Living the DREAM – seminar series
LUT University

Wednesday **June 3rd, 2026** at 9:55 – 15:20 in **Room 1325**, Lappeenranta campus

Zoom information

Meeting ID: 694 0092 8755
Passcode: 696262
`https://lut.zoom.us/j/69400928755?pwd=nS3iba4Qgx5NDQ9oGfIh6dboK5AoAk.1`

Schedule

	Session 1	
9:55	Welcome	
10:00	Talk 1	Oula Kekäläinen: *”Relating geometric inverse problems for different families of generalized geodesics”*
10:15	Talk 2	Ville-Petteri Manninen: *”Bayesian inverse eigenvalue problems via surrogate modeling with sparse stochastic collocation”*
10:30	Talk 3	Emma Hannula: *”Scalable Bayesian deep learning via amortized inference and partial stochasticity”*
10:45	Coffee (30 min)	
	Session 2	
11:15	Talk 4	Subhendu Pramanick: *”Hyperparameter estimation in Gaussian process modelling”*
11:30	Talk 5	Ammar Kheder: *”Deep learning for air quality forecasting”*
11:45	Talk 6	Samuel Repka: *”Mineral segmentation for scanning electron microscopy”*
12:00	Talk 7	Jarmo Flander: *”Mathematical phase retrieval”*
12:15	Lunch (1h)	
	Session 3	
13:15	Talk 8	Sara Heikkinen: *”Remote sensing applications on lake quality monitoring”*
13:30	Talk 9	Sachitha Alwis Weerasinghe: *”Smart mobility and sports performance analytics”*
13:45	Talk 10	Jenni Köykkä: *”Lagrangian modeling of aerosol transport and chemistry in the Arctic”*
14:00	Talk 11	Samuel Agenorwoth: *”Uncertainty and computation in acoustic imaging of pipes”*
14:15	Coffee (30 min)	
	Session 4	
14:45	Talk 12	Olivier Dondjio: *”Bayesian optimization with inhomogeneous smoothness”*
15:00	Talk 13	Akseli Suutari: *”Resource-efficient machine learning for particle jet tagging”*
15:15	Closing words	

Attached are 13 abstracts in alphabetical order.

Uncertainty and Computation in Acoustic Imaging of Pipes.

Emilia Blåsten[1,2], Samuel Agenorwoth[1,2]

[1] Computational Engineering, Lappeenranta-Lahti University of Science and Technology, Lahti, Finland [2] Atmospheric Modelling Centre, Lahti, Finland

Abstract

We compare two inverse algorithms for reconstructing the cross-sectional area of a pipe from boundary measurements: the SG approach, adapted from vocal-tract imaging, and the KLO boundary control method, adapted from the wave speed recovery.

Using randomly generated area profiles and two complementary forward models, we evaluate reconstruction accuracy under varying noise levels. Our statistical analysis based on relative errors and paired comparisons examines how the relative performance of the two methods varies with the data-generation scenario and noise level.

We conclude with a summary of the key numerical findings and visual examples that illustrate the behavior of both algorithms.

Bayesian Optimization with Inhomogeneous Smoothness

Toni Karvonen[1], Olivier Dondjio[2]

[1 2] Computational Engineering, LUT University, Lappeenranta, Finland

Abstract

Bayesian optimization (BO) is a popular global optimization technique that uses a Gaussian Process (GP) as a surrogate model for the objective function. Traditional BO methods often assume a stationary GP kernel, meaning that the function is assumed to have the same level of smoothness everywhere in the input space. However, many real-world objective functions exhibit varying smoothness across the input space. To address this limitation, we propose a BO framework that employs locally adaptive estimation of GP kernel hyperparameters, allowing the model to adapt to local variations in smoothness and improve optimization performance. Preliminary benchmark results show that our proposed framework successfully adapts to varying smoothness, consistently outperforming standard stationary baselines.

Mathematical Phase Retrieval

Jarmo Flander[1]

[1] Department of Computational Engineering, LUT University, Lappeenranta, Finland

Abstract

In many imaging problems, measurements record only the intensity of a wave, while the phase information is lost. This missing phase is often essential for reconstructing the original signal, which leads to the mathematical problem known as phase retrieval. The problem appears in areas such as X-ray crystallography, diffraction imaging, and microscopy.

