Proceedings of the International BASP Frontiers Conference 2025

January 26–31, 2025

Villars-sur-Ollon, Switzerland

BASP Frontiers Conference 2025

Chair

Prof. Yves Wiaux, Heriot-Watt University, UK

Steering Committee

Dr. Kazunori Akiyama, Haystack Observatory, MIT, USA
Prof. Florian Knoll, Friedrich-Alexander University Erlangen Nuremberg, Germany
Prof. Philip Schniter, Ohio State University, USA
Prof. Yves Wiaux, Heriot-Watt University, UK

Session Organisers

Computational imaging theory and methods

Dr. Audrey Repetti, Heriot-Watt University, UK
Prof. Jong Chul Ye, Korea Advanced Institute of Science & Technology, South Korea
Prof. Mujdat Cetin, University of Rochester, USA

Astronomical imaging

Dr. Stefan Wijnholds, Astron, The Netherlands
Dr. Aviad Levis, University of Toronto, Canada
Prof. Laurence Perreault-Levasseur, University of Toronto, Canada

Medical imaging

Prof. Julia Schnabel, Technical University of Munich & Helmholtz Munich, Germany
Prof. Andrew Webb, Leiden University Medical Center, The Netherlands
Prof. Daniel Sodickson, New York University Grossman School of Medicine, USA

Logistics

Dr. Arwa Dabbech, Heriot-Watt University, UK
Taylor Chu, Heriot-Watt University, UK
Ayoub Tajja, EPFL, Switzerland
Prof. Jean-Philippe Thiran, EPFL, Switzerland
Prof. Yves Wiaux, Heriot-Watt University, UK

Foreword

Astronomy and biomedical sciences find common challenges in their need to process data into interpretable signals or images. The complexity of the data acquired and processed is constantly increasing, thereby calling for new advances in signal processing methodology. Data coming in larger volumes every day can be multi-modal, multi-spectral, scalar- or tensor-valued, living in high-dimensional geometries, on graphs, etc. The international Biomedical and Astronomical Signal Processing (BASP) Frontiers conference was created to promote synergies between the astronomy, biomedical sciences, and signal processing communities, with a core focus on computational imaging.

It has been 14 years ... Building on the success of the first six events in 2011, 2013, 2015, 2017, 2019, (Covid-cancelled in 2021), and 2023, BASP Frontiers 2025 will gather about 90 participants in the unique environment of the Swiss Alps, and open its floor to intense interdisciplinary discussions on topics ranging from the latest theoretical developments in signal processing and machine learning, to their application for imaging inverse problems. The most fruitful discussions often take place after the sessions themselves, during breakfast, coffee breaks, or dinner. Not to mention the ski slopes, which further promote creativity.

Major Fourier imaging modalities such as aperture synthesis by radio interferometry in astronomy and magnetic resonance imaging in medicine have contributed to define the ethos of BASP Frontiers since its inception. As of today, a variety of other imaging modalities contribute to shaping the landscape of the conference. Interestingly, over the years, we have witnessed both algorithms and researchers "crossing the bridge" between astronomical and medical imaging.

On Behalf of the Steering Committee

Professor Yves Wiaux

Conference Chair

Contents

Part I

Lightning Keynote talks on Modern Perspectives on Imaging

Perspectives on the future of medical imaging

Greg Zaharchuk*
*Stanford University, USA

Abstract Artificial intelligence (AI) – defined here as deep learning and transformer-based computer vision and language models – has been proposed as a revolutionary technology for radiology, possibly more impactful than past revolutions such as new imaging modalities and or the adoption of PACS. However, equally loud are voices that suggest that AI is overhyped and may not meaningfully address the core challenges facing radiology. Medicine is an inherently conservative field, and rapid technology adoption is the exception rather than the norm. In other specialities, changes are often triggered by landmark clinical trials or new guidelines. In radiology, where much of our work occurs upstream of clinical decision-making, linking practice changes to improved patient outcomes is difficult. Without clear evidence of clinical impact, financial incentives will likely drive technology adoption. This is already visible, where products that have made inroads are largely those have led to improved efficiency or cost-effectiveness. Conversely, adoption has been slower for technologies promising improved diagnostic accuracy. In this talk, I will explore what the next decade may hold for AI in radiology, drawing on current applications while also exploring the history of similar medical technologies to better understand the potential forms AI might take and how it will be integrated into practice.

