AI & Mathematics

Thursday June 4 **University of Groningen**

9:00 – 10:15	Walk-in with coffee/registration
10:15 –10:30	Opening

Keynote Session - Chair: Tim van Erven (UvA)

10:30 –11:30	**Lars Grüne** (University of Bayreuth)
11:30 – 12:00	Coffee break

Graph Signal Processing and AI - Chair: Palina Salanevich (UU)

12:00 – 12:30	**Kubilay Atasu** (TU Delft)
12:30 – 14:00	Lunch break
14:00 – 14:30	**Gaurav Rattan** (University of Twente)
14:30 – 15:00	**Peter Bloem** (VU Amsterdam)
15:00 – 15:30	**Clara Stegehuis** (University of Twente)
15:30 – 16:00	Coffee break

Sequential decision making - Chair: Wouter M. Koolen (UT)

16:00 – 16:30	**Rianne de Heide** (University of Twente)
16:30 – 17:00	**Jilles S. Dibangoye** (University of Groningen)
17:00 – 17:15	Coffee break
17:15 – 17:45	**George Iosifidis** (TU Delft)
17:45 – 18:15	**Matthijs Spaan** (TU Delft)
19:00	Workshop dinner

αiM
AI & Mathematics

Friday June 5 **University of Groningen**

8:45 – 9:00 Walk-in with coffee/registration

Keynote Session - Chair: Tim van Erven (UvA)

9:00 – 10:00 **Peter Tino** (University of Birmingham)

Inverse Problems and Dynamical Systems – Chair: Kerstin Bunte (RuG)

10:00 – 10:30 **Kanat Camlibel** (University of Groningen)
10:30 – 11:00 **Vivi Rottschäfer** (Leiden University)

11:00 – 11:30 Coffee break

Explainable and formal reasoning in AI – Chair: Kerstin Bunte (RuG)

11:30 – 12:00 **Tim van Erven** (University of Amsterdam)
12:00 – 12:30 **Fabian Gloeckle** (Institut Polytechnique de Paris)

12:30 – 14:00 Lunch break

Scientific Machine Learning – Chair: Dimitris Loukrezis (CWI)

14:00 – 14:30 **Peter Förster** (TU Eindhoven)
14:30 – 15:00 **Nikolaj Takata Mücke** (TU Delft)

15:00 – 15:30 Coffee break

15:30 – 16:00 **Taniya Kapoor** (Wageningen University)
16:00 – 16:30 **Alexander Heinlein** (TU Delft)

16:30 - 17:30 Closing and drinks

Titles and abstracts

Lars Grüne (University of Bayreuth)

Title: Can neural networks solve high dimensional optimal feedback control problems?

Abstract: Deep Reinforcement Learning has established itself as a standard method for solving nonlinear optimal feedback control problems. In this method, the optimal value function (and in some variants also the optimal feedback law) is stored using a deep neural network. Hence, the applicability of this approach to high-dimensional problems crucially relies on the network's ability to store a high-dimensional function. It is known that for general high-dimensional functions, neural networks suffer from the same exponential growth of the number of coefficients as traditional grid based methods, the so-called curse of dimensionality. In this talk, we use methods from distributed optimal control to describe optimal control problems in which this problem does not occur.

Peter Tino (University of Birmingham)

Title: Learning Temporal Structures with Adaptive State Space Models - From Non-Autonomous Dynamical Systems to Temporal Feature Spaces and Back.

Abstract: Parametrized state space models cast, for example, as recurrent neural networks of various flavors, have been often used in machine learning to learn from temporal data. I will present a general view of such systems as driven dynamical systems, followed by a brief historical overview of attempts to understand their learning and representations. Real-world time series often exhibit the so-called fading-memory property, that can be captured by much simplified classes of state space models known as reservoir computers. I will show that for reservoir models there is a deep connection between their state space and the notion of feature space in kernel machines (a well-known class of machine learning methods). This viewpoint will lead to several rather surprising results regarding reservoir model structures. I will also present statements regarding their universal approximation capabilities.

Kubilay Atasu (TU Delft)

Title: Graph Neural Networks for Edge-Attributed Multigraphs

Abstract: Graph neural networks have emerged as Swiss Army knives for machine learning on graph-structured data, yet standard message-passing GNNs exhibit important theoretical and practical limitations in how they handle multigraphs with rich edge attributes, such as those found in financial and social networks. In this talk, I will present some of our recent advances introducing specialised GNN architectures for edge-attributed multigraphs and demonstrate their effectiveness in tasks including subgraph pattern detection, social interaction modelling, and financial crime analysis. I will also discuss several open challenges, particularly regarding scalability, privacy, and robustness.