This talk offers an accessible introduction to mathematical phase retrieval. It explains what the phase retrieval problem is, why the loss of phase makes reconstruction difficult, and which questions mathematics seeks to answer. In particular, the focus is on the uniqueness and stability of the signal reconstruction when the measurements are affected by Gaussian white noise.

Scalable Bayesian Deep Learning via Amortized Inference and Partial Stochasticity

Emma Hannula[1]

[1] Computational Engineering, LUT, Lappeenranta, Finland

Abstract

Bayesian deep learning aims to provide neural networks with uncertainty quantification, moving from estimates toward full posterior distributions over model parameters. Despite considerable progress, challenges of scalability, computational efficiency, and expressive posterior approximation still remains. In this work, we present two approaches to this problem from different angles, each targeting a distinct bottleneck in the deployment of Bayesian methods for deep learning.

In the first part, we employ BayesFlow, a simulation-based inference framework based on normalizing flows, to perform amortized posterior estimation over the parameters of simulator models. By training a conditional normalizing flow on simulated parameter-data pairs, the framework learns a reusable posterior approximator that does not require additional simulation at inference time. This amortization property makes the approach particularly beneficial to settings where the same model structure is repeatedly fitted to new observations, such as in computational neuroscience or epidemiological modelling.

In the second part, we investigate partially stochastic neural networks, in which a carefully chosen subset of network parameters are treated as random variables while the remainder are kept deterministic. This selective stochasticity offers a practical middle ground between computationally expensive fully Bayesian networks and overconfident deterministic models, preserving meaningful uncertainty estimates at a significantly reduced computational cost.

Remote sensing applications on lake quality monitoring

Sara Heikkinen[1], Zina-Sabrina Duma[1], Tuomas Sihvonen[1], Jesse Railo[1], Satu-Pia Reinikainen[1]

[1] Computational Engineering, Computational Spectroscopy Group, LUT University, Lappeenranta, Finland

Abstract

Remote sensing is a widely utilized methodology in environmental monitoring due to its extensive coverage from a single image. This approach enables monitoring in hard-to-reach areas, primarily through satellite data acquisition. Satellites are equipped with sensors capturing data across multiple wavelengths, resulting in multi- or hyperspectral datasets. Such spectral data provide comprehensive information, facilitating effective modelling when combined with ground truth measurements. This study employs imagery from the Landsat 8-9 satellite system as it is widely applied in environmental monitoring applications.

Finland employs automatic measurement stations in select lakes to collect hourly data, which can be correlated with satellite imagery to calibrate retrieval models for specific water quality traits. These models enable estimation of trait concentrations at various sites of interest. This study utilizes data from two Finnish lakes, with measurements focusing on Chlorophyll-a and turbidity to assess lake quality.

This study evaluates the impact of pansharpening on Chlorophyll-a concentration mapping for Lake Saimaa, Finland, by using original and pansharpened Landsat 8-9 images. Pansharpening is a data fusion technique used to enhance spatial resolution in multi- or hyperspectral satellite imagery. Additionally, correlations between Chlorophyll-a levels, turbidity concentrations, and satellite band-based indices are examined to elucidate the spatial variability across different regions of Lake Vesijärvi, located in Lahti, Finland.

Relating geometric inverse problems for different families of generalized geodesics

Oula Kekäläinen

LENS, LUT University, Lappeenranta, Finland

Abstract

Geometric inverse problems are a subset of inverse problems in which the goal is to recover geometric information through indirect measurements of a mathematical object. A classical example of an inverse problem is X-ray computed tomography (CT), where attenuation data of X-rays is measured passing through an object and the goal is to reconstruct an image of the inside of the object. This problem can be generalized to the geometric setting by replacing the straight lines of the X-ray with geodesics on a Riemannian manifold (M, g).

The motivation for this project comes from studying the magnetic potential ($\mathcal{MP}$) system given by a Lorentzian force and a potential. One can show that a geodesic of the $\mathcal{MP}$ system can be understood as reparameterization of a geodesic for a different simpler system, where the metric is obtained from g by a conformal change. Using this fact, one can solve a geometric inverse problem for the $\mathcal{MP}$ system using existing results of the simpler system.

