—and variable spatial resolution.

Currently, no solution combines the high temporal and spatial resolution of Sentinel-1 with the precision of CS2-based delineation. Our project addresses this gap by developing a deep learning-based method that fuses radar (Sentinel-1) and altimetry (CS2) data to delineate ice shelf coastlines more accurately. The study focuses on three key Antarctic ice shelves—Larsen-C, Ronne, and Filchner—which are expansive, have experienced recent calving events, and present complex conditions due to nearby glaciers and islands. The study time period is 2018-2022.

Our approach uses a refinement network where an initial guess coming from the CS2 data is fed (after processing) to the network in addition to the HH and HV S1 feature maps, and when the network outputs the initial ISCL line, it is fed back to the network as a refined guess. That loop commences in multiple iterations until saturation on inference, making an iterative refinement network that iteratively refines the output ISCL until it fits the ice shelf, minimizing the error.

This work is part of the ESA's Antarctica Ice Shelf Climate Change Initiative (AIS CCI+), running until 2026. Preliminary results are presented.

4. Life Cycle Assessment (LCA) for the use of Artificial Intelligence in small and medium enterprises

Kushtrim Visoka¹, Andrea Alessandro Gasparini^{1,2}

1. Norwegian Artificial Intelligence Research Consortium (NORA), Oslo, Norway

2. Department of Informatics, University of Oslo, Oslo, Norway

Sustainability is often omitted when new Artificial Intelligence-based services are launched, and the public and the market are overwhelmed by all the possibilities of use and effectiveness. This poster proposes the use of the Life Cycle Assessment (LCA) approach to address the impact Artificial Intelligence (AI) may have on sustainability. Despite increasing attention on the environmental footprint of AI, there is limited guidance on how small and medium-sized enterprises (SMEs) and public sector can practically apply LCA methodologies to their AI services and infrastructure.

LCA is used primarily to understand the consumption of resources, and the environmental effects on society in all the stages of a product. The holistic view of a life cycle helps companies and organizations address and correct possible impacts their own products may have on the environment and avoid their own effects on climate change.

An additional use of LCA has emerged for organizations to analyze the effects that activities may have on the environment. The Organizational LCA is then used to find “environmental hotspots”, and the long-term environmental performance of the organization which includes those activities usually not directly connected to production.

Adopting LCA on AI is still in its early stages, however, two major directions have emerged. One to assess ready-to-use-systems and one to assess the entire process of developing AI-based resources or services. The former can be relevant for small to medium companies and organizations and using LCA, the result is a complex structure highlighting each step of the process when using AI.

Finally, after each of these points is discussed in the context of small and medium enterprises and public institutions, the poster will conclude why LCA for AI can be a relevant asset. By adopting LCA principles, SMEs and public sector can not only mitigate negative environmental impacts but also position themselves as responsible innovators.

5. The environmental footprint of AI for sustainability purposes

Lina Plataniti¹, Synnøve Rubach², Kari-Anne Lyng¹

1. NORSUS AS, Fredrikstad, Norway

2. Becour AS, Fredrikstad, Norway

AI-based technologies are built to support sustainable development across various sectors, like energy consumption optimisation or precision agriculture. At the same time, the development and deployment of AI is not without environmental concerns. AI systems are often energy-intensive, and their high energy demands have direct climatic impacts. However, energy consumption and its connection to climate change are not the only issues. Environmental impacts related to material resources and emissions, water and land use, and end-of-life management should also be considered from a lifecycle perspective. To enable this, the methodological paths to assess the overall environmental footprint of AI systems need to be identified, and the results should be interpreted and balanced with the benefits of AI for sustainability.

Life Cycle Assessment (LCA) is a well-established methodology for evaluating the overall spectrum of the environmental impacts of a system throughout its lifecycle. According to literature findings, LCA can be applied to AI systems, however, there are several challenges to be addressed. In this poster, we present the potential of the use of LCA as a core methodology for assessing the environmental footprint of AI systems and the methodological frameworks that may be applied. We also discuss the methodological challenges related to system boundaries definition, selection of functional unit, and access to reliable data. We also address the necessity of outcomes that balance the benefits and impacts of AI systems through a case study of an AI-based assistant developed by an energy-tech company helping companies to reach net-zero emission targets.

The need for further work in the field is evident, and therefore we are inviting the AI community to a cross-disciplinary dialogue to establish a deeper understanding of the environmental aspects of AI systems, balancing the benefits with the negative impacts.

6. Towards Energy-Efficient Distributed Training: Empirical Evaluation of Gradient Sparsification Schemes

Preben Castberg¹

1. Department of Informatics, UiO, Oslo, Norway

The rapid scaling of deep learning models like GPT-3 and LLaMA 2 has led to significant environmental costs, with training emissions exceeding 500 tons of CO₂ equivalent per model. These models also face scalability challenges when increasing the number of GPUs, as communication bottlenecks during gradient aggregation introduce high waiting times that prevent linear speedups. This thesis investigates layer-wise gradient sparsification techniques to address these bottlenecks in distributed training.

