5th IMA Conference on Inverse Problems from Theory to Application

Abstracts

8th – 10th September 2026

INSTITUTE OF
MATHEMATICS &
ITS APPLICATIONS

Contents

Keynote Audrey Repetti, A preconditioned forward-backward algorithm for minimizing nonconvex composite functions: Application to radio-astronomical imaging

Over the last decade, non-convex and non-smooth optimisation have been instrumental tools to develop algorithms for solving high-dimensional inverse imaging problems. In particular, the forward-backward (FB) algorithm that aims to minimise a sum of a differentiable function and a non-smooth function, has shown to be particularly flexible and robust. In this presentation, we will focus on a preconditioned version of the FB algorithm for solving the challenging case where the non-smooth function is non-convex and results from the composition between a strictly increasing, concave, differentiable function and a convex non-smooth function. The proposed algorithm circumvents the explicit,and often challenging, computation of the proximity operator of the composite functions through a majorize-minimize approach. Precisely, each composite function is majorised using a linear approximation of the differentiable function, which allows one to apply the proximity step only to the convex nons-mooth function. We show that the proposed approach is a generalisation of reweighting methods, with convergence guarantees. We further show that this algorithm is particularly suitable in the context of radio-interferometric imaging in astronomy. This modality aims to provide high-dynamic range images of the sky by processing terabyte-sized data obtained from radio-telescopes. The resulting problem appears to be a great playground for nonconvex non-smooth optimisation : it necessitates the use of reweighting methods to achieve a high dynamic range, and the preconditioning enables to mainly work on image-size rather data-size. Based on this, an imaging software dubbed Africanus has been developed to handle radio-astronomical data from modern telescopes. We will show a validation of the method on terabyte-sized data from the MeerKAT telescope.

Adaptive and Efficient Residual Whiteness Principle via ADMM and Proximal Mappings

This talk introduces the new Proximal Residual Whiteness Principle (PRWP), a highly efficient adaptive strategy to choose the regularisation parameter in generic variational regularisation methods for inverse problems in computational imaging. Assuming additive white Gaussian noise, the standard Residual Whiteness Principle (RWP) offers a robust automatic strategy to select an appropriate regularisation parameter by minimising a measure of non-whiteness (such as generalised kurtosis) in the residual image, typically requiring the usually computationally expensive repeated solution of the same regularised problem for a potentially large range of regularisation parameter values. PRWP overcomes this computational bottleneck by embedding the RWP criterion within iterative optimisation methods for the considered regularised formulation ((outer) iterations), and by reformulating the kurtosis minimisation as a constrained optimisation problem, which is then handled by ADMM (inner step). The introduction of a novel strategy to compute the proximity operator of generalised kurtosis measures, coupled with inexact updates in the ADMM solver, makes the resulting scheme computationally convenient, especially when the discretised forward operator lacks any exploitable structure. The stability, robustness, and general applicability of the new PRWP is assessed across large-scale imaging tasks, including denoising, sparse deconvolution, X-ray computed tomography (CT) reconstruction, and seismic imaging.

The SCD semismooth* Newton method for the efficient minimization of Tikhonov functionals

Simon Hubmer, Helmut Gfrerer, Ronny Ramlau

We consider the efficient numerical minimization of Tikhonov functionals with nonlinear operators and non-smooth and non-convex penalty terms, which appear e.g. in variational regularization. For this, we consider a new class of SCD semismooth* Newton methods, which are based on a novel concept of graphical derivatives, and exhibit locally superlinear convergence. We present a detailed description of these methods, and provide explicit algorithms in the case of sparsity (ℓ_p, $0 \leq p < \infty$) and TV penalty terms. The numerical performance of these methods is then illustrated on a number of tomographic imaging problems.

Weighted Total Variation Regularization for Inverse Problems

Ole Løseth Elvetun

Based on joint work with Bjørn Fredrik Nielsen and Martin Burger

Total variation (TV) regularization is a natural choice for inverse problems whose solutions are spatially extended and approximately piecewise constant. However, when the forward operator has a large nullspace—as in inverse source problems arising, e.g., in inverse ECG and EEG problems – standard TV produces systematic artifacts: reconstructed sources are biased toward the observation boundary, and their location and size are poorly recovered.

In this talk, we present weighted TV regularization schemes designed to counteract this bias. The spatially varying weights are derived from a sensitivity analysis of the forward operator, expressed in terms of an associated Green's function, and rebalance the TV penalty to compensate for the operator's inhomogeneous sensitivity. We introduce two related models: an isotropic weighted TV functional, in which the weight reflects worst-case directional sensitivity, and an anisotropic refinement, in which the weighting depends on the actual jump direction of the solution. Both variational-type and basis-pursuit formulations are considered. We establish some exact recovery results under suitable structural assumptions, and we give an interpretation on why the constructed weights remove spatial bias.

Numerical experiments for inverse source problems demonstrate that the weighted formulations substantially reduce the artifacts of unweighted TV, yielding markedly improved localization and size recovery. We also present simulations for a hybrid weighted-sparsity/TV approach that captures both small and extended sources.

An Arnoldi-Tikhonov Method for the Efficient Regularization of Linear Ill-Posed Problems

Ronny Ramlau, Industrial Mathematics Institute
Johannes Kepler, University Linz
Johann Radon Institute for Computational and Applied Mathematics (RICAM)

Email:
ronny.ramlau@jku.at
ronny.ramlau@ricam.oeaw.ac.at

We consider the linear operator equation $Ax = y$, where $A: X \to Y$, X, Y are Hilbert spaces and A is a compact operator. In order to compute a solution of this ill-posed equation, regularization methods are commonly used. In particular, Tikhonov regularization is very popular. For a numerical realization, A has to be discretized which might lead to the task of computing a solution of a large scale linear system. We propose to use the Arnoldi method for decomposing the matrix into $AV_m = V_{m+1} \bar{H}_{m+1,m}$ with a rather low dimensional matrix $\bar{H}_{m+1,m}$, which is consecutively used to compute the minimizer of the Tikhonov functional, resulting in a highly efficient method. We present both an analysis of the method and numerical examples.