Gaurav Rattan (University of Twente)

Title: Learning on Graphs with Homomorphism Tensors

Abstract: Graph Neural Networks (GNNs) are deep learning models for graphs and networks. The structural theory of graph homomorphisms has recently emerged as a promising candidate for analyzing the expressive power of GNNs. Inspired by earlier work of Lovasz and others, we present a novel algebraic framework, namely Graph Homomorphism Tensors, as a unified language for learning on graphs. Our framework allows us to translate a wide variety of combinatorial operations in graphs as algebraic operations on tensors and vice-versa. More importantly, these tensors serve as numerical analogs of Boolean logical formulas on graphs, hence facilitating a better understanding of expressivity and generalization in GNNs.

Peter Bloem (VU Amsterdam)

Title: Towards a foundation model for knowledge graphs

Abstract: Knowledge graphs are a natural way to capture structured, multi-modal data. They can help us understand data in many modalities, including natural language and images, as well as the way these data interrelate. We propose to build a foundation model in the vein of GPT and TabPFN, which natively understands multi-modal knowledge graphs. This talk will go over the technical challenges: training large models when little data is available, generalizing to unseen graph representations, learning generative models for knowledge graphs, and defining a framework in which many useful tasks can be solved in-context, including ones that are not necessarily related to knowledge graphs.

Clara Stegehuis (University of Twente)

Title: Colored interactions: from graphs to hypergraphs

Abstract: Traditional social network analysis often models homophily, the tendency of similar individuals to form connections. using a single parameter. We will show that in many important applications, such as hypergraphs or temporal contact networks, homophily occurs at several different scales. We present a model that integrates these different homophily values through a random graph model with a maximum entropy approach. We demonstrate that the interaction between different levels of homophily results in complex percolation thresholds. Furthermore, we show that our model fits remarkably well on a wide range of data sets, capturing their homophily patterns accurately.

Rianne de Heide (University of Twente)

Title: Multiple testing with e-processes

Abstract: Sequential decision making calls for statistical tools that remain valid under optional stopping, adaptivity, and evolving testing goals. In this talk, I will discuss recent advances in multiple testing with e-values and e-processes.
I will first introduce the e-Closure Principle, a general necessary and sufficient principle for constructing multiple testing procedures. I will then discuss a challenge that arises when not testing a single hypothesis sequentially, but multiple: one can "regret" having collected more data for one hypothesis, as it can "undo" a rejection of another. Finally, I will present ongoing work on a subtle but important issue in multiple testing with e-processes: the difficulty of aligning filtrations across hypothesis streams, and the consequences this has for establishing not only valid but also practical multiple testing guarantees in adaptive sequential environments.

References:

Xu, Z., Solari, A., Fischer, L., de Heide, R., Ramdas, A., & Goeman, J. (2025). Bringing closure to false discovery rate control: A general principle for multiple testing. arXiv preprint arXiv:2509.02517.

Tavyrikov, Y., Goeman, J. J., & de Heide, R. (2025). Carefree multiple testing with e-processes. arXiv preprint arXiv:2501.19360.

Jilles S. Dibangoye (University of Groningen)

Title: Bellman Recursion and Convex Geometry in Partially Observable Stochastic Games

Abstract: Partially observable stochastic games (POSGs) combine temporal propagation, imperfect information, strategic response, and decentralized execution. In standard formulations, these roles are intertwined, which makes Bellman structure difficult to expose. In this talk, I will present a lossless recursive reformulation for finite-horizon POSGs that separates Bellman propagation, strategic admissibility, and execution feasibility.

This separation yields a common recursive foundation across several strategic settings. It also reveals a geometric structure of the same Bellman object: after resolving admissibility and rewriting the value in suitable conditional-mixture coordinates, one obtains a convex representation of the selected value under the assumptions of the geometric analysis. I will discuss how general-sum, zero-sum, common-payoff, and Stackelberg models fit into this picture as different semantic resolutions of a shared Bellman backbone.

The emphasis of the talk will be structural rather than algorithmic: the goal is to identify the representational layer on which Bellman recursion closes, and on which principled approximation may later be built.