In this talk, we discuss building a framework under which we can move from one dynamical system to another and understand how geometric inverse problems of one system are related to the other.

Deep Learning for Air Quality Forecasting

$PM_{2.5}$ are fine particles suspended in the air, emitted by traffic, industry, and wildfires. Because of their tiny size, they penetrate deep into the lungs and bloodstream, and long-term exposure is strongly linked to cardiovascular and respiratory disease. Forecasting their concentration one to several days ahead helps public health authorities issue warnings and supports policy decisions on emission controls. However, existing operational models operate at coarse spatial resolution, making it impossible to identify pollution hotspots at the neighbourhood level.

My doctoral research develops **physics-guided deep learning** architectures that embed physical knowledge about how wind transports pollution and how terrain blocks or accumulates particles directly into neural network structure, enabling high-resolution $PM_{2.5}$ forecasting over large domains.

The work follows a progressive trajectory. An initial contribution, **AQ-Net** [1], introduced a deep spatio-temporal architecture for air quality reanalysis over Finland, establishing the foundations of the approach (published, SCIA 2025). Building on this, **TopoFlow** [2] encodes terrain and wind direction into the attention mechanism of a Vision Transformer to predict $PM_{2.5}$ over China, significantly improving the capture of pollution accumulation in complex topography (last stage of review, *npj Climate and Atmospheric Science*). **CRAN-PM** [3] then scales this approach to continental Europe via a dual-branch architecture that bridges coarse meteorological data with fine-scale local observations through physically-informed cross-attention (under review, ECCV 2026).

The next step of this research is to extend the framework to a global scale, moving towards a worldwide high-resolution air quality forecasting system.

Keywords: $PM_{2.5}$, air quality forecasting, physics-guided deep learning, Vision Transformer

References

[1] A. Kheder et al., “TopoFlow: Physics-guided Neural Networks for High-Resolution Air Quality Prediction,” arXiv:2602.16821, 2026. Last stage of review, npj Climate and Atmospheric Science.

[2] A. Kheder et al., “CRAN-PM: Cross-Resolution Attention Network for High-Resolution $PM_{2.5}$ Prediction,” Under review, ECCV 2026.

[1] A. Kheder, B. Foreback, L. Wang, Z. Liu, M. Boy, “Deep Spatio-Temporal Neural Network for Air Quality Reanalysis,” SCIA 2025. LNCS, vol. 15725, pp. 74–87, 2025.

Lagrangian Modeling of Aerosol Transport and Chemistry in the Arctic

Jenni Köykkä[1,2], Robin Wollesen de Jonge[3], Petri Clusius[2,3], Zhi-Song Liu[1,2], Michael Boy[1,2,3]

[1] School of Engineering Sciences, Lappeenranta-Lahti University of Technology LUT, Lahti FI-15110, Finland
[2] Atmospheric Modelling Centre – Lahti, Lahti University Campus, FI-15110 Lahti, Finland
[3] Institute for Atmospheric and Earth Systems Research, University of Helsinki, FI-00014 Helsinki, Finland

Abstract

Aerosols play a key role in the Arctic climate system by affecting radiative balance, cloud properties, and Arctic-specific feedback loops.They originate from sea salt emitted by the ocean and blowing snow, long-range transport of sulfate and black carbon from mid-latitudes, and secondary formation from natural and anthropogenic precursors. Their effects depend strongly on composition: black carbon and mineral dust absorb sunlight and reduce snow and ice albedo, enhancing warming, while sulfate aerosols scatter sunlight and cool the surface. Still, Arctic aerosol sources and formation pathways remain uncertain.

To address this, we investigate the sources and chemical evolution of Arctic aerosols using two Lagrangian 1D chemistry models, ADCHEM and SOSAA-FP, which differ in emission handling and particle-phase chemistry. Comparing the models allows us to quantify key uncertainties, and we do this by analyzing two Greenland case studies. At Disko Island, west coast of Greenland (August–September 2023), aerosols are dominated by secondary organic aerosol, DMS-derived particles, and black carbon, including contributions from Canadian wildfires. At Villum Research Station, northeastern Greenland (March–August 2024), particle formation and growth are largely driven by DMS from phytoplankton blooms in the Norwegian, Greenland, and Barents Seas.