Keynote Samuli Siltanen, Tomography, PDEs, and Sustainable Learned Reconstruction

Two examples of severely ill-posed mathematical inverse problems are discussed, both arising from medical imaging. Ill-posedness means that the task of reconstructing the image from data is very sensitive to measurement noise and modelling errors.

The first problem, limited-angle X-ray tomography, is linear. The aim is to recover the X-ray attenuation function from a collection of line integrals, whose orientation is restricted to a small opening angle. This happens for example in 3D mammography. Using theory of microlocal analysis and suitable computational tools including complex wavelets, it is possible to apply targeted deep learning and achieve results of unprecedented quality.

The second example, electrical impedance tomography (EIT), is nonlinear. There one probes the human body with harmless electric currents, measures voltages at the skin, and produces an image of the internal electric conductivity distribution. EIT works wonders in lung imaging and shows promise in stroke monitoring and detection of cancer. A recent breakthrough study showed that the EIT reconstruction process can be decomposed into parts where most of the ill-posedness is confined into deblurring a linear sinogram (Radon transform). This in turn paves the way for targeted deep learning where only the sinogram-cleaning is learned and all else done analytically.

The overall message of the presentation is this. Ill-posed inverse problems are tough to solve even for modern artificial intelligence, requiring large neural network models and lots of training data. This talk presents approaches where traditional mathematics of inverse problems is used for extracting robust information that are in turn used as inputs for machine learning. Such a combination allows harnessing the power of AI in a more sustainable way, with small networks and reasonable amounts of training data.

QVaR: a Quantum Variational Regularization method for Linear Inverse Problems

Co-authors: Siiri Rautio, Hjørdis Schlüter, Babak Maboudi Afkham

We present a tailored framework for solving regularized linear inverse problems using quantum optimization methods. By discretizing the solution space and encoding data fidelity and regularization terms into quadratic unconstrained binary optimization (QUBO) models, we formulate both Tikhonov- and sparsity-promoting regularized inverse problems within a unified quantum optimization framework. We further introduce a notion of quantum sensitivity that characterizes the effect of perturbations arising from approximate quantum evolution and discretization. We derive bounds relating these perturbations to stability properties of the underlying inverse problem, thereby establishing a theoretical connection between quantum solution variability and classical notions of ill-posedness. The framework is extended to variational inverse problems through wavelet-based representations and complemented by reduced-order modeling strategies in both parameter and Hamiltonian spaces to mitigate current hardware limitations. Numerical experiments on simulated and physical quantum hardware demonstrate agreement between the theoretical sensitivity analysis and the observed distributions of low-energy solutions, illustrating both the potential and the present limitations of quantum optimization for inverse problems.

Radiotherapy Planning as an Inverse Problem: Methods, Models, and Emerging Directions

Joana Matos Dias

INESC Coimbra, University of Coimbra, Portugal

joana@fe.uc.pt

Radiotherapy treatment planning is a fundamental step in cancer treatment, where radiation beams are used to deliver a prescribed radiation dose to tumour regions while minimizing exposure to surrounding healthy tissues and organs-at-risk. From a mathematical perspective, this process is naturally formulated as an inverse optimization problem: the desired clinical outcome is known in advance through dose–volume criteria, but the beam configurations and intensity patterns capable of achieving the clinical prescription must be determined.

This inverse planning problem gives rise to a highly constrained, nonlinear, and frequently nonconvex optimization problem, further complicated by the conflicting nature of clinical objectives (preserving normal tissue but, at the same time, properly irradiating the volumes to treat). Over the years, our research team has explored several optimization methodologies to address these challenges.

Initial approaches relied on quadratic optimization formulations, where deviations from prescribed dose distributions are penalized. Although these models remain widely used in clinical practice due to their computational efficiency, they typically require extensive trial-and-error parameter tuning to obtain clinically acceptable solutions. To overcome these limitations, we investigated heuristic and metaheuristic optimization methods, particularly suited to exploring large and nonconvex search spaces arising from beam angle optimization and other combinatorial aspects of treatment planning. We also incorporated fuzzy inference systems to avoid human trial and error parameter tuning. More recently, machine learning models have been integrated into the optimization framework, functioning as objective functions and enabling patient-specific informed planning strategies, also accounting for more than dose-volume objectives. These developments reflect an important transition towards approaches that combine optimization theory, artificial intelligence, and clinical expertise.

Radiotherapy treatment planning exemplifies how inverse optimization problems can effectively bring together mathematical modeling, advanced computational algorithms, and impactful real-world clinical applications.

Analytical Inversion of Modulo Radon Transform

M. Beckmann

University of Hamburg
Germany

In the recent years, the topic of high dynamic range (HDR) X-ray imaging and tomography has started to gather attention due to advances in hardware technology. The issue is that registering high-intensity projections that exceed the dynamic range of the detector cause sensor saturation, which, in turn, leads to a loss of information. Inspired by the multi-exposure fusion strategy in computational photography, a common approach is to acquire multiple Radon projections at different exposure levels that are algorithmically fused to facilitate HDR reconstructions.