Perspectives on the Past, Present, and Future of Astronomical Imaging

Jean-Luc Starck*
*CEA-Saclay, France

Abstract Radio astronomers pioneered the development of deconvolution techniques to address imaging challenges unique to radio waves, whose long wavelengths create inherently incomplete, noisy, and low-resolution data. To reconstruct this missing information and improve image clarity, methods like CLEAN and Maximum Entropy Method were created. These techniques laid foundational principles for handling data sparsity and nonlinear reconstruction challenges, later benefiting optical and infrared astronomy and even medical imaging. A major leap forward came with the 1990 launch of the Hubble Space Telescope, which suffered from spherical aberration. This optical flaw spurred a wave of innovation in deconvolution, particularly through wavelet-based techniques that allowed multi-scale analysis, critical for addressing both fine and diffuse structures blurred by HST's optics. Wavelet methods quickly became integral in astronomical imaging, allowing scientists to recover details across various spatial scales. In the mid-2000s, the field advanced again with the introduction of compressed sensing, a theoretical framework demonstrating that high-quality images can be reconstructed from fewer measurements if they are sparse in a basis, such as wavelets. This principle, pioneered by researchers like Candés and Donoho, significantly improved the efficiency of data reconstruction across astronomy and beyond. Recently, deep learning has revolutionized astronomical imaging. Convolutional neural networks and generative models offer advanced solutions for image restoration, noise reduction, and feature extraction, which are essential for handling modern astronomy's vast data. While deep learning unlocks powerful new capabilities, it also brings challenges, including interpretability, reliability, and the need for rigorous validation, all critical for trusted scientific discovery.

Machine learned regularisation for inverse problems - the dos and don'ts

Carola-Bibiane Schönlieb*
*University of Cambridge, UK

Abstract Inverse problems are about the reconstruction of an unknown physical quantity from indirect measurements. Most inverse problems of interest are ill-posed and require appropriate mathematical treatment for recovering meaningful solutions. Regularisation is one of the main mechanisms to turn inverse problems into well-posed ones by adding prior information about the unknown quantity to the problem, often in the form of assumed regularity of solutions. Classically, such regularisation approaches are handcrafted. Examples include Tikhonov regularisation, the total variation and several sparsity-promoting regularisers such as the L1 norm of Wavelet coefficients of the solution.

While such handcrafted approaches deliver mathematically and computationally robust solutions to inverse problems, providing a universal approach to their solution, they are also limited by our ability to model solution properties and to realise these regularisation approaches computationally. Recently, a new paradigm has been introduced to the regularisation of inverse problems, which derives regularisation approaches for inverse problems in a data driven way. Here, regularisation is not mathematically modelled in the classical sense, but modelled by highly over-parametrised models, typically deep neural networks, that are adapted to the inverse problems at hand by appropriately selected (and usually plenty of) training data. In this talk, I will review some machine learning based regularisation techniques, present some work on unsupervised and deeply learned weakly convex regularisers and their application to image reconstruction from tomographic measurements, and finish by discussing some open mathematical problems.

Part II

Computational imaging theory and methods

Session 1

On the Interface of Optimization and Deep Learning for Computational Imaging

Organiser: Audrey Repetti

Summary *This session focuses on computational imaging methods at the interface of optimization and deep learning. Recent advances aiming at pairing these two areas has led to significant advancements in image reconstruction and enhancement. Optimization techniques traditionally provided mathematically rigorous solutions for imaging problems, but they often struggled with high-dimensional data and non-linearities. Deep learning, particularly neural networks, has transformed this landscape by learning complex mappings from data, enabling faster and higher expressivity in image solutions. When combined, these approaches harness optimization algorithms to refine and guide the training of deep networks, resulting in robust models that are highly efficient for performing computational imaging tasks.*

On the Interface of Optimization and Deep Learning for Computational Imaging

Audrey Repetti*
*Heriot-Watt University, UK

Abstract Inverse imaging problems aim to find an estimate of an unknown image from degraded measurements. In this context, proximal iterative algorithms have been state-of-the-art methods for several decades. They consist in solving a regularised variational problem including a data-fidelity term, associated with the forward model, and a regularisation term, to incorporate prior information on the target image. However, with the rise of neural networks, proximal algorithms have been evolving toward hybrid methods mixing optimisation and deep learning strategies. On the one hand, these hybrid methods have shown to improve the reconstruction quality of the output solution compared to traditional proximal algorithms. On the other hand, they enable a more robust reconstruction process compared to pure deep learning approaches, as well as a more theoretically sound interpretation of the output results. These observations created a feedback loop for both optimisation and deep learning, that led to extensive research developments in the last few years, in particular for computational imaging. The first part of this presentation will be dedicated to introducing these methods and discussing their advantages and challenges. In the second part of the presentation, a short overview will be given on building neural network architectures by unrolling proximal algorithm iterations. The resulting architectures are often referred to unfolded neural networks. Compared to traditional networks (such as Unet or DnCNN), they have a more theoretically sound interpretation, gain in robustness, and have a lower complexity cost (making them easier to train and cheaper). We will focus on networks grounded on the dual forward-backward and the Chambolle-Pock algorithms to build Gaussian denoisers, showing that accelerated inertial versions of these algorithms enable skip connections, and proposing different learning strategies. Associated article is available at the DOI: 10.1109/TIP.2024.3437219.