While a range of compression mechanisms—such as quantization, sparsification, and low-rank approximation—have been developed, most apply compression uniformly across DNN layers, overlooking their heterogeneity in parameter count and sensitivity to compression. Building on recent theoretical work in layer-wise top-k gradient sparsification for distributed stochastic gradient descent (SGD), this thesis aims to evaluate and benchmark these methods in practice. The objective is to compare the energy and accuracy trade-offs of top-k sparsification against other state-of-the-art compression schemes through controlled experiments on multi-node systems.

This work contributes to the broader goal of sustainable machine learning by demonstrating that building energy-efficient models requires not only algorithmic innovation, but also a deeper understanding of how training dynamics interact with hardware constraints. Through empirical evaluation, I aim to inform the design of scalable training systems that are both efficient and environmentally responsible.

7. Learning and predicting fishing activities from AIS data

Klaus Johannsen¹, Xue-Cheng Tai¹, Gro Fønnes¹, Junyong You¹

1. NORCE Norwegian Research Centre AS, Nygårdsgaten 112, 5008 Bergen, Norway

Marine life is vital for providing food, oxygen, and biodiversity. Monitoring illegal fishing is crucial for sustainable marine management. This study employs Automatic Identification System (AIS) data and Norwegian Catch Reports to detect fishing activities. Data from vessels over 15 meters, covering January 2022 to November 2023, was processed into hour-long time series tagged with Catch Report data. The dataset was resampled to create 100

balanced bootstrap sets, resulting in 30 million data points across 2.5 million time series. Using classification methods based on random forests and deep convolution

networks, we achieved 92% to 93.7% accuracy. This study demonstrates the effectiveness of AIS and Catch Report data for automated fishing activity monitoring.

8. Marine Intelligence: Innovations to Enhance Underwater Exploration

Arangan Subramaniam¹, Changkyu Choi², Nils Olav Handegard³, Robert Jenssen², Ali Ramezani-Kebrya⁴

1. University of Oslo (Department of Physics, Oslo, Norway)

2. UiT The Arctic University of Norway (Department of Physics and Technology, Tromsø, Norway)

3. Norwegian Institute of Marine Research (Bergen, Norway)

4. University of Oslo (Department of Informatics, Oslo, Norway)

Foundation models (FMs), large neural networks pretrained on massive image collections, offer powerful, general-purpose representations that can shape many fields. Yet underwater acoustics poses a unique challenge: multi-frequency echosounder outputs six-channel data, while a standard Vision Transformer (ViT) expects three-channel RGB input.

To bridge this gap, we propose a lightweight FM adapter. It learns to compress and remap all six acoustic channels into three semantically meaningful image channels, preserving key scattering features while matching ViT's input format. Because the adapter is tiny compared to the FM backbone, we can train-or even update-, it without retraining the full model. This enables continual learning as ocean conditions shift. Echosounders record returns at six distinct frequencies, capturing the size and location of marine life in high-dimensional, multi-layered data. By translating these volumes into ViT-compatible images, our adapter unlocks the FM's pattern recognition power for real-time ecological monitoring.

At NORA conference, we will present our ongoing work using a ViT as an off-the-shelf FM. Our goal is to explore how a lightweight FM adapter can translate six-channel echosounder data into a format the ViT can understand. We test different adapter designs that reduce and combine the six acoustic frequency channels into three channels that fit the RGB format expected by ViTs, and study how well this setup works for typical marine monitoring tasks.

9. A multimodal transformer model for fishing vessel destination prediction

Helge Fredriksen¹

1. UiT - The Arctic University of Norway

A recent development in the realm of deep learning trajectory prediction for marine traffic is the TrAISformer (Nguyen and Fablet, 2024), which seems competitive in performance, and a well-designed adoption towards the AIS data structure and properties of ship tracks in general.

To enhance the predictive capabilities of the TrAISformer model in vessel destination forecasting, we propose the integration of static vessel and route features into the learning process. Currently, the TrAISformer model does not incorporate such information, despite its potential to improve accuracy. Notable static features include known destination harbors, vessel characteristics (e.g., length, fishing gear type), and historically preferred ports. Moreover, we aim to reframe the task as a classification problem towards predicting known landing harbors for fish catches. This approach aligns with recent work (e.g., Wang et al., 2022) demonstrating improved performance through joint modeling of vessel movements and destination inference. Furthermore, inspired by advancements in Multimodal Foundational Models (Bommasani et al., 2022; Liang et al., 2024), we will consider a future AIS prediction framework that not only integrates static features but also incorporates additional dynamic modalities, such as weather. The transformer architecture's inherent flexibility makes it a promising foundation for such a multimodal predictive system.

10. Examination of Computational Performance of Biological Neural Networks with Reservoir Computing

Felix S. Reimers¹, Ola Huse Ramstad², Axel Sandvig³, Ioanna Sandvig³, Stefano Nichele^{1,2}.