As opposed to this, in our current line of work, a single-shot approach to HDR tomography has been proposed based on a co-design of hardware and algorithms and modelled by the so-called Modulo Radon Transform (MRT), a novel generalization of the conventional Radon transform (RT). In the hardware step, RT projections are folded via a modulo non-linearity, which allows HDR values to be mapped into the dynamic range of the sensor and, thus, avoids saturation. Thereon, recovery is performed by algorithmically inverting the modulo folding so that the MRT measurements are mapped back to the original RT range.

While current MRT reconstruction approaches typically require a low-pass pre-filtering, this talk presents a novel analytical inversion formula for the MRT avoiding bandlimitedness assumptions. It is based on the solution of a Poisson problem linking the Laplacian of the RT of a function to its MRT in combination with the classical filtered back projection formula for inverting the RT. Discretizing the inversion formula using Fourier techniques leads to our novel Laplacian Modulo Unfolding -- Filtered Back Projection algorithm, in short LMU-FBP, to recover a function from fully discrete MRT data. Our theoretical findings are finally supported by numerical experiments showing the effectiveness of our approach.

Bayesian Inversion for a Hamilton-Jacobi Equation with a Star-Shaped Prior

Viet Ha Hoang

We study the Bayesian inverse problem for the Hamilton-Jacobi equation $H(\nabla T(x)) = s(x)$ where the function $s(x)$ is binary. The domain contains an unknown star-shaped inclusion which we wish to infer. The prior setting assumes that the inclusion's centre point is uniformly distributed in the domain, and the radial function follows a log-normal model. We show the existence and well posedness of the Bayesian inverse problem by employing approximations of the forward equation's viscosity solution using sup-convolution of the discontinuous function $s(x)$. Numerical approximation of the posterior measure is studied. We quantify how the forward solver's convergence rate depends on the log-normal radial function. This implies an explicit rate of convergence for the posterior's numerical approximation. The results are used to determine optimal sampling parameters for Multilevel Markov Chain Monte Carlo approximations of expectations with respect to the posterior measure for quantities of interest. Numerical experiments verify the theoretical convergence rates and show that the ground truth for our quantities of interest, which are the solution to the forward equation or the inclusion, can be accurately recovered.

Keynote Roland Potthast, Physics and AI for State Estimation

How can physics and AI be brought into coherence in inverse problems? This talk addresses this question in the context of remote sensing, data assimilation, and numerical weather prediction. Our central thesis is that AI becomes most effective not when it replaces physics, but when it is designed to work coherently with physical structure, physical operators, and physically motivated optimization principles.

We first illustrate the role of physics in AI through the emulation and improvement of observation operators. Recent work on data-driven surrogate and correction models for satellite observations shows how machine learning can emulate radiative transfer and, even more importantly, how hybrid residual models can improve established operators such as RTTOV while keeping physics at their core (Buono et al., 2026). In this way, AI is used not as a black-box replacement, but as a structured extension of physical modeling that improves accuracy in challenging regimes while preserving interpretability and operational relevance.

We then turn to the role of physics in the loss function and optimization principles underlying inversion for state estimation. Here, the focus is on variational and ensemble-based data assimilation, where physical consistency enters through the formulation of the estimation problem itself. This perspective connects to recent work on AI-VAR, where neural networks learn analysis mappings inspired by variational data assimilation (Keller and Potthast, 2024), and to recent work on the AI particle filter, aimed at extending AI-based inversion to nonlinear and non-Gaussian settings. Together, these developments point toward a unified view of inversion in which physics enters both through the forward operators and through the objective functions, constraints, and uncertainty representations that guide learning.

The talk argues that the future of state estimation lies in this coherent integration of physics and AI: combining physical insight, learned operators, and uncertainty-aware inference into methods that are both powerful and trustworthy.

Generative AI for Almost-Nearly-Perfect Bayesian Computational Imaging

Marcelo Pereyra

This talk introduces a novel mathematical and computational framework for Bayesian image and video restoration that leverages distilled generative denoising diffusion models as powerful and flexible image/video priors. Our approach is built on two key innovations.

First, we develop a new plug-and-play paradigm for embedding distilled generative models within Langevin diffusion processes, enabling the formulation of generative restoration techniques with clear, interpretable components that separately model statistical image priors and data sensing. This structure allows instrumental and degradation models to be specified at inference time and seamlessly integrated with pre-trained generative priors such as modern Stable Diffusion models.

The second innovation combines Langevin sampling with deep unfolding to design new neural network architectures tailored for Bayesian imaging. When coupled with modern adversarial distillation techniques, this leads to highly efficient algorithms that achieve state-of-the-art restoration performance with as few as 4 to 8 neural function evaluations.

We validate our framework through comprehensive experiments, demonstrating that it not only matches or exceeds the accuracy of current generative restoration methods but also offers significant improvements in computational efficiency and adaptability across diverse sensing scenarios.

RELATED WORKS

The material for this talk is drawn from recent work on generative AI for image restoration, exploring distilled

latent Stable Diffusion XL image-text models in zero-shot settings, video extensions, and the integration of deep

unfolding with diffusion distillation techniques.