True and False Monotone Neural Networks

Jean-Christophe Pesquet*, Audrey Repetti†, Hugues Talbot* and Younes Belkouchi‡
*CentraleSupélec, Université Paris-Saclay, France, †Heriot-Watt University, UK, ‡Gleamer, France

Abstract There has recently been a growing interest in learning monotone operators. Depending on how monotonicity is defined, different approaches can be employed using appropriate penalization losses. The proposed method is particularly effective for solving monotone inclusion problems, which frequently arise in image processing. Proximal algorithms can be utilized to address these challenges, offering a solution even when the Lipschitz constant of the neural network is unknown. Notably, the FBF algorithm provides convergence guarantees as long as the learned operator remains monotone. Building on plug-and-play methodologies, our goal is to apply these newly learned operators to solve non-linear inverse problems. To this end, we first formulate the problem as a variational inclusion problem. Next, we train a monotone neural network to approximate an operator that may not naturally exhibit monotonicity. Using the FBF algorithm, we present simulation examples demonstrating the successful resolution of the non-linear inverse problem.

Data-Driven Spatial Adaptivity for Regularising Inverse Problems

Sebastian Neumayer*
*Chemnitz University of Technology, Germany

Abstract In this talk, I will discuss various ways to introduce spatial adaptivity into filter-based regularisation functionals. With this adaptivity, we are able to cancel out the filter responses to structure. Hence, we can interpret it as boosting the initial regulariser based on the data. A direct question if is we can repeat this process to get even better solutions. If we try this naively, the answer is sadly no for many cases. However, we can instead train the model in such a way that this procedure works, which opens up many interesting relations with other image reconstruction approaches. Our numerical results are on par with other approaches that rely on spatial adaptivity. Underpinning preprint is available on arXiv: 2407.06608.

Deep image regularisation for Poisson inverse problems via mirror descent

Christian Daniele*, **Luca Calatroni***†‡, Silvia Villa◇, and Samuel Vaiter*
*CNRS, France, †UCA, France, ‡Inria, France, ◇Universitá degli studi di Genova, Italy

Abstract We consider the framework of Deep Equilibrium Models for learning the regularisation term in the context of Poisson imaging inverse problems. As an underlying non-convex optimisation scheme, we consider a convergent Mirror Descent algorithm defined in terms of a potential function well suited to the Kullback-Leibler data term. Due to the lack of contractivity of the underlying fixed-point iterations, we resort to an efficient Jacobian-free technique to perform effectively the back-propagation step. Different choices of the deep architecture enforcing regularisation (e.g., ICNN, DnCNN...) are compared and validated on several numerical examples on standard numerical image datasets.

(Deep) machine learning for exoplanet detection in direct imaging at high contrast

Julien Mairal*, Théo Bodrito*, Olivier Flasseur†, Jean Ponce‡, Maud Langlois†, and Anne-Marie Lagrange†
*INRIA, France, †CNRS, France, ‡ENS, France

Abstract Deep learning has revolutionized image processing and is often considered to outperform classical approaches based on accurate modeling of the image formation process. In this presentation, we will discuss the interplay between model-based and learning-based paradigms, and show that hybrid approaches show great promises for scientific imaging, where robustness and uncertainty quantification is important. We will focus on the problem of exoplanet detection from direct high-contrast imaging, which is particularly challenging due to the high contrast between the planet and the star luminosities. In addition to tailored instrumental facilities implementing adaptive optics and coronagraphy, post-processing methods combining several images are needed to attenuate the nuisances corrupting the signals of interest. Most of these post-processing methods build a model of the nuisances from the target observations themselves, resulting in strongly limited detection sensitivity at short angular separations due to the lack of angular diversity. To address this issue, we propose to build the nuisance model from an archive of multiple observations by leveraging supervised deep learning techniques. The proposed approach casts the detection problem as a reconstruction task and captures the structure of the nuisance from two complementary representations of the data. The proposed approach also encompasses statistical modelling of learnable spatial features. The latter is beneficial to improve both the detection sensitivity and the robustness against heterogeneous data. We apply the proposed algorithm to several data sets from the VLT/SPHERE instrument, and demonstrate a superior precision-recall trade-off compared to recent algorithms currently used in practice.

Convergence analysis of Plug-and-Play methods for image inverse problems

Samuel Hurault*
*Institut de Mathématiques de Bordeaux, France

Abstract Plug-and-play methods constitute a class of iterative algorithms for imaging inverse problems where regularization is performed by an off-the-shelf Gaussian denoiser. These methods have demonstrated impressive visual results, particularly when the denoiser is parametrized by deep neural networks. However, the theoretical convergence of PnP methods has yet to be fully established. This talk begins with an overview of the literature on PnP algorithms, followed by the introduction of new convergence results for these methods when paired with specific denoisers, known as Gradient-Step and Proximal denoisers. Finally, in order to address images corrupted by Poisson noise, we will introduce a novel and convergent Bregman version of Plug-and-Play algorithms.