1. Department of Computer Science and Communication, Østfold University College, Halden, Norway

2. Department of Computer Science, Oslo Metropolitan University, Oslo, Norway

3. Department of Neuromedicine and Movement Science, NTNU, Trondheim, Norway

Modern Deep Learning tools like artificial neural networks (ANNs) can perform remarkable computations in several domains. While the human brain is incapable of outperforming ANNs on specific tasks, it is able to perform complex computations as well and works in a plethora of domains, all while using a fraction of the energy required for artificial neural networks. Modern biotechnologies allow to grow biological neural networks (BNNs) in a lab, and the interfacing of them with a computer through Micro-Electrode Arrays (MEAs). As such, it is possible to have bidirectional communication with in-vitro biological neural networks. However, it is rather difficult to control the growth of biological neural networks towards specific topologies or architectures. Additionally, neural cells move and change their connections throughout their lifetimes, making such networks highly dynamical systems. This may cause their computational capacity as well as the optimal inputs and outputs to change over time. Furthermore, the application of specific designer drugs can be used to inhibit connections. Recent neural technologies also allow triggering the development of both healthy networks and networks with specific neural diseases. Recording the neural activity of biological networks on MEAs allows the extraction of functional connectivity graphs. These can be implemented within the Reservoir Computing (RC) framework, where the network parameters are untrained and only a readout module, possibly in the form of a linear layer, is trained. Transforming BNNs into ANNs within the RC framework allows the assessment of their performance in solving problems with regards to their developmental stage, healthiness, dynamical state, and other relevant network states while working on digital snapshots outside the lab. In this work, results from different

reservoir benchmarks on graphs extracted from many neural cultures under different conditions are discussed.

11. Low dimensional representations in boolean reservoir computing

Christopher Michael Skeide Vibe¹, Mikkel Lepperød², Solve Sæbo³, Stefano Nichele^{1,4}

1. Østfold University College (HIOF, Norway)
2. Norwegian University of Life Sciences (NMBU, Norway)
3. Simula Research Laboratory (Simula, Norway)
4. Oslo Metropolitan University (OsloMet, Norway)

The manifold hypothesis in machine learning posits that complex models often solve tasks using simplified internal representations, despite their high complexity. For example, a simple globe might suffice for world navigation, yet an AI architecture may contain trillions of nodes and connections, making it difficult to extract such representations from the complex model. Notably, neuroscience research led by Nobel Prize winners May-Britt and Edvard Moser has revealed that both rodent and human brains represent spatial navigation on the surface of a toroidal shape (donut), facilitated by grid cells wired together in a Continuous Attractor Network (CAN). This discovery indicates that grid cells function as a type of place cell, providing a sense of place.

In the preliminary phase of this research, a spatial navigation task known as path integration is being employed, wiring an artificial intelligence system to emulate the strategy used by grid cells. The primary aim is not merely to solve the task, but to understand the underlying mechanisms by which the task is solved. This research aspires to yield more interpretable AI models and offer insights into neuroscience.

Research Questions:

1. Can complex systems be understood by low-dimensional representations?
2. Would a random graph/wiring converge to certain low-dimensional representations?
3. What graph properties are needed to effectively solve the task?
4. Is the CAN wiring optimal for solving tasks that require spatial representations?
5. Do the connectomes of simple model organisms exhibit low-dimensional representations?

By addressing these questions, this research attempts to bridge gaps between machine learning and neuroscience, providing interpretable AI models and enhancing our understanding of biological neural networks.

12. New Ways to Design Deep Neural Networks

Xue-Cheng Tai¹, Hao Liu², Raymond H. Chan³, Lingfeng Li⁴,

1. Norwegian Research Centre, Nygårdsgaten 112, 5008 Bergen, Norway
2. Department of Mathematics, Hong Kong Baptist University, Kowloon Tong, Kowloon, Hong Kong SAR
3. Lingnan University, Tuen Mun, Hong Kong SAR
4. Hong Kong Centre for Cerebro-Cardiovascular Health Engineering, Hong Kong SAR

In this work, we aim to present a general framework for designing neural networks based on dynamic differential equations using operator-splitting technique. The basic idea is discretizing a continuous control problem using the dynamic system as a constrain which is then approximated by some operator-splitting schemes. We unroll this splitting scheme as a neural network with some control variables parameterized as learnable modules. We also illustrate this approach through examples of the UNet and the PottsMGNet.

13. Probabilistic Kolmogorov-Arnold Networks with Gaussian Process Activations for Uncertainty Quantification

Aaron de Leyos¹, Alexander Johannes Stasik², Signe Riemer-Sørensen²

1. Department of Data Science, NMBU, Ås, Norway

2. Department of Mathematics and Cybernetics, SINTEF Digital, Oslo, Norway

Kolmogorov-Arnold Networks (KANs) have recently been proposed as an alternative to classical feed forward networks (FNN) such as the Multilayer Perceptron (MLP). Inspired by the Kolmogorov-Arnold representation theorem, KANs do not possess learnable linear weights. Instead, they possess learnable activation functions – typically modeled as one-dimensional B-splines. The architecture introduces the possibility for interpretability of the activation functions and parameter efficiency compared to the MLP.

Despite these advantages, standard KANs are deterministic and do not provide a mechanism for uncertainty quantification – a key requirement for applications in science, engineering and safety-critical domains. In this work, we introduce a probabilistic extension of KANs by replacing their spline-based activations with one-dimensional Gaussian processes (GPs). We preserve the core architecture of KANs while endowing them with the ability to express uncertainty through Bayesian inference.