https://arxiv.org/abs/2503.12615

https://arxiv.org/abs/2510.01339

https://arxiv.org/pdf/2507.02686

Approximating interpretable posterior correlation structure in inverse problems

Teo Deveney

Data fitting under prior regularisation has become one of the prominent approaches to inverse problems. In the statistical community this is typically posed as a Bayesian inference procedure where we seek to identify properties of the posterior distribution given some prior distribution and observation likelihood. The most common target is the posterior mode, as this coincides with the solution to the variational regularisation problem widely studied in optimisation literature. However, seeking further information from the posterior can be tricky as explicit modelling of high dimensional probability distributions is prohibitively expensive. This typically restricts further statistical approaches to implicit inference techniques (such as sampling), or explicit inference with structure breaking assumptions (for example assuming that the pixels in an image are independent). In this work we describe a tractable correlated uncertainty approach to provide richer explicit posterior approximations in image reconstruction tasks. We use sparse Gaussian parameterisations over image distributions such that correlations between pixels are captured. These models become the basis of our posterior estimates, resulting in realistic correlations over localised structures such as texture or edges. In contrast to sampling techniques our model returns an analytic posterior estimate which allows us to quantify and plot correlation structure in a visually interpretable way (for example the collapse of the nullspace uncertainty can be seen with increasing data). Posterior statistics such as pixelwise variance or inter-pixel correlations can also be read off the resulting distribution. Moreover, the computational cost of reconstructing this posterior approximation is comparable to the equivalent MAP optimisation problem, making it a practical addition to existing approaches.

Accelerating diffusion posterior sampling in large-scale Bayesian inverse problems

Duc-Lam Duong

University of Oulu, Finland

duc-lam.duong@oulu.fi

Abstract: Diffusion models have emerged as a powerful framework for solving Bayesian inverse problems, but applying them to large-scale, infinite-dimensional settings remains computationally challenging. Standard diffusion posterior sampling methods require evaluating the forward model—or its adjoint—at every single reverse time step, leading to severe computational bottlenecks in large-scale inverse problems. In this talk, I will present our recent work in an attempt to overcome this limitation by introducing an unconditional representation of the conditional score (UCoS) for linear inverse problems. We show that the conditional score function can be exactly decomposed into an unconditional score and a deterministically updated affine term. This key theoretical insight allows us to completely shift the evaluations of the forward model to an offline training phase. During the online sampling stage, our method generates posterior samples without multiple evaluations of the forward operator, significantly accelerating the inference process. Furthermore, because our formulation is derived directly in infinite dimensions (Hilbert spaces), the resulting algorithms are inherently discretization-invariant and scale robustly to high-dimensional imaging and functional data. We demonstrate the efficiency and accuracy of our approach through numerical experiments with CT, deblurring and diffuse optical tomography. This talk is based mostly on the papers arXiv:2405.15643 and arXiv:2602.03449.

Beyond Image Priors: Gradient Step Denoisers on the Kernel Domain for Blind Inverse Problems

Andrea Sebastiani1

1 University of Genoa, andrea.sebastiani@edu.unige.it

lug-and-play (PnP) methods have become a powerful and flexible framework for solving imaging inverse problems, replacing handcrafted regularization with an off-the-shelf denoiser. Among the available parametrizations, gradient step denoisers are especially attractive since they are built as the gradient descent step of an explicit regularization functional. For this reason they reconcile the empirical strength of learned priors with rigorous convergence guarantees. While such denoisers are now well understood as image priors, far less is known about how to design them on other domains, which carry their own geometry and constraints. In this talk, we focus on gradient step denoisers acting on the kernel domain, motivated by blind deconvolution, where the blurring kernel is unknown and must be estimated jointly with the image. Unlike natural images, blur kernels are low-dimensional and subject to strong structural constraints—nonnegativity, normalization, smoothness, sparsity—so a prior suited to this domain cannot simply be transferred from the image setting. We show how to train a gradient step denoiser that encodes this kernel structure while still corresponding to the gradient of an explicit, learned potential, yielding a convergent learned regularizer dedicated to the kernel. We then embed this kernel-domain denoiser, together with an image-domain counterpart, in an asymmetric block-coordinate PnP scheme for the resulting nonconvex blind inverse problem. After a concise overview of the PnP framework and its connections to variational regularization, we present the theoretical guarantees of the proposed approach and illustrate its performance on representative blind image restoration experiments.

Keynote Julian Tachella, Towards foundation models in computational imaging

Most existing learning-based methods for solving imaging inverse problems can be roughly divided into two classes: iterative algorithms, such as plug-and-play and diffusion methods leveraging pretrained denoisers, and unrolled architectures that are trained end-to-end for specific imaging problems. Iterative methods in the first class are computationally costly and often yield suboptimal reconstruction performance, whereas unrolled architectures are generally problem-specific and require expensive training. In this talk, I will present a novel non-iterative, lightweight architecture that incorporates knowledge about the forward operator (acquisition physics and noise parameters) without relying on unrolling. The model is trained to solve a very wide range of inverse problems, such as deblurring, magnetic resonance imaging, computed tomography, inpainting, and super-resolution, and handles arbitrary image sizes and channels, such as grayscale, complex, and color data. In addition, the model can be easily adapted to unseen inverse problems or datasets with a few fine-tuning steps (up to a few images) in a self-supervised way, without ground-truth references. I will present results in various imaging modalities, from medical imaging to low-photon imaging and microscopy.

Structured importance-weighted inference for uncertainty quantification in high-dimensional deformable image registration

Ivor Simpson

Dense deformable image registration is a severely ill-posed high-dimensional inverse problem: many deformation fields can explain the observed images almost equally well, so reliable solutions require uncertainty quantification rather than point estimates alone. Existing amortised probabilistic registration methods are efficient, but typically rely on restrictive variational approximations that struggle to represent multi-modal posteriors and often yield poorly calibrated uncertainty.

We introduce Structured SIR, an importance-weighted inference framework for dense 3D registration that combines sampled importance resampling with a structured Gaussian proposal distribution. The proposal covariance is parameterised as the sum of a low-rank component and a sparse spatially structured Cholesky precision factor with cross-directional correlations, enabling tractable sampling and density evaluation while capturing both local and long-range spatial dependencies.