Stochastic Deep Restoration Priors for Imaging Inverse Problems

Yuyang Hu*, **Chicago Park***, Albert Peng*, Weijie Gan*, Peyman Milanfar†, Mauricio Delbracio†, Ulugbek Kamilov*
*Washington University, Saint-Louis, USA, †Google, USA

Abstract Deep neural networks trained as image denoisers are widely used as priors for solving imaging inverse problems. While Gaussian denoising is thought sufficient for learning image priors, we show that priors from deep models pre-trained as more general restoration operators can perform better. We introduce Stochastic deep Restoration Priors (ShaRP), a novel method that leverages an ensemble of such restoration models to regularize inverse problems. ShaRP improves upon methods using Gaussian denoiser priors by better handling structured artifacts and enabling self-supervised training even without fully sampled data. We prove ShaRP minimizes an objective function involving a regularizer derived from the score functions of minimum mean square error (MMSE) restoration operators, and theoretically analyze its convergence. Empirically, ShaRP achieves state-of-the-art performance on tasks such as magnetic resonance imaging reconstruction and single-image super-resolution, surpassing both denoiser- and diffusion-model-based methods without requiring retraining.

LoFi: Scalable Local Image Reconstruction with Implicit Neural Representation

Amirehsan Khorashadizadeh*, Tobias Liaudat†, Jason McEwen‡, and Ivan Dokmanic*
*University of Basel, Switzerland, †CEA, France, ‡University College London, UK

Abstract Neural fields or implicit neural representations (INRs) have attracted significant attention in machine learning and signal processing due to their efficient continuous representation of images and 3D volumes. In this work, we build on INRs and introduce a coordinate-based local processing framework for solving imaging inverse problems, termed LoFi (Local Field). Unlike conventional methods for image reconstruction, LoFi processes local information at each coordinate separately by multi-layer perceptrons (MLPs), recovering the object at that specific coordinate. Similar to INRs, LoFi can recover images at any continuous coordinate, enabling image reconstruction at multiple resolutions. With comparable or better performance than standard CNNs for image reconstruction, LoFi achieves excellent generalization to out-of-distribution data and memory usage almost independent of image resolution. Remarkably, training on 1024×1024 images requires just 3GB of memory—over 20 times less than the memory typically needed by standard CNNs. Additionally, LoFi's local design allows it to train on extremely small datasets with less than 10 samples, without overfitting or the need for regularization or early stopping. Finally, we use LoFi as a denoising prior in a plug-and-play framework for solving general inverse problems to benefit from its continuous image representation and strong generalization. Although trained on low-resolution images, LoFi can be used as a low-dimensional prior to solve inverse problems at any resolution. We validate our framework across a variety of imaging modalities, from low-dose computed tomography to radio interferometric imaging. Underpinning preprint is available on arXiv: 2411.04995v1.

Equivariant plug-and-play image reconstruction

Thomas Moreau*
*Inria, France

Abstract Plug-and-play algorithms constitute a popular framework for solving inverse imaging problems that rely on the implicit definition of an image prior via a denoiser. These algorithms can leverage powerful pre-trained denoisers to solve a wide range of imaging tasks, circumventing the necessity to train models on a per-task basis. Unfortunately, plug-and-play methods often show unstable behaviours, hampering their promise of versatility and leading to suboptimal quality of reconstructed images. In this work, we show that enforcing equivariance to certain groups of transformations (rotations, reflections, and/or translations) on the denoiser strongly improves the stability of the algorithm as well as its reconstruction quality. We provide a theoretical analysis that illustrates the role of equivariance on better performance and stability. We present a simple algorithm that enforces equivariance on any existing denoiser by simply applying a random transformation to the input of the denoiser and the inverse transformation to the output at each iteration of the algorithm. Experiments on multiple imaging modalities and denoising networks show that the equivariant plug-and-play algorithm improves both the reconstruction performance and the stability compared to their non-equivariant counterparts. Underpinning preprint is available on arXiv: 2312.01831.

Variational reconstruction of parameter maps in medical imaging and materials science

Gabriele Scrivanti*
*CNRS, France

Abstract Variational image reconstruction is a powerful tool for recovering high-quality images from degraded data. By formulating the reconstruction task as a variational optimisation problem, this approach leverages prior knowledge about the image and the degradation process to produce detailed and accurate results. This framework can be extended to the reconstruction of hidden parameter maps, which represent spatially varying properties of an image or a physical system. These maps can encode essential information for various applications. For instance, in medical imaging, parameter maps can aid in tissue identification in ultrasound images. In materials science, they can provide characterising information, such as the carrier lifetime of semiconductors. We present two original contributions in these fields. First, we start from the well-established assumption that the class of generalised Gaussian (GG) distributions effectively describes ultrasound images, with different regions and tissues viewed as realisations of random variables following GG distributions. We then propose a novel model that jointly reconstructs an image from blurry observations and estimates its GG parameters. This enables the segmentation of different tissues or regions within the image. Second, we employ a simple mono-exponential decay model to describe the time-resolved photo-luminescence intensity of photovoltaic materials. By using the so-called "structure tensor total variation" regularisation, we extract spatially correlated estimates of the unknown parameters, including the carrier lifetime. Future research could explore the integration of these models with self-supervised learning techniques (not requiring the use of ground truth data) to further enhance the model's ability to capture complex image patterns and improve the performance of parameter estimation techniques. Associated articles are available at the DOIs: 10.1007/s10851-024-01184-z, 10.1002/adfm.202402343.