We present a detailed study of the behavior of the combined KAN with GPs (GP-KAN), evaluating their predictive performance and uncertainty calibration across synthetic and real-world tasks. We compare our novel GP-KAN to other established probabilistic modeling approaches, including Bayesian neural networks and standard Gaussian processes, focusing on predictive error and uncertainty estimation quality. Our results show that the combined GP-KANs retain the structural benefits of KANs while offering well-calibrated, flexible uncertainty estimates. The use of one-dimensional activation functions yields adaptive priors that allow for smooth and interpretable uncertainty propagation. This work highlights the potential as a compelling hybrid between kernel-based methods and neural architectures, advancing the development of building interpretable and probabilistic machine learning models.

14. An ubiquitous and pervasive Wireless Sensor Network (WSN) for Active and Assisted Living (AAL)

Matteo Iervasi¹, Florenc Demrozi¹

1. University of Stavanger (Department of Electrical Engineering and Computer Science, University of Stavanger, Stavanger, Norway)

This PhD project focuses on the design and development of a ubiquitous and pervasive Wireless Sensor Network (WSN) for continuous monitoring and management in the Active and Assisted Living (AAL) domain. The system integrates various sensors (e.g., motion, physiological, environmental, and radio-based) with actuators (e.g., speakers and vibration motors) embedded in wearable devices such as smartwatches and smart glasses. Its primary goal is to assist elderly individuals and people with special needs in their daily activities without requiring direct intervention from healthcare professionals.

Healthcare-related scenarios, including Parkinson's disease (PD) management and elderly care, will serve as primary test cases for evaluating the system's effectiveness. In the context of PD, the proposed WSN will enable real-time symptom monitoring, addressing key challenges faced by patients. To mitigate these challenges, the system will provide responsive stimulation and cueing, a technique that employs external signals to enhance movement. This approach aims to help patients overcome motor symptoms such as Freezing of Gait, thereby improving their ability to perform daily activities.

Additionally, machine learning algorithms will be deployed on edge devices to detect specific movement patterns and determine the appropriate cueing strategy in real time. Ultimately, this project aspires to develop a personalized, technologically advanced solution that leverages real-time monitoring and machine learning techniques to enhance the quality of life for the elderly and individuals with special needs.

15. Cross-Linguistic Semantic Analysis for Mild Cognitive Impairment Detection

Arezo Shakeri¹, Mina Farmanbar¹

University of Stavanger

This study explores cross-linguistic semantic and lexical markers for detecting Mild Cognitive Impairment (MCI) through natural language processing. MCI, a prodromal stage of dementia, presents cognitive and linguistic deficits, yet current diagnostic methods are costly and invasive. The research aims to identify language-independent linguistic features indicative of MCI across English and Chinese. Utilizing the TAUADIAL Challenge dataset, comprising 507 image-based speech recordings (261 Chinese, 246 English) from MCI and neurologically normal (NC) participants aged 60–90, the study employs Whisper large-v3 for transcription and language identification. Following preprocessing and language-specific tokenization, twenty linguistic features were extracted, focusing on semantic and lexical aspects such as word repetition, referential ambiguity, lexical diversity, coherence, and semantic drift.

A one-way ANOVA was performed to examine associations between linguistic features and language, task type, and diagnostic category. Significant differences emerged in all features across language and task types, indicating task-dependent language usage. Notably, 11 out of 20 features significantly distinguished MCI from NC groups in both

languages. Features with particularly low p-values included referential_ambiguity, pronoun_to_noun_ratio, pause_frequency, and semantic_word_embedding_distance, reflecting reduced clarity, coherence, and vocabulary complexity in MCI speech. These findings underscore the efficacy of language-based, non-invasive biomarkers for early MCI detection across linguistic contexts. The study highlights the diagnostic value of semantic and lexical analysis in multilingual populations and advocates future research into acoustic features for enhanced multimodal assessment.

16. Longitudinal Analysis of Imaging Data Using Gradient Boosting with Random Effects and Shapley Values

Julia Kropiunig¹, Øystein Sørensen¹

1. LCBC, Department of Psychology, Oslo Norway

To bridge the gap between the interpretability of traditional statistical methods and the flexibility of advanced machine learning, we propose an analytical approach based on GPBoost, a hybrid modelling framework that combines gradient tree boosting with participant-specific random intercepts. This allows for the modelling of complex non-linear and interaction effects, as well as within-subject variability.

In a second step, local dependence-aware Shapley values are computed to quantify each feature's contribution to individual predictions while accounting for feature correlations. These contributions are statistically tested for association with the cognitive outcome, allowing global patterns to emerge while distinguishing between model-learned noise and true signal.

We applied this approach to longitudinal data from the Center for Lifespan Changes in Brain and Cognition (LCBC), consisting of 911 observations from 311 participants aged 20 to 93, each with at least two structural MRI scans and cognitive assessments at least one year apart. Brain measures were decomposed into participant-specific mean and observation-specific change, allowing us to relate both stable differences and temporal changes to cognition.

In particular, mean volumes of the hippocampus, thalamus and nucleus accumbens as well as temporal changes in the nucleus accumbens and caudate nucleus, were predictive of cognitive performance, consistent with the literature. These results show that our approach facilitates interpretable, high-dimensional modelling of brain-behavior relationships.