We evaluate this method on brain MRI registration, where the posterior distribution spans approximately 1.6 million variables. Structured SIR yields substantially better-calibrated uncertainty than variational baselines and deep variational ensembles, while maintaining competitive average registration accuracy. Crucially, the inferred posterior exhibits coherent multi-modal structure: high-probability samples produce anatomically meaningful corrections beyond the posterior mean, indicating that the method captures plausible alternative registrations rather than unstructured noise. Our experiments further highlight the value of the structured covariance components, yielding substantial gains in best-sample accuracy over comparator methods. Inference remains computationally efficient: proposal generation, weighting of 1,200 samples, and posterior characterisation are completed in under five seconds on a single GPU. We also observe stable learning from a cold start. These results position Structured SIR as a practical framework for scalable, expressive posterior characterisation and uncertainty quantification in high-dimensional imaging inverse problems.

Neural Network-Based Motion-Aware Reconstruction of Dynamic MRI Data

Martin Holler[1], Matthias Höfler[1], Hendrik Kleikamp[1], and Martin Uecker[2]

[1)] IDea_Lab
University of Graz
Leechgasse 34, 8010 Graz, Austria
martin.holler@uni-graz.at
matthias.hoefler@uni-graz.at
hendrik.kleikamp@uni-graz.at
web page: https://amml-research.at

[2)] Institute of Biomedical Imaging
Graz University of Technology
Stremayrgasse 16/III, 8010 Graz, Austria
e-mail: uecker@tugraz.at

Magnetic resonance imaging (MRI) has become a standard diagnostic tool in medical assessments. However, long acquisition times remain one of its main drawbacks. This makes dynamic imaging, as needed in cineMRI, particularly challenging: motion artefacts and limited data acquisition per time point lead to an inherently ill-posed reconstruction problem. Typical solution strategies include breath-holding by the patient to limit induced motion or manual binning of frames according to an external signal.

We propose a method that incorporates motion signal information directly in the image discretisation. Inspired by the deep-image prior [1], invertible residual networks [2], Lipschitz-constrained networks via Householder reflections, and previous work on motion disentanglement [3], we design a multi-component neural network-based architecture for image reconstruction. The motion signals, in this case representative of the breathing and cardiac cycle, are used as conditioning information for the network. This framework not only allows for retrospective modification and suppression of motion but also enables the reconstruction of incomplete motion information, such as when no breathing belt is used.

In our presentation, we provide a sound formulation of the reconstruction problem, together with general features of the proposed network architecture. These are complemented by numerical test cases highlighting capabilities such as cine reconstruction, retrospective motion suppression, feature tracking, displacement extraction, and resolution independence.

[1] Ulyanov, D., Vedaldi, A., Lempitsky, V., *Deep image prior*, Proceedings of the IEEE conference on computer vision and pattern recognition, pp. 9446-9454, 2018.

[2] Behrmann, J., Grathwohl, W., Chen, R. T., Duvenaud, D., Jacobsen, J. H., *Invertible residual networks, International conference on machine learning*, pp. 573-582, 2019.

[3] Abdullah, A., Holler, M., Kunisch, K., Landman, M. S., *Latent-space disentanglement with untrained generator networks for the isolation of different motion types in video data*, International conference on scale space and variational methods in computer vision, pp. 326-338, 2023.

Learned regularisation and uncertainty quantification for image deconvolution in fluorescence microscopy

Sofia Agostoni

MalGa, DIBRIS, University of Genoa, Italy

Image Scanning Microscopy (ISM) is a modern fluorescence microscopy technique that enhances spatial resolution by combining multiple detector measurements, making it particularly well-suited to low-photon imaging conditions. We address the associated reconstruction task by formulating it as a Poisson inverse problem and adopting a Bayesian framework to recover the underlying image from degraded observations.

To go beyond model-based design of image regularisation, we learn the prior directly from data. This is achieved through a Plug-and-Play (PnP) approach, in which a deep neural network denoiser, trained on representative microscopy images, is used as an implicit prior within an iterative reconstruction scheme. In this setting, the denoiser acts as a proximal operator associated with a learned regulariser, and the overall algorithm can be interpreted as computing a MAP estimate under this data-driven prior. This allows the reconstruction to leverage complex and realistic image statistics that hand-crafted models such as total variation or Gaussian priors fail to capture.

While the PnP approach provides high-quality point estimates, it does not characterise the full posterior distribution. To address this limitation, we further investigate sampling-based methods that go beyond MAP estimation and enable full posterior inference. Specifically, we consider algorithms grounded in discretizations of the overdamped Langevin diffusion, a stochastic process whose stationary distribution coincides with the target posterior. Within this family, we explore the PnP - Unadjusted Langevin Algorithm (PnP -ULA). By generating samples from the posterior, these methods allow the computation of the MMSE estimate as well as uncertainty quantification, providing a richer and more principled characterisation of the reconstruction.

Solving Inverse Problems with Flow-based Models via Model Predictive Control

Alexander Denker

Flow-based generative models provide strong unconditional priors for inverse problems, but guiding their dynamics for conditional generation remains challenging.

Recent work casts training-free conditional generation in flow models as an optimal control problem; however, solving the resulting trajectory optimisation is computationally and memory intensive, requiring differentiation through the flow dynamics or adjoint solves.

We propose MPC-Flow, a model predictive control framework that formulates inverse problem solving with flow-based generative models as a sequence of control sub-problems, enabling practical optimal control-based guidance at inference time.

We provide theoretical analysis linking MPC-Flow to the underlying optimal control objective and show how different algorithmic choices yield a spectrum of guidance algorithms, including regimes that avoid backpropagation through the generative model trajectory.