Session 2

Foundation and Multimodal Models in Computational Imaging

Organiser: Jong Chul Ye

Summary *Recently, foundation models, which are large scale models pre-trained on massive amounts of diverse data in a self-supervised manner, have been quickly replacing the existing CNN based end-to-end and/or supervised-learning approaches due to their superior performance. As a result of large-scale self-supervised pre-training, foundation models have excellent generalization capabilities with intriguing emergent properties. This session targets recent advances in foundation models for computational imaging, such as text+image models, as well as diffusion models for inverse problems, to understand their advantages and opportunities.*

Diffusion Foundation Models for Inverse Problems

Jong Chul Ye*
*Korea Advanced Institute of Science & Technology, South Korea

Abstract The recent advent of diffusion models has led to significant progress in solving inverse problems, leveraging these models as effective generative priors. Nonetheless, challenges related to the ill-posed nature of such problems remain, such as 3D and video extension and overcoming inherent ambiguities in measurements. In this talk, we will review the recent theoretical and algorithmic advances in diffusion models for inverse problems.

Datasets Design for Improved and Reliable Deep Learning-Based Imaging

Reinhard Heckel*
*Technical University of Munich, Germany

Abstract In this talk, I will discuss methods and tools for building better datasets for deep learning based imaging, with a focus on magnetic resonance imaging. Better datasets enable to train models that perform better and are more robust under distribution shifts. Better datasets not only improve model performance, but also model efficiently. Smaller models trained on better datasets can achieve the performance of their larger counterparts.

MediConfusion: Can you trust your AI radiologist? Probing the reliability of multimodal medical foundation models

Mahdi Soltanolkotabi*
*University of Southern California, USA

Abstract Multimodal Large Language Models (MLLMs) have tremendous potential to improve the accuracy, availability, and cost-effectiveness of healthcare by providing automated solutions or serving as aids to medical professionals. Despite promising first steps in developing medical MLLMs in the past few years, their capabilities and limitations are not well-understood. Recently, many benchmark datasets have been proposed that test the general medical knowledge of such models across a variety of medical areas. However, the systematic failure modes and vulnerabilities of such models are severely underexplored with most medical benchmarks failing to expose the shortcomings of existing models in this safety-critical domain. In this paper, we introduce MediConfusion, a challenging medical Visual Question Answering (VQA) benchmark dataset, that probes the failure modes of medical MLLMs from a vision perspective. We reveal that state-of-the-art models are easily confused by image pairs that are otherwise visually dissimilar and clearly distinct for medical experts. Strikingly, all available models (open-source or proprietary) achieve performance below random guessing on MediConfusion, raising serious concerns about the reliability of existing medical MLLMs for healthcare deployment. We also extract common patterns of model failure that may help the design of a new generation of more trustworthy and reliable MLLMs in healthcare.

Towards Multi-Modal Foundation Models for 3D Medical Imaging

Akshay Chaudhari*
*Stanford University, USA

Abstract Over 20 million abdominal computed tomography (CT) scans are performed annually in the US. Given the current shortage radiologists, artificial intelligence may help alleviate the burden of interpreting these complex imaging studies while simultaneously using the images to extract novel physiological insights. Prior approaches for automated medical image interpretation leverage vision language models (VLMs), but are generally limited to 2D images and short reports, and do not leverage electronic health record (EHR) data for supervision. To overcome these shortcomings for abdominal CT interpretation, we introduce Merlin - a 3D VLM that leverages both structured EHR and unstructured radiology reports for supervision without requiring additional manual annotations. We train Merlin using a high-quality clinical dataset of paired CT scans (6+ million images from 15,331 CTs), EHR diagnosis

codes (1.8+ million codes), and radiology reports (6+ million tokens). We evaluate Merlin on 6 task types and 752 individual tasks. The non-adapted (off-the-shelf) tasks include zero-shot findings classification (31 findings), phenotype classification (692 phenotypes), and zero-shot cross-modal retrieval, while model adapted tasks include 5-year chronic disease prediction (6 diseases), radiology report generation, and 3D semantic segmentation (20 organs). We perform internal validation on a test set of 5,137 CTs, and external validation on 7,000 clinical CTs and on two public CT datasets. We further assess the efficacy of various network architectures and training strategies to depict that Merlin has favourable performance to existing task-specific baselines. We derive data scaling laws to empirically assess training data needs for requisite downstream task performance. We perform all training on a single GPU. We plan to release our trained models, code, and dataset. Underpinning preprint is available on arXiv:2406.06512.