17. Deep Learning-Based Arrhythmia Segmentation of Porcine ECG Using Implanted Devices

Bjørn-Jostein Singstad^{1,2,3}, Semra Oztemel Sari⁴, Mathis Korseberg Stokke^{4,5,6}

1. Vestfold Hospital Trust (Research and Innovation Department, Tønsberg, Norway)

2. Akershus University Hospital (Medical Technology and E-Health, Lørenskog, Norway)

3. University of Oslo (Institute of Clinical Medicine, Oslo, Norway)

4. Oslo University Hospital and University of Oslo (Institute for Experimental Medical Research, Oslo, Norway)
5. University of Oslo (KG Jebsen Centre for Cardiac Research, Oslo, Norway)
6. Oslo University Hospital (Department of Cardiology, Rikshospitalet, Oslo, Norway)

The identification and quantification of arrhythmic episodes in electrocardiogram (ECG) signals are essential for effectively managing heart diseases. However, manually reviewing long-term recordings for arrhythmic episodes is time-consuming and exhausting. In this work, we segment arrhythmic occurrences from long-term ECG recordings obtained from porcine models. These recordings are acquired using an implanted ECG monitoring device, providing stable measurements compared to non invasive methods.

The ECG data were collected over seven days from 12 individual porcine models with a sample frequency of 1000 Hz. A cardiologist then annotated a one-minute ECG segment for each pig each day. The cardiologist labeled each sample in the one-minute sequence with either Sinus rhythm, Idioventricular Rhythm (IVR), Ventricular Ectopic Beats (VES), noise and Other arrhythmias. The processed and annotated data were then split into 65% training, 15% validation, and 20% test. A 1-dimensional convolutional Unet with temporal self-attention was optimized using a dual loss strategy that integrates Categorical Focal Crossentropy with temporal smoothing loss.

Our model's evaluation on the test dataset produced an F1-score of 0.89 and a dice coefficient of 0.74. The Sinus rhythm class demonstrated the highest accuracy, with an F1-score of 0.96, showing robust detection of normal rhythms. The IVR and Other classes also performed well, with F1-scores of 0.81 and 0.82, while noise and VES classes achieved F1-scores of 0.58 and 0.40.

Our unique dataset offers a best-case scenario for developing arrhythmia segmentation algorithms, as implantable ECG devices in stationary animals yield minimal noise and artifacts. This approach shows potential for arrhythmia monitoring and diagnosis, providing a scalable strategy for automated cardiac event detection. Further research will focus on expanding the annotated training dataset through self-supervised and active learning.

18. Transfer retrieval: A new paradigm for personalised medicine through AI-driven genomic database search

Arina Surko¹, Hasan Oğul¹, Sinan Uğur Umu²

1. Østfold University College (Department of Computer Science and Communication, Halden, Norway)

2. University of Oslo (Department of Pathology, Oslo, Norway)

Genetic diseases can differ significantly among patients due to the variations in the human genome. Personalised medicine aims to facilitate the treatment's efficacy by tailoring it to each patient based on their unique genetic profile. Transfer retrieval is a novel concept for personalised medicine, based on assessing the similarity between

patients' profiles to enable the content-based search of relevant genomic profiles in large, clinically annotated databases. A challenge in developing transfer retrieval is the presence of a variety of technologies and data modalities of gene expression profiles. This diversity complicates the similarity assessment between the profiles. To address this issue, transfer retrieval will enable the similarity assessment by transforming each raw measurement data into a common and biologically interpretable concept. This approach will increase the possibilities of transfer retrieval and, also, keep it agile for future discoveries. More advanced sequencing technologies are emerging, such as single-cell RNA sequencing (sc-RNAseq), which allows analysing each cell from the sample individually. Currently, we are building and testing the transfer retrieval models for the sc-RNAseq data. As for today, two transfer retrieval models are tested. The first model was built as a baseline, and it is based on the average expression per gene across all cells of the sample, imitating bulk RNAseq data. The second one is based on the differential expression data for a cancer dataset. The tested distance metrics include Euclidean and MSE distances, Cosine and Jaccard similarities. The retrieval results are assessed via the Area Under Curve (AUC) metric.

19. Explainable-AI-ready materials modelling data and metadata documentation

Martin T. Horsch¹, Fadi Al Machot¹, Maria Bashir¹, Heinz A. Preisig², Shailendra Singh¹

1. Material Theory and Informatics Group, Norwegian University of Life Sciences (NMBU), Ås, Norway

2. Department of Chemical Engineering, Norwegian University of Science and Technology (NTNU), Trondheim, Norway

This work presents efforts of the Knowledge Graph Alliance (KGA) working group for explainable-AI-ready data and metadata principles in view of requirements for AI systems that comply with the transparency requirements of the AI Act by documenting epistemic metadata, i.e., an annotation that helps establish the knowledge status of data [1].

We discuss how mid-level ontology development for epistemic metadata [2] can help address such requirements and, beyond this, identify pre-existing semantic artefacts and technologies [3] and for implementing systems that provide the digital product passport as specified by the Batteries Regulation, the Ecodesign for Sustainable Products Regulation, and its future delegated acts. Digital twins in manufacturing, e.g., of batteries [4], must be designed to provide the needed data.

Based on the KGA's synopsis of core concepts for XAI-readiness [5], we draft actionable core concepts, i.e., definitions jointly with roles/protocols for such concepts "in action," to help establish both semantic and pragmatic interoperability [6].