We evaluate MPC-Flow on benchmark image restoration tasks, spanning linear and non-linear settings such as in-painting, deblurring, and super-resolution, and demonstrate strong performance and scalability to massive state-of-the-art architectures via training-free guidance of FLUX.2 (32B) in a quantised setting on consumer hardware.

Keynote Serena Morigi, Model-Based Geometric Deep Learning for Nonlinear Inverse Problems: Integrating Forward Models and Mesh-Based Regularization

serena.morigi@unibo.it

The reconstruction of unknown parameters from indirect and noisy measurements in nonlinear inverse problems, such as Electrical Impedance Tomography (EIT), requires a delicate balance between physical consistency and structural prior knowledge. Traditional variational methods, while mathematically rigorous, often struggle to adapt to the complex morphological priors required by real-world applications.

In this talk, we present a hybrid framework that bridges classical variational optimization and Geometric Deep Learning ensuring that the learning process is strictly guided by the underlying physical model. We explore the integration of Graph Neural Networks (GNNs) as learned proximal operators acting directly on irregular triangular meshes, and we analyze how informative priors can be utilized to guide data sparsity and stabilize the inversion of ill-posed operators.

By operating on the native geometric domain, this framework allows the model to preserve the underlying physics while enforcing crucial structural constraints. Our results demonstrate how embedding the geometry of the problem through a 'physics-informed' inductive bias enables high-fidelity reconstructions even in data-scarce scenarios, effectively overcoming the limitations of purely data-driven black-box models or purely variational methods.

Two-dimensional strain field imaging by anisotropic electrical impedance tomography

Mikko Räsänen

Electrical impedance tomography (EIT) imaging of anisotropic conductivities is challenging compared to scalar conductivity imaging due to the increased number of unknown parameters, and in addition, due to anisotropic conductivities related by a boundary-fixing diffeomorphism producing the same EIT data. So far, studies on reconstructing anisotropic conductivities in either two or three dimensions using real data are limited in number. However, successful application of anisotropic EIT imaging has potential uses in the imaging of mechanical strain fields based on EIT data, and also in the imaging of materials that are naturally anisotropic.

In this talk, we present results of imaging anisotropic conductivity changes caused by planar strain in a two-dimensional setup using real data and the Bayesian approach to inverse problems. The EIT reconstructions represent qualitative features of the components of the strain field, and the results are compared with strain field reconstructions obtained directly using Digital Image Correlation (DIC).

This is joint work with Moe Pourghaz (North Carolina State University) and Aku Ursin (University of Eastern Finland).

Recovering Dipoles in EEG Using Weighted Group LASSO

Ole Løseth Elvetun, Bjørn Fredrik Nielsen and Niranjana Sudheer
Norwegian University of Life Sciences

A classical inverse problem in EEG is to recover one or a small number of dipole sources from voltage recordings on the scalp. This problem is naturally suited to sparsity-based techniques. However, each dipole has both a spatial location and an orientation, requiring three “basis” dipoles at each position to represent a freely oriented source. This complicates the application of standard sparsity regularization. In addition, without appropriate weighting, the solution is affected by a strong depth bias, where deeper sources are misinterpreted as superficial ones.

To address these challenges, we propose a weighted group sparsity approach, also known as the weighted group LASSO, in which each group comprises the three basis dipoles at a given position. This formulation yields a rotation-invariant penalty at each location. We show that, using this approach, any single dipole can be exactly recovered independently of its orientation. Furthermore, under a suitable disjointness condition, multiple dipoles can also be recovered; this condition effectively requires that the sources be sufficiently well separated in space. Finally, we illustrate the theoretical results with several numerical experiments.

Hidden Heat-Sources Estimation with a Deep Kalman Filter

Marco Dell'orto and Fabio Marcuzzi

This presentation deals with identifying internal heat sources inside a solid object, from knowledge of boundary temperature measurements. We assume the distributed heat sources have an unknown location inside the body and their intensity is time-varying. Some applications could be the identification of internal cavities or material inhomogeneities [1], hidden corroded surfaces, etc. It is an effective tool for performing diagnostic tasks or detecting material changes through indirect measurements, useful e.g. for non-destructive testing or embedded digital twins.

To estimate heat-sources from boundary temperature measurements is a inverse problem, that here will be conducted inside a Data Assimilation procedure enhanced by Deep Learning (aka, the Deep Kalman Filter - DKF [2]) .

In this perspective, it can be seen as an inverse problem involving a dynamical system where the model is only partially known in a Markovian form, and requires a non-Markovian closure [3]. We will analyze three possible closures of this model, without a prior model of heat-sources dynamics:

- given a (space-)discretization of the heat source, we associate to each discretization-node a parameter of the predictor model set in the DKF;
- with additional input variables, time-varying, representing the heat-source discretization, a point-source response model and using the maximum principle for heat equation;
- with a Recurrent Neural Network (RNN) and following the Mori-Zwanzig formalism [3].