3D Computer Vision in the Age of Deep Learning

Daniel Cremers*
*Technical University of Munich, Germany

Abstract Since 2012, deep learning has revolutionized image and data analysis. In my presentation, I will present some of our work to extend the power of neural networks beyond image classification to problems like protein prediction, shape analysis and visual SLAM. Moreover, I will highlight recent efforts to combine 3D computer vision with large language models.

Learning Generative Models from Corrupted Data

Giannis Daras*
*University of Texas at Austin, USA

Abstract In scientific applications, generative models are used to regularize solutions to inverse problems. The quality of the models depends on the quality of the data on which they are trained. While natural images are abundant, in scientific applications access to high-quality data is scarce, expensive, or even impossible. For example, in MRI the quality of the scan is proportional to the time spent in the scanner and in black-hole imaging, we can only access lossy measurements. Contrary to high-quality data, noisy samples are generally more accessible. If we had a method to transform noisy points into clean ones, e.g., by sampling from the posterior, we could address these challenges. A standard approach would be to use a pre-trained generative model as a prior. But how can we train these priors in the first place without having access to data? We show that one can escape this chicken-egg problem using diffusion-based algorithms that account for the corruption at training time. We present the first algorithm that provably recovers the distribution given only noisy samples of a fixed variance [1]. We extend our algorithm to account for heterogeneous data where each training sample has a different noise level [2]. The underlying mathematical tools can be generalized to linear measurements with the potential of accelerating MRI [3]. Our method has deep connections to the literature on learning supervised models from corrupted data, such as SURE and Noise2X. Our framework opens exciting possibilities for generative modeling in data-constrained scientific applications. We are actively working on applying this to denoise proteins from the Protein Data Bank and we present some preliminary results of this endeavor. Associated articles are available at the DOI: [1] 10.5555/3692070.3692471 and on arXiv: [2] 2411.02780, [3] 2305.19256.

Diffusion Foundation Models for Inverse Problems

Matt Bendel*, Rizwan Ahmad*, and **Philip Schniter***
*Ohio State University, USA

Abstract When using diffusion algorithms to solve imaging inverse problems, the main technical challenge is computing the measurement-conditional score function. Although many methods have been proposed to approximate the conditional score function, we demonstrate that they can be relatively inaccurate. In response, we propose a novel approximation that uses an iterative scheme with a "renoising" step. The renoising step ensures that the unconditional score function sees white noise, in accordance to how it is trained. Finally, we embed our iterative approximation into the denoising diffusion implicit model framework. We demonstrate excellent performance on a wide range of linear and non-linear inverse problems over a wide range of computational budgets. Underpinning preprint is available on arXiv: 2501.17468.

AI4EYE: Translation and Fusion of Multimodal Retinal Imaging

Xin Tian*, Nantheera Anantrasirichai†, Lindsay Nicholson†, and **Alin Achim**†
*University of Oxford, UK, †University of Bristol, UK

Abstract Optical coherence tomography (OCT) and confocal microscopy are pivotal in retinal imaging, offering distinct advantages and limitations. In vivo OCT offers rapid, non-invasive imaging but can suffer from clarity issues and motion artifacts, while ex vivo confocal microscopy, providing high-resolution, cellular-detailed colour images, is invasive and raises ethical concerns. To bridge the benefits of both modalities, we propose a novel framework based on unsupervised 3D CycleGAN for translating unpaired in vivo OCT to ex vivo confocal microscopy images. This helps exploit the inherent 3D information of OCT and translate it into the rich, detailed colour domain of confocal microscopy. Our approach effectively synthesizes colour information from 3D confocal images, closely approximating target outcomes and suggesting enhanced potential for diagnostic and monitoring applications in ophthalmology. While translating OCT into confocal microscopy is of value in its own right, it can also serve as a preliminary step before image fusion, enabling single modality – like fusion approaches. With this in mind, we also discuss here the Topology-Aware Graph Attention Network (TaGAT) for multi-modal retinal image fusion, leveraging a novel Topology-Aware Encoder (TAE) with Graph Attention Networks (GAT) to effectively enhance spatial features with retinal vasculature's graph topology across modalities. The TAE encodes the base and detail features, extracted via a Long-short Range (LSR) encoder from retinal images, into the graph extracted from the retinal vessel. Within the TAE, the GAT-based Graph Information Update block refines and aggregates the node features to generate topology-aware graph features. These are combined with base and detail features and decoded as a fused image. Our model outperforms state-of-the-art methods in OCT and confocal microscopy retinal image fusion. Associated preprints are available on arXiv: 2408.04091, 2407.14188.