This work was funded within Horizon Europe through GA no. 101137725 (BatCAT) and 101138510 (DigiPass CSA).

[1] Horsch, M. T. et al., Epistemic metadata for computational engineering information systems. In Proc. FOIS 2023, IOS, 2023.

- [2] Horsch, M. et al., Exploration of core concepts required for mid- and domain-level ontology development to facilitate explainable-AI-readiness of data and models. In Proc. DAO-XAI 2024, CEUR-WS, 2024.
- [3] Preisig, H. A. et al., Battery manufacturing knowledge infrastructure requirements for multicriteria optimization based decision support in design of simulation. In Proc. SeMatS 2024, CEUR-WS, 2024.
- [5] Al Machot, F. et al. (eds.), Proc. DCLXVI 2024, Springer, 2025.
- [6] Horsch, M. T. et al. Pragmatic interoperability and translation of industrial engineering problems into modelling and simulation solutions. In Proc. DAMDID 2020, Springer, 2021.

20. LangGraph-Powered Conversational AI for Intelligent Sports Highlight Discovery

Mehdi HoushmandSarkhoosh^{1,2}, Cise Midoglu², Saeed Shafiee Sabet², Tomas Kupka², Michael Alexander Riegler^{1,3}, Pål Halvorsen^{1,2,4}

1. Oslo Metropolitan University (Department of Computer Science, Oslo, Norway)

2. Forzasys AS (Oslo, Norway)

3. Simula Research Laboratory (Department of Cyber Security, Oslo, Norway)

4. Simula Metropolitan Center for Digital Engineering (Department of Holistic Systems, Oslo, Norway)

We present a conversational AI system for real-time sports highlight discovery, extending beyond standard RAG and LangChain approaches through structured reasoning, persistent memory, and domain-specific tools within a graph-based architecture. Built on LangGraph and integrated with Forzasys' comprehensive sports databases, the system enables users to intuitively access precise, ranked video highlights using natural language and sustained conversational context. Users can seamlessly pose queries such as "Show Brann's best goals from last month" or "What about goals by Ole Sæter against Bodø/Glimt?", prompting the system to interpret intent, resolve temporal expressions, handle fuzzy matching of entities, and generate structured SQL queries with semantic precision. Unlike traditional RAG systems focused on retrieving unstructured text documents, our framework constructs executable queries directly against structured event databases—containing detailed match statistics, player profiles, and indexed video timestamps. Core modules include advanced temporal parsing, specialized entity disambiguation, matchup resolution, and personalized ranking algorithms prioritizing relevance, match context, and viewer engagement metrics. The architecture explicitly supports explainability, scalability, and real-time media interaction by linking conversational queries directly to curated video assets in Forzasys' content repositories. By combining modular LLM-driven reasoning nodes, robust conversational state management, and specialized sports analytics tools tailored to structured data, our system provides an innovative, production-ready solution for content creators, broadcasters, and fans, enhancing sports storytelling and redefining how audiences experience events such as Ole Sæter's top Eliteserien moments through AI-driven interaction.

21. Information-Theoretic Approach to Learn Causal Relationship Between Question and Answering in DocVQA

Kjetil Indrehus¹, Changkyu Choi², Ali Ramezanikebrya³

1. University of Oslo (Department of Informatics, Oslo, Norway)

2. UiT The Arctic University of Norway (Department of Physics and Technology, Tromsø, Norway)

3. University of Oslo (Department of Informatics, Oslo, Norway)

Deep learning has become the industry standard for solving visual question answering (VQA) tasks. Document Visual Question Answering (DocVQA) extends VQA to complex document layouts with text, tables and graphics. Accurate and accountable DocVQA systems could accelerate tasks in legal review, financial auditing and medical records analysis. However, current models operate as black boxes without explanation methods.

Explainable artificial intelligence (XAI) tackles this through two main strategies. Post-hoc methods is applied to a trained model, but can misrepresent the models true decision making. Self-explainable methods redesign the network to generate explanations with the prediction. We will utilize the second method in the DocVQA setting.

We propose a self-explainable strategy by introducing a causal DocVQA architecture that learns two heatmaps. The answer heatmap localises the textual span predicted as the answer, while the question heatmap isolates document regions that provide semantic evidence. During training we enforce a causal constraint based on information theoretic learning. This ensures that explanations reflect genuine causal influences rather than superficial correlations. Our backbone is Pix2Struct, an image-encoder-text-decoder. We fine tune the encoder on DocVQA and augment it with dual heatmap prediction heads. The total loss combines answer cross entropy with a scaled mutual information regularization term.

22. Monte Carlo Dropout for Model Uncertainty Estimation

Aslak Djupskås¹, Alexander Johannes Stasik^{1,2}, Signe Riemer-Sørensen²

1. Faculty of Science and Technology, Norwegian University of Life Sciences, Ås, Norway

2. SINTEF Digital, SINTEF, Oslo, Norway

In practical applications, particularly those involving safety-critical decisions, machine learning models should not only provide point predictions but also reliable uncertainty estimates. Bayesian machine learning provides a framework for capturing model uncertainty through posterior distributions over parameters or functions. However, exact Bayesian inference in neural networks remains computationally intensive and often intractable in practice.