[1] Giusteri, G.G., Marcuzzi, F., Rinaldi, L.: Replacing voids and localized parameter changes with fictitious forcing terms in boundary-value problems. Results in Applied Mathematics 20 (2023). https://doi.org/10.1016/j.rinam.2023.100402

[2] Chinellato, E., Marcuzzi, F.: State, parameters and hidden dynamics estimation with the Deep Kalman Filter: regularization strategies. Journal of Computational Science 87, 102569 (2025). https://doi.org/https://doi.org/10.1007/978-3-031-63775-9

[3] Levine, M.E., Stuart, A.M.: A framework for machine learning of model error in dynamical systems, Communications of the American Mathematical Society, Volume 2, Issue 07, (2022), DOI 10.1090/cams/10

A self-supervised approach for quantitative lifetime estimation in fluorescence microscopy for semiconductor and photovoltaic materials

Gabriele Scrivanti[1], joint work with Luca Calatroni[1] and Stefania Cacovich[2]

[1] MaLGa Center, DIBRIS, Università di Genova - Genova, Italy
[2] IPVF IPVF, UMR 9006, CNRS, Ecole Polytechnique,IP Paris, Chimie Paristech, PSL - Palaiseau, France

We propose an unsupervised deep-learning framework for the reconstruction of carrier lifetime maps from noisy time-resolved fluorescence imaging (TR-FLIM) measurements, with applications to semiconductor and photovoltaic device characterisation. The problem is formulated as a nonlinear inverse problem, where the goal is to recover spatially distributed lifetime parameters from a series of temporally resolved, photon-count-limited observations corrupted by a challenging combination of noise sources whose joint statistical modelling is computationally intractable. Accurate recovery of these parameters is critical for predicting solar cell efficiency and identifying early signatures of material degradation.

Our approach builds on the Noise2Noise (N2N) statistical learning framework, which exploits the theoretical result that a model trained to map one noisy realisation of a signal to another independent noisy realisation of the same signal converges, in expectation, to the underlying clean signal without ever requiring access to ground-truth data. We incorporate a physics-driven log-linear decay model into the N2N loss function, reflecting the known exponential decay structure of the emitted photoluminescence and enabling the simultaneous estimation of spatially varying amplitude and lifetime maps. The reconstruction is performed in a zero-shot regime, meaning no pre-training on external datasets is required: optimisation is carried out directly on the set of experimental noisy observations available for a single scene. The unknown parameter maps are parameterised as the outputs of a neural network, whose design serves as an implicit regulariser on the solution space.

This setting is particularly well suited to low signal-to-noise ratio regimes arising from short acquisition protocols, which minimise potential damage caused by prolonged exposure to the laser excitation source. The method is further demonstrated to be effective for the monitoring of lifetime evolution under controlled degradation conditions, including thermal stress and humidity exposure, offering a rigorous and computationally efficient tool for materials characterisation.

Accessibility and Usability of Model-Based Iterative Reconstruction Algorithms at the Science and Technology Facilities Council

Margaret Duff – margaret.duff@stfc.ac.uk

The Science and Technology Facilities Council (STFC) is a UK research council that funds research across particle physics, astronomy, and nuclear physics and, perhaps less well known, funds and delivers the National Facilities: the UK's large-scale physical science infrastructure. These include Diamond Light Source (an X-ray synchrotron), the ISIS Neutron and Muon Source, the Central Laser Facility, the Boulby Underground Laboratory, and the National Satellite Test Facility, among others. Scientific experiments conducted at these facilities span a wide range of modalities, scales, and applications, generating a rich set of inverse problems in reconstruction and analysis.

These problems are often ill-posed, with data characterised by noise, missing measurements, hyperspectral structure, and non-standard acquisition geometries. Variational regularisation, combined with iterative optimisation methods, provides a flexible framework for incorporating prior knowledge to address this ill-posedness; however, adoption is often hindered by issues of accessibility and usability for large datasets by non-specialists.

This work presents an overview of inverse problem applications at STFC, drawing on recent work across the facilities. We discuss advances aimed at improving usability, including approaches that reduce or eliminate manual regularisation parameter tuning, automated decision-making strategies, and the use of stochastic optimisation algorithms to accelerate convergence in large-scale settings. We also outline the development of open-source software designed to support reproducible and scalable reconstruction workflows for non-expert users. Finally, we consider emerging challenges in integrating deep learning approaches, including hybrid methods that combine physics-based modelling with data-driven techniques, and the associated demands for robustness and interpretability.

This presentation aims to highlight and begin bridging the gap between inverse problems research and practical deployment in real-world scientific environments.

Safeguarded iterative reconstructions for nuclear fuel safeguard inspections

Tommi Heikkilä (tommi.heikkila@lut.fi), LUT University, Finland

Sara Heikkinen, LUT University, Finland

Riina Rimppi, Nuclear Safety Authority (STUK), Finland

Tapio Helin, LUT University, Finland

Passive Gamma Emission Tomography (PGET) is an IAEA-approved technique for verifying spent nuclear fuel assemblies prior to geological disposal. Reconstructing the emission and attenuation maps from PGET measurements is a nonlinear ill-posed inverse problem, currently solved with a Levenberg-Marquardt (LM) scheme that requires 10-20 iterations to achieve sufficient accuracy. We propose an accelerated iterative solver that combines the LM algorithm with a Deep Gauss-Newton step, in which a learned operator refines the update proposed by the deterministic algorithm at each iteration if a trust-region based sadefuard condition is satisfied. This ensures that the accelerated iterates perform no worse than LM and retain convergence to a critical point of the regularized objective function.

Within this setup we compare three architectures for the learned component: an encoder-decoder-style convolutional neural network, Fourier Neural Operators, and Wavelet Neural Operators. Each is trained on a small set of coarsely simulated 9x9 assemblies. Initial experiments on simulated and real measurements from Finnish nuclear power plants suggest that the proposed scheme reaches LM-quality reconstructions faster, and relatively robustly against outliers.

Bi-level regularized inversion with application to data-driven PDEs

Tram Nguyen

We present the novel bi-level regularization methods for ill-posed inverse problems governed by nonlinear time-dependent PDEs. By a bi-level Landweber scheme, one alternates between upper-level parameter reconstruction and lower-level state approximation, bypassing exact PDE solves. We derive adaptive stopping rules for both levels and prove convergence, stability and acceleration of the bi-level algorithms. The bi-level approach illustrates its universality through several reaction-diffusion applications, in which the nonlinear reaction law needs to be determined.