Blind Image Restoration via Fast Diffusion Inversion

Hamadi Chihaoui*, Abdelhak Lemkhenter†, and Paolo Favaro*
*University of Bern, Switzerland, †Microsoft Research

Abstract Image Restoration (IR) methods based on a pre-trained diffusion model have demonstrated state-of-the-art performance. However, they have two fundamental limitations: 1) they often assume that the degradation operator is completely known and 2) they alter the diffusion sampling process, which may result in restored images that do not lie onto the data manifold. To address these issues, we propose Blind Image Restoration via fast Diffusion inversion (BIRD) a blind IR method that jointly optimizes for the degradation model parameters and the restored image. To ensure that the restored images lie onto the data manifold, we propose a novel sampling technique on a pre-trained diffusion model. A key idea in our method is not to modify the reverse sampling, i.e., not to alter all the intermediate latents, once an initial noise is sampled. This is ultimately equivalent to casting the IR task as an optimization problem in the space of the input noise. Moreover, to mitigate the computational cost associated with inverting a fully unrolled diffusion model, we leverage the inherent capability of these models to skip ahead in the forward diffusion process using large time steps. We experimentally validate BIRD on several image restoration tasks and show that it achieves state of the art performance. Underpinning preprint is available on arXiv: 2405.19572.

Session 3

Uncertainty Quantification in Computational Imaging

Organiser: Mujdat Cetin

Summary *This session explores recent advances in the theory and methods of uncertainty quantification for computational imaging. As computational imaging involves solving ill-posed inverse problems using complicated estimators, quantifying uncertainties associated with formed imagery, including both aleatoric (stochastic) and epistemic (systematic or modelling) uncertainties, has been an important topic of interest. Recent emergence of deep learning-based image formation methods, including generative models, has both increased the need for proper uncertainty quantification with an eye towards trustworthiness, and also provided tools that can possibly enable progress in that direction. This session highlights a variety of statistical and learning-driven recent work in this area, including estimation and visualization of epistemic and aleatoric uncertainties, posterior variance-based error quantification, distribution-free uncertainty quantification using conformal prediction, as well as task-driven uncertainty quantification for computational imaging, among others.*

Does your computational imaging algorithm know what it doesn't know?

Mujdat Cetin* and Canberk Ekmekci*
*University of Rochester, USA

Abstract Computational imaging involves solving ill-posed inverse problems using complicated estimators. Accordingly, quantifying uncertainties associated with formed imagery has been an important topic of interest. This includes both aleatoric (stochastic) and epistemic (systematic or modeling) uncertainties. We consider this problem in the context of deep learning-based image formation methods. Emergence of such methods has both increased the need for proper uncertainty quantification, and also provided tools that can possibly enable progress in that direction. The need stems from the broader goal of achieving trustworthy learning-based computational imaging, where uncertainty quantification plays a key role. Recent advances in generative modeling and Bayesian neural networks have provided computational tools that can contribute to the development of uncertainty-aware imaging methods. In this talk, we discuss uncertainty quantification for computational imaging methods that leverage physics-based information about the observation process, while using deep learning to construct implicit, yet powerful, priors/regularizers. Specifically, we consider several image reconstruction paradigms including plug-and-play (PnP) regularization, deep unrolling, and generative models, and describe how one can add uncertainty quantification capability to such methods. Using examples from magnetic resonance imaging, computed tomography, and image deblurring, we discuss how uncertainty quantification can point to training data limitations, out-of-distribution and abnormal test samples, and test data quality limitations. We also discuss how calibration methods/metrics, scoring rules, and distribution-free uncertainty quantification methods such as conformal prediction play a key role in the process of rigorous uncertainty estimation. We also draw connections to other work to be presented in the session. Related works are available on arXiv: 2207.00698, and openreview.net/forum?id=VqFHhTYonI.

Black box guarantees for modern methods in inverse problems

Jeremias Sulam*
*Johns Hopkins University, USA

Abstract This talk presents two complementary developments at the intersection of generative modeling and inverse problems. First, we introduce K-RCPS, a conformal prediction framework for diffusion models that offers entrywise uncertainty guarantees and controls user-specified risk. By leveraging a novel convex optimization approach, K-RCPS achieves multi-dimensional risk control with minimal prediction intervals, demonstrated on real-world image denoising for faces and abdominal CT scans. Second, we will introduce Learned Proximal Networks (LPN), which guarantee exact proximal operators for a data-driven nonconvex regularizer. A new training scheme, proximal matching, recovers the log-prior of the data distribution, thereby ensuring both convergence in inverse problems and interpretability of the learned priors. Together, these methods provide principled, uncertainty-aware reconstruction with state-of-the-art performance.