In the paper “Dropout as a Bayesian Approximation: Representing Model Uncertainty in Deep Learning”, by Yarin Gal and Zoubin Ghahramani from 2015, it was demonstrated that a neural network with dropout applied at inference, before each weight layer, is

mathematically equivalent to an approximation of a deep Gaussian process. However, questions about model calibration remain open.

In this work, we explore the Monte Carlo Dropout (MCD) by analyzing the effects of different dropout rates, regularization techniques, and network sizes. Furthermore, we compare MCD predictions with those from Bayesian neural networks (BNN) and Gaussian processes (GP). For the GP, we test various kernels to identify the most suitable for this specific problem. The BNN model is inferred using Markov Chain Monte Carlo (MCMC) with the No-U-Turn Sampler (NUTS).

To compare these models, we evaluate the results on a simple noisy function under varying amounts of training data, assessing both interpolation and extrapolation performance. In addition, we analyze the covariance matrices and the autocorrelation of the posterior distributions for all three models.

We discovered that the performance of Monte Carlo Dropout is highly dependent on the architectural choices and training configurations. Lastly, MCD results in unjustified high confidence in both interpolation and extrapolation, unlike the GP and BNN.

23. SAIL-K Framework For Secure AI Applications

Ryan Marinelli¹

1. Department of Informatics, University of Oslo, Oslo, Norway

In this work, a new framework, SAIL-K for engaging with AI systems is proposed. The goal of this framework is to augment social blueprints to orient the focus of developers to create more robust and secure AI systems. Key issues are identified per layer with a sketch of solutions to prime intervention.

24. Deconstructing Obfuscation: A Four-Dimensional Framework for LLMs in Assembly Code Deobfuscation

Anton Tkachenko¹, Dmitrij Suskevici¹, Benjamin Adolphi¹

1. Security Research Department, Promon, Oslo, Norway

This research systematically evaluates eight state-of-the-art commercial Large Language Models (LLMs) on their ability to analyze and deobfuscate assembly code from obfuscated binaries. Using a known C program obfuscated with Obfuscator-LLVM, we tested GPT-3o Mini High, GPT-4o, GPT-4.5, O1 Pro Mode, DeepSeekR1, Grok3, Grok2, and Claude 3.7 Sonnet against four obfuscation techniques: bogus control flow, instruction substitution, control flow flattening, and combined techniques. We quantified human intervention required for successful deobfuscation on a six-level scale, from fully autonomous (Level 0) to beyond expert correction (Level 5).

Our findings reveal significant performance variations across techniques and models. While several LLMs successfully deobfuscated individual techniques with minimal guidance, all models universally failed against combined obfuscation. We introduce a

novel four-dimensional theoretical framework—Reasoning Depth, Pattern Recognition, Noise Filtering, and Context Integration—that systematically explains these variations and pinpoints specific capability limitations in current AI systems.

We identified and classified five recurring error patterns: predicate misinterpretation, structural mapping errors, control flow misinterpretation, arithmetic transformation errors, and constant propagation errors. These patterns reveal fundamental limitations in how LLMs process obfuscated code, particularly when mathematical reasoning and complex pattern recognition are required simultaneously.

This research has dual implications for cybersecurity: it highlights significant vulnerabilities in legitimate software protection mechanisms while informing potential defensive applications in malware analysis. Our empirically-derived three-tier resistance model provides actionable insights for developing both more resilient obfuscation techniques and improved automated analysis tools in an increasingly AI-augmented cybersecurity landscape.

25. Detecting AI Influence in Student Writing: Toward Reliable and Interpretable Classifiers

Ibrahim Riza Hallac¹, Abdelaziz Qassi¹, Hasan Ogul¹

1. Department of Computer Science and Communication (Østfold University College, Halden 1757, Norway)

Large language models are now widely used in student writing, raising new challenges for academic integrity, authorship, and fairness. This work develops reliable classifiers to detect AI influence in English-language student essays, distinguishing between texts written fully by students, edited with AI, or generated mostly by AI. We augment publicly available student writing datasets with no known AI involvement by generating AI-edited and AI-generated versions using controlled prompting with large language models. This process yields an augmented dataset for detection studies.

We train transformer-based language models for three-way classification and compare their accuracy and inference-time efficiency with open-source baselines and commercial detectors such as GPTZero. Teachers review a small set of challenging cases to enhance model interpretability and guide alignment. Active learning is used to prioritize uncertain predictions for human feedback, and preference-based tuning helps adjust model outputs to reflect teacher judgments. This work contributes an augmented dataset for detection studies and a comparative evaluation of detection strategies, offering insights into detection effectiveness, model generalization, and the trade-offs involved in educational deployment.

26. Sequential Preference Learning with the Bayesian Mallows Model

Øystein Sørensen¹, Anja Stein², Waldir Leoncio Netto³, David S. Leslie²

1. Department of Psychology, University of Oslo, Oslo, Norway

2. School of Mathematical Sciences, Lancaster University, Lancaster, United Kingdom

3. Oslo Centre for Biostatistics and Epidemiology, University of Oslo, Oslo, Norway

Data in the form of rankings or preferences, such as click data, consumer preferences, social hierarchies, and voting records, are common across various applications, and are challenging to analyze due to their discrete nature.