George Iosifidis (TU Delft)

Title: Fast online learning with unreliable forecasters

Abstract: Modern AI systems increasingly require a principled combination of offline and online learning: offline data and pretrained models provide strong priors, while online adaptation is essential to cope with distribution shifts and evolving environments. Bridging these two regimes in a mathematically rigorous way remains a central challenge, particularly when aiming for both high empirical performance and reliable and interpretable learning guarantees. This tutorial presents optimistic learning as a unifying framework to address this challenge. The key idea is to incorporate black-box predictions, learned offline or inferred from historical data, into online learning algorithms, thereby accelerating learning when predictions are accurate, while retaining worst-case convergence guarantees when they are not. This approach offers a systematic way to combine data-driven models with adaptive decision-making algorithms. The talk will conclude with applications of optimistic learning to selected key problems in modern communication systems.

Matthijs Spaan (TU Delft)

Title: Exploiting Epistemic Uncertainty for Deep Exploration in Reinforcement Learning

Abstract: Reinforcement Learning (RL) allows an autonomous agent to optimize its decision making based on data it gathers while exploring its environment. Given limited and possibly inaccurate data, the agent is uncertain regarding its state of knowledge, which is referred to as epistemic uncertainty. Estimates of such epistemic uncertainty can guide an agent's decision making, notably where to focus its exploration of the environment. The principled embedding of epistemic uncertainty in present-data reinforcement learning is an important open problem.

In this talk, I will present recent work on exploiting epistemic uncertainty estimates in hard-exploration problems. First, our approach called Sequential Monte-Carlo for Deep Q-Learning studies uncertainty quantification for the value function in a model-free RL algorithm by training an ensemble of models to resemble the Bayesian posterior. Second, our Projection-Ensemble DQN algorithm focuses on the distributional RL setting and increases the diversity of an ensemble of distributional value functions by employing different projections of value distributions for different ensemble members. Third, our Epistemic Monte Carlo Tree Search methodology incorporates epistemic uncertainty into model-based RL by estimating the epistemic uncertainty associated with predictions at every node in the MCTS planning tree.

We demonstrate our algorithms on a variety of hard-exploration benchmarks, showing that they succeed in outperforming state-of-the-art baselines and highlighting how exploiting epistemic uncertainty brings about these improvements.

Kanat Camlibel (University of Groningen)

Title: Learning controllers with provable guarantees: The data informativity framework

Abstract: In this talk, we will present a data-driven framework for control design that provides provable guarantees directly from data without requiring an explicit model of the system. The approach is based on the notion of data informativity which characterizes what can be reliably inferred or achieved given a finite dataset. Rather than identifying a single model, the framework considers the entire set of systems consistent with the data and derives conditions under which a desired property, such as stability or performance, holds uniformly across this set. This leads to controller design methods with formal guarantees of correctness and robustness even in situations where system identification is not possible. From an AI perspective, the framework can be viewed as a principled way of extracting certifiable knowledge from data enabling safe and reliable decision-making in control tasks.

Vivi Rottschäfer (Leiden University)

Title: Can you always determine parameters from data in a unique way?

Abstract: In using numerical simulations of models of differential equations to fit experimental data and estimate these parameters, the theory of the identifiability properties of the system is often overlooked. In this talk, I will consider the question of whether parameters in a mathematical model can be determined {\it uniquely} from available data. I will introduce structural identifiability analysis (SIA) and give a brief overview of three classical SIA methods (transfer function, Taylor Series and similarity transformation) and apply these to drug-receptor binding models.

Tim van Erven (University of Amsterdam)

Title: Formal Results in Explainable AI

Abstract: Since most machine learning systems are not inherently interpretable, explainable machine learning tries to generate explanations that communicate relevant aspects of their internal workings. This is a relatively young subfield, which is generating a lot of excitement, but it is proving very difficult to lay down proper foundations: What is a good explanation? When can we trust explanations? Most of the work in this area is based on empirical evaluation, but there are some noteworthy mathematical results. I will highlight several formal results, including negative results about so-called Shapley values from cooperative game theory, which are often used in practice, but are not nearly as well supported by theory as is commonly thought.