Finally, we aim to extend this work using machine learning. While particulate matter with diameter less than 2.5 μm (PM2.5) is a major contributor to premature mortality, its mass concentration alone does not capture toxicity. Instead, particle size distribution (PSD) is critical, as smaller particles penetrate deeper into tissues. Our goal is to train a machine learning model that predicts the PSD from PM2.5 measurements.

Bayesian Inverse Eigenvalue Problems via Surrogate Modeling with Sparse Stochastic Collocation

Tapio Helin[1] Nuutti Hyvönen[2] Ville-Petteri Manninen[1]
Mikael Saarikangas[3] Matti Lassas[4]

[1] Computational Engineering, LUT, Lappeenranta, Finland
[2] Department of Mathematics and Systems Analysis, Aalto University, Espoo, Finland [3] Metso Oyj, Helsinki, Finland
[4] Department of Mathematics and Statistics, University of Helsinki, Helsinki, Finland

Abstract

Inverse problems are a class of optimization problems where the objective is to recover unmeasurable quantities, for example, of a physical phenomenon, through accessible information related to the quantity of interest. Mathematically, this can be expressed as a function $G : x \mapsto y$, where x is the input and y is a measurable output, leading to the objective of recovering x given a noisy measurement of the output y.

The Bayesian framework is one of the standard approaches to solving inverse problems and estimating uncertainty, but the methods require substantial computational effort to produce the desired quantities. Our research aims to demonstrate the use of an approximative (surrogate) model $\tilde{G} \approx G$ in the setting of Bayesian inverse eigenvalue problems, to greatly reduce the computational cost while maintaining accuracy.

Our method of choice is to utilize stochastic finite elements to construct a polynomial approximation of the forward map via sparse stochastic collocation. Prior research on the topic has established the approximation capabilities of the approach, providing a solid foundation for its application to inverse problems.

We aim to demonstrate the capabilities of the methodology, both in theory and in practice, using the 1D inverse Sturm–Liouville problem as an example. The result of our research provides an overview for application in Bayesian inverse eigenvalue problems, with spatially continuous underlying parameters.

Hyperparameter Estimation in Gaussian Process Modelling

Subhendu Pramanick[1]

[1] Computational Engineering, Lappeenranta–Lahti University of Technology, Lappeenranta, Finland

Abstract

As a data-driven approach, the Gaussian Process (GP) modelling is a powerful Bayesian framework for nonlinear regression that requires minimal prior functional knowledge. The model's flexibility and predictive power are primarily encoded in its covariance function (kernel), which defines the probability distribution over sample functions. The geometric properties of these functions—such as mean-square differentiability (smoothness), periodicity, and characteristic length-scales—are governed by a set of latent parameters known as hyperparameters. Consequently, hyperparameter estimation is the fundamental process that aligns the GP with the observed data. By aligning the kernel with underlying data patterns, accurate estimation ensures reliable uncertainty quantification in GP modeling for risk-aware decision-making.

Mineral segmentation for scanning electron microscopy

Samuel Repka[1,2], Tuomas Eerola[1], Pavel Zemčík[1,2], Matej Koscelník[2], David Motl[3], Jakub Výravský[3]

[1] Computational Engineering, Lappeenranta-Lahti University of Technology, Lappeenranta, Finland
[2] Faculty of Information Technology, Brno University of Technology, Brno, Czechia
[3] TESCAN GROUP, Brno, Czechia