Data-driven projection-based regularization for galaxy classification

Manuel Cañizares

The orbits of stars in a galaxy can reveal information about its past. In particular, orbital structures give us information about previous merger processes between galaxies. Finding these structures from astronomical observations requires computationally heavy dynamical modeling. Furthermore, it is an ill-posed problem, which amplifies measurement noise.

However, orbital structures can be observed for simulated galaxies, in which individual stars can be tracked. By modeling the relationship between measurements and orbit densities as an abstract linear operator, we propose a projection-based method to use these simulated data as a learning set that allows us to regularize the problem.

Furthermore, in this data-driven approach we learn the operator from the training data, which allows us to avoid the dynamical modeling, and resulting in extremely fast reconstructions.

This is joint work with Otmar Scherzer, Ioannis Martikos, Glenn van de Ven, Prashin Jethwa and Stefanie Reiter, from University of Vienna. The approach is based in papers by A. Aspri, Y. Korolev and O. Scherzer, and by M. Hanke and O. Scherzer.

Posters Presented for Day 1 - 8th September 2026	
Sayana K. Jacob	Regularization Methods in the Context of Pseudo-Differential Operators
Matheus Loures	A Fast Boundary Method with Uncertainty Quantification: Method of Fundamental Solutions
Pietro Maurino	Iterated Graph Laplacian for Image Restoration Problems
Letizia Protopapa	Fitbenchmarking: Comparing Minimizers for Inverse Problems
Sebastian Eberle	Uncertainty Quantification for Phase Retrieval with Applications to Near-Field Holography
Yıldız Bozkurt	Layer Adapted Time of Flight Calculation using Interpolation for Medical Ultrasound Imaging
Andreas Chatziafratis	Analytical Inversion Mechanisms for Evolution PDE Via the Fokas Unified Transform Method
Chaima Abid	An Alternating Approach to Reconstruct PROMPT GAMMA Distribution and Hadron Velocity from Time-Of-Flight Measurements
Teresa Klatzer	Latent Sliced Wasserstein Flows for Scientific Inverse Problems
Bryce Shirley	Parallel And High Order Algorithms for Thick Sample Ptychography
Robert Johnson	Mathematical Algorithms for Ptychography Imaging

Posters Presented for Day 2 - 9th September 2026	
Emma Erkocevic	Data-Driven Katsevich Filtered Backprojection
Alexandra Valavanis	Learning Regularisation Parameters Via Fenchel-Young Minimisation
Daniel Burrows	Motion-Enabled Tomography Via Gaussian Mixture Models
Stepan Bogdanov	Machine-Learning-Assisted Inverse Reconstruction of Micro- and Nanoplastics Using Structured-Light Topological Dynamics
Jehan Alswaihli	A Neural Field Theory Framework for EEG Inversion: Connecting Dynamics with Data
Christian Daniele	Noise And Task-Adaptive Deep Equilibrium Mirror-Descent for Poisson Image Reconstruction
Thomas Cheetham	Bayesian Inverse Periodic Navier-Stokes Problems: Applications To 4D Flow MRI
Riccardo Barbano	Solving Inverse Problems with Flow-Based Models Via Model Predictive Control
Xialu Liu	Sparse Factor Models for High Dimensional Time Series
Patrick Fahy	Optimisation-Aware Learned Regularisers for Imaging Inverse Problems
Luca Trautmann	Inferring Forward Operator Distributions in Linear Blind Inverse Problems Via Diffusion Priors
Jose Rodrigo Rojo Garcia	Bayesian Optimal Experimental Design with Wasserstein Information Criteria
Nicolas Goeman	A Hierarchical Likelihood Model for Non-Linear Inverse Problems Under Additive and Multiplicative Noises
Maren Casfor	Model Error Propagation in Bayesian Inversion

Posters Presented for Day 3 - 10th September 2026	
Oscar Humphreys	Gaussian Processes for Probabilistic Seismic Wavefield Simulation
Karoliina Puronhaara	Utilising A Learned Forward Operator in the Inverse Problem of Photoacoustic Tomography
Lena Dunst	Large-Scale Model-Based 3D Image Reconstruction for Raster-Scan Optoacoustic Mesoscopy
Claire Namuroy	Bayesian Inversion and Uncertainty Quantification of RANS Turbulence Models in Flow MRI
Jarjish Rahaman	Spatial And Temporal Modulation in Time-Domain Diffuse Optical Tomography
Santeri Kaupinmäki	Sparse Reconstructions of Sound Speed and Absorption Using Phase-Insensitive Sensors in Ultrasound Tomography
Ivan Hasenohr	Online Parameter Reconstruction for Bilinear Control Systems with Applications To MRI
Anssi Manninen	Towards Robust Experimental Reconstruction Methods for Quantitative Photoacoustic Tomography Via Learned Iterative Methods
Keigo Oyama	Bayesian Model Selection for Detectability Limits Between Sphere and Core–Shell Models Using Synthetic SAXS Data
Davit Gyulamiryan	Application of Diffuse Domain Method to Inverse Navier–Stokes Problem for Flow-MRI Reconstruction
Nazanin Forghaniallahabadi	Regularised Recovery of Latent Market Regimes in Cryptocurrency Returns: An Inverse-Problem Approach
Clara Hawkins	Multi-Modality Imaging for Ptychography and X-Ray Fluorescence
Maciej Kaczorek	Sound Reconstruction from Neural Activity Patterns
Azhir Mahmood	A Spectral Approach to Adversarial Regularisation for Inverse Problems