The Bayesian Mallows model is a flexible tool for analyzing data in the form of complete or partial rankings, and transitive or intransitive pairwise preferences. A particular advantage of the Bayesian Mallows model compared to other alternatives is that it is fully probabilistic, and properly quantifies the uncertainty of all quantities of interest. This is particularly important when using the model output for decision making.

In many applications of preference learning, data arrive sequentially and it is of practical interest to update posteriors and predictions efficiently on-line, based on the currently available data. Despite this, most algorithms proposed so far have focused on batch inference.

We present an algorithm for sequentially estimating the posterior distributions of the Bayesian Mallows model using nested sequential Monte Carlo. This algorithm requires minimal user input for tuning parameters and is straightforward to parallelize. Additionally, it returns the marginal likelihood as a direct byproduct of estimation, making it a viable alternative to Markov chain Monte Carlo techniques in batch estimation settings.

We will present simulation experiments investigating the accuracy of the algorithm and how it scales with the number of CPUs and the size of the data, for a range of realistic settings.

Paper reference:

Sørensen, Ø., Stein, A., Netto, W. L., & Leslie, D. S. (2024). Sequential Rank and Preference Learning with the Bayesian Mallows Model (No. arXiv:2412.13644). arXiv. <https://doi.org/10.48550/arXiv.2412.13644>

27. Comparative Analysis of Bitcoin Price Forecasting Using PatchTST, Traditional Models, and Sentiment Integration

Shakiba Sadat Mirbagheri¹, Claudio Sartori¹, Mehrzad Abdi khalife², Mina Farmanbar³

1. Department of Computer Science and Engineering, University of Bologna, Bologna, Italy

2. Ticker Bell Kft. Company, Budapest, Hungary

3. Department of Electrical Engineering and Computer Science, University of Stavanger, Stavanger, Norway

Predicting Bitcoin prices remains a significant challenge due to the market's high volatility and complex, non-linear dynamics. This study evaluates the performance of a Transformer-based model, PatchTST, in forecasting Bitcoin price movements. Its results are compared against popular statistical and machine learning models, including

ARIMA, Prophet, XGBoost, and various LSTM architectures. We used 10 years of daily Bitcoin data, including price along with technical indicators such as open, close, volume, and the Relative Strength Index (RSI). RSI is integrated both as an input feature and a prediction target to improve the model's ability to capture price momentum and market sentiment indicators, such as overbought or oversold conditions. This dual use of RSI adds interpretability and aligns the model's logic with technical trading strategies commonly employed by human analysts. Experiments were conducted using multiple forecasting windows (24, 36, 48, and 60), with the 24-step history consistently yielding the most accurate results. Data preprocessing steps, including normalization and time-aware splitting, were applied to ensure fair evaluation and prevent data leakage. PatchTST, which breaks time series data into smaller parts for better pattern recognition, gave the best results achieving an R^2 of 0.988, MSE of 0.021, and F1-score of 78.6%.

We also explored integrating sentiment data from Twitter into the models. While sentiment analysis slightly improved some metrics, it also increased certain errors, highlighting the difficulty of effectively using social media signals in price prediction. The findings imply that sentiment analysis has promise, but for improved prediction performance, it needs to be further developed and integrated with better and diverse datasets.

The findings highlight the importance of model choice, data traits, and input features, and suggest hybrid models using technical and sentiment data can improve crypto forecasting.

28. Improving Forecast Resolution in Vector Autoregressive Processes

Luis M. Lopez-Ramos¹, Emilio Ruiz-Moreno², Baltasar Beferull-Lozano²

1. HOST Department, Simula Metropolitan Center for Digital Engineering, Oslo, Norway

2. SIGIPRO Department, Simula Metropolitan Center for Digital Engineering, Oslo, Norway

Vector autoregressive (VAR) models are used to characterize the dynamics of a collection of variables. Any given VAR model is linked to an implicit assumption of the highest rate, or resolution, at which those variables may vary. Hence, those models do not directly allow to predict dynamic variations at finer rates. While it is possible to apply interpolation methods (e.g. splines) to predictions at the original rate, such an approach is prone to overfitting or oversmoothing. Alternative modelling techniques such as mixed-frequency VAR or Stochastic differential-equation models are less interpretable and computationally expensive. To overcome these drawbacks, we propose a method that, given any VAR process and a user-defined rate increase factor, it computes the parameters of a new VAR process and the entries of an upsampling matrix. With these two new elements, one can use the resulting VAR model to predict at the desired resolution while ensuring that i) the process mean is unchanged, ii) the forecast estimates at the rates given by the original resolution do not change, and iii) the long-run associated credible intervals are closely preserved.

To realize this idea, we formulate an optimization problem where the future values predicted by the desired model are enforced to be close to those of the original model

(averaged over process realizations) up to a furthest predictive horizon H . H is a hyperparameter of the proposed method. We performed experiments that (apart from confirming the viability of the proposed method) establish that, given a 2-dimensional VAR(2) process and the goal of generating a VAR(3) process operating at a resolution 3 times higher, it suffices with setting $H = 3$ for the predictions obtained with the new model to closely match the corresponding samples predicted by the original model in the long term, whereas $H < 3$ incurred overfitting. Future work includes studying nonlinear imputation methods to replace the upsampling matrices.