Fabian Gloeckle (Institut Polytechnique de Paris)

Title: Formalizing 1000 Textbooks

Abstract: Multi-agent LLM systems are fundamentally changing the formalization landscape. We recently demonstrated this by formalizing a 500 page graduate textbook in algebraic combinatorics to Lean within one week, using 30K Claude 4.5 Opus agents collaborating in parallel on a shared repository. The resulting 130K lines of code and 5,900 theorems simultaneously set a record in multi-agent software engineering. The inference cost matched or undercut the estimated salaries of a comparable human expert team, with significant efficiencies still on the table. In this talk, I will situate this result within the broader research landscape, discuss best practices and open challenges for LLM-driven formalization, and describe what a world might look like in which we'd have 1000x'd the volume of formalized mathematics.

Peter Förster (TU Eindhoven)

Title: A non-intrusive and physics-consistent machine learning framework for electric circuits

Abstract: Electric circuits are commonly modeled using modified nodal analysis (MNA). One of the peculiarities of MNA is that it can be classified as a differential-algebraic equation (DAE), i.e. a system of implicit differential equations also containing hidden algebraic constraints. (Think e.g. Kirchhoff's laws in the context of circuits.) This DAE structure allows one to decouple MNA into its differential and algebraic parts, describing the dynamics and (hidden) constraints of the system, respectively. We introduce an approach to compute this decoupling using only topological information about the circuit, in particular enabling the non-intrusive identification of the differential and algebraic variables, respectively associated with the differential and algebraic parts of the DAE. In a machine learning setting, one can then only learn the differential variables, while reconstructing the algebraic variables in a physics-consistent manner by exploiting the decoupling. Moreover, this reconstruction can also be achieved non-intrusively in the context of electric circuits, by leveraging properties of the implicit Euler method. The entire workflow is demonstrated using an example from power electronics.

Nikolaj Takata Mücke (TU Delft)

Title: Generative Modeling for Physics-Aware Forecasting and Data Assimilation via Stochastic Interpolants

Abstract: Stochastic interpolants provide a flexible generative modeling framework for probabilistic forecasting by transforming samples between arbitrary distributions via stochastic differential equations. In this talk, I present two complementary developments that leverage this framework for scientific applications: physics-aware probabilistic forecasting and posterior sampling for data assimilation. First, I address a key limitation of existing generative models for fluid dynamics, namely their disregard for underlying physical laws. I introduce an energy-consistent stochastic interpolant which is optimized to conserve kinetic energy in expectation, combined with a divergence-free projection step during inference. Evaluated on a two-dimensional turbulent flow, this approach achieves stable rollouts over thousands of time steps. Second, I show how a pre-trained stochastic interpolant can be adapted for Bayesian data assimilation without retraining, by deriving a posterior SDE that incorporates observational information through a likelihood score. To make this practical, I introduce an observation interpolant that enables tractable likelihood score computation. I demonstrate the method on turbulent flows simulated with the stochastic Navier-Stokes equations and on turbulent atmospheric flows in an urban environment. Together, these contributions illustrate the versatility of stochastic interpolants as a unifying framework that can generate physically consistent forecasts and seamlessly incorporate observational data for uncertainty quantification in high-dimensional dynamical systems.

Taniya Kapoor (Wageningen University)

Title: Less Data, Less Compute, More Physics: Sustainable Scientific Machine Learning

Abstract: Machine learning has proven effective for problems involving complex systems, for example in weather prediction, fluid flow, and material design. But many modern AI methods still need huge datasets and lots of computing power, which can be expensive and waste energy. In this talk, I share a more sustainable idea: use physics to make machine learning smarter. Instead of learning everything only from data, we can add known physical rules such as conservation of energy, symmetry, and governing equations. This helps models learn faster, need less training data, and give more reliable results, especially in situations where we don't have many measurements. I will briefly introduce popular approaches like Physics-Informed Neural Networks (PINNs), neural operators, and hybrid methods that combine ML with traditional simulations. The main message is simple: when we add physics, we can reduce data and compute while improving accuracy. This makes scientific AI greener, cheaper, and more accurate for real-world science and engineering.

Alexander Heinlein (TU Delft)

Title: Localization and Efficient Training of Scientific Machine Learning Models

Abstract: Scientific machine learning (SciML) combines scientific computing and machine learning to solve complex physical problems. In this talk, we discuss neural network and neural operator approaches for differential equations, trained using data and/or physics-based loss functions. The focus of this talk is on two aspects. First, localization techniques based on domain decomposition are used to improve scalability, reduce computational costs, and help the models capture challenging spatio-temporal scales. Second, we discuss the training of these models, which constitutes the major computational cost. We investigate the training dynamics and how training may be improved. Numerical experiments on representative academic model problems (including multiscale problems) illustrate the performance of the discussed approaches.