Abstract

Segmentation of mineral images, i.e., classifying pixels into specific mineral types, is a fundamental task in the geological sciences. In this work, we focus on data from Scanning Electron Microscopes (SEMs) and utilize two distinct data sources to segment the image: Backscattered Electrons (BSE) image and sparse Energy-Dispersive X-Ray Spectroscopy (EDS) data. While the BSE image is fast to acquire, it lacks sufficient discriminatory power for accurate mineral classification. In contrast, the EDS measurements provide comprehensive information about the chemical composition but are time-consuming to acquire. These complementary characteristics motivate the fusion of sparse EDS data with dense BSE images to achieve both fast data acquisition and processing, as well as compositional EDS information on critical locations. This introduces two challenges: (1) how to select the most informative sampling locations for EDS while minimizing the number of measurements, and (2) how to fuse the BSE image data with pointwise EDS measurements for segmentation. This work proposes a new way to select the most informative sampling points and leverages graph neural networks for data fusion and segmentation. The method is capable of effectively working with sparse EDS data and is faster than previous state-of-the-art methods. The approach accelerates the mineral segmentation process in two ways: by enabling the use of sparser data to achieve the same accuracy and by reducing data-processing time.

Resource-Efficient Machine Learning for Particle Jet Tagging

Akseli Suutari[1], Alex Karonen[1], Henri Petrow[1],
Santeri Laurila[2,3],Henning Kirschenmann[1,2], Lasse Lensu[1]

[1] Department of Computational Engineering, LUT University, Lappeenranta, Finland
[2] Helsinki Institute of Physics, Finland
[3] European Organization for Nuclear Research (CERN), Geneva, Switzerland

Abstract

Due to the increase in data volume caused by the high-luminosity upgrade of the CERN Large Hadron Collider (LHC), there is a need for methods for real-time data reduction with strict latency constraints. High-interest particles such as the Higgs boson and W and Z bosons create boosted jets, which creates a motivation to identify these jets at the Compact Muon Solenoid (CMS) detector's Level-1 trigger (L1T). Jet tagging is traditionally performed by large offline models, but methods such as distributed arithmetic optimization with `da4ml` and quantization-aware training enable the efficient use of machine learning models on field-programmable gate arrays (FPGAs). Leveraging these methods, we present a model for tagging boosted Higgs jets that meets the latency and resource constraints of the FPGAs in CMS L1T.

Smart Mobility and Sports Performance Analytics

Sachintha Alwis Weerasinghe

Department of Mechanical Engineering, LUT University
Yliopistonkatu 34, 53850 Lappeenranta, Finland

Abstract

Quantitative assessment of propulsion and pole-use mechanics in real-world sport environments remains challenging because many performance-related variables cannot be directly observed during training or competition. This work proposes smart wheelchair and ski-pole prototypes as instrumented sensing platforms for mobility and sports performance analytics, with particular focus on para wheelchair racing and cross-country skiing. The main objective is to transform these sports interfaces into practical measurement systems that capture force, motion, and temporal patterns during natural use, enabling more informative analysis than conventional video-only or laboratory-based assessments.

From a **sensing** perspective, the proposed systems integrate embedded sensors such as inertial measurement units, force-sensitive elements, and timing-related measurements to capture propulsion cycles, pole-ground interaction, asymmetries, and movement variability. Raw signals are pre-processed through filtering and synchronization, and relevant features are extracted in both time and frequency domains. For example, Fourier-based spectral analysis can be used to identify dominant propulsion or pole-plant frequencies, while transient features can reveal technique changes, fatigue, or instability.

From a **modelling** perspective, the sensor data are linked to meaningful performance descriptors through calibration models, sensor fusion, and data-driven estimation. The general modelling problem can be expressed as

$$y = f(x, \theta) + \epsilon,$$

where x represents multimodal sensor measurements, θ denotes unknown system or athlete-specific parameters, y represents target performance quantities, and ϵ accounts for noise and uncertainty. The framework supports regression and classification tasks for estimating quantities such as stroke rhythm, propulsion symmetry, contact timing, cycle consistency, and technique phase. The emphasis is on robust models that remain reliable under outdoor disturbances and athlete-to-athlete variability.

The **industry relevance** of the work lies in its potential to support next-generation sports equipment, athlete monitoring systems, and intelligent training feedback tools. Instrumented wheelchairs and ski poles could provide manufacturers, coaches, and sports technology companies with practical pathways toward performance-oriented product development, personalized training analytics, and field-deployable sensing solutions. Overall, the work positions smart sports equipment as a bridge between mechatronic sensing, applied modelling, and real-world athletic performance analysis.