Uncertainty Visualization via Posterior PCA and Posterior Hierarchical Trees

Tomer Michaeli*
*Hebrew University of Jerusalem, Israel

Abstract When solving ill-posed inverse problems, it is often desirable to communicate to the user the set of likely solutions rather than presenting only a single plausible reconstruction. In this work, we present several highly efficient methods for constructing summarizations of the posterior distribution, providing informative visualizations of uncertainty. Our techniques can wrap around any model trained to predict the minimum MSE (MMSE) estimate (i.e., the posterior mean). Our first approach visualizes posteriors across multiple levels of granularity. It does so by predicting a tree-valued hierarchical summarization of the posterior distribution for any input measurement, in a single forward pass of a neural network. Our second approach predicts the principal components (PCs) of the posterior distribution. These PCs capture the main modes of variation around the posterior mean, so that adding or subtracting them from the MMSE prediction visualizes the principal uncertainty directions. For the task of Gaussian denoising, we show that the posterior PCs can be extracted in closed form from a pretrained MMSE model. Our derivation also allows us to compute the entire marginal distribution along each PC. For general inverse problems, we present a method for training a neural network to directly predict the top K PCs of the posterior in a single forward pass, as well as a method to efficiently predict the entire marginal distribution within the subspace spanned by those PCs. We showcase the efficacy of our approaches across diverse datasets and image restoration tasks. We demonstrate that our methods perform at least on par with baselines that rely on diffusion-based posterior samplers, while being orders of magnitude faster. Associated articles are available on arXiv:2405.15719, 2309.13598, 2309.15533. Code is available on GitHub.

Estimating Epistemic and Aleatoric Uncertainty with a Single Model

Matthew Chan*, Maria Molina*, and **Christopher Metzler***
*University of Maryland, USA

Abstract Estimating and disentangling epistemic uncertainty (uncertainty that can be reduced with more training data) and aleatoric uncertainty (uncertainty that is inherent to the task at hand) is critically important when applying machine learning to high-stakes applications such as medical imaging and weather forecasting. Conditional diffusion models' breakthrough ability to accurately and efficiently sample from the posterior distribution of a dataset now makes uncertainty estimation conceptually straightforward: One need only train and sample from a large ensemble of diffusion models. Unfortunately, training such an ensemble becomes computationally intractable as the complexity of the model architecture grows. In this work we introduce a new approach to ensembling, hyper-diffusion models (HyperDM), which allows one to accurately estimate both epistemic and aleatoric uncertainty with a single model. Unlike existing single-model uncertainty methods like Monte Carlo dropout and Bayesian neural networks, HyperDM offers prediction accuracy on-par with, and in some cases superior to, multi-model ensembles. Furthermore, our proposed approach scales to modern network architectures such as Attention U-Net and yields more accurate uncertainty estimates when compared to existing methods. We validate our method on two distinct real-world tasks: x-ray computed tomography reconstruction and weather temperature forecasting. Underpinning preprint is available on arXiv: 2402.03478.

Generalized locality for robust, lightweight imaging architectures

Ivan Dokmanić*, Amirehsan Khorashadizadeh*, Valentin Debarnot*, and Vinith Kishore*
*University of Basel, Switzerland

Abstract Many recent imaging deep nets leverage locality in the form of patches to reduce sample and parameter complexity. I will show that the idea of locality extends much beyond patches in a way that dovetails with the underlying imaging geometry. This can be used to build lightweight deep reconstructors with excellent out-of-distribution generalization. An exciting application is end-to-end learning for 3D imaging. I will show high-quality end-to-end reconstructions in cryo-electron tomography where OOD performance is crucial to mitigate hallucinations. We'll peek under the hood at the intriguing theoretical bridges to geometric machine learning and microlocal analysis, and analyze a stylized probabilistic model of CT to prove that localizing reconstruction improves sample efficiency and OOD generalization without sacrificing in-distribution performance.

Posterior Variance Based Error Quantification in Imaging

Andreas Habring*
*Graz University of Technology, Austria

Abstract In this poster, a method for obtaining pixel-wise error bounds in Bayesian regularization of inverse imaging problems is presented. The proposed method employs estimates of the posterior variance together with techniques from conformal prediction in order to obtain coverage guarantees for error bounds, without making any assumption on the underlying data distribution. It is generally applicable to Bayesian regularization approaches, independent, e.g., of the concrete choice of the prior. Furthermore, the cover-age guarantees can also be obtained in case only approximate sampling from the posterior is possible. With this, in particular, the proposed framework is able to incorporate any learned prior in a black-box manner. Guaranteed coverage without assumptions on the underlying distributions is only achievable since the magnitude of the error bounds is, in general, unknown in advance.

Nevertheless, experiments with multiple regularization approaches confirm that in practice, the obtained error bounds are rather tight. For realizing the numerical experiments, a novel primal-dual Langevin algorithm for sampling from non-smooth distributions is introduced. Underpinning preprint is available on arXiv: 2212.12499.

Image-to-Image Regression with Distribution-Free Uncertainty Quantification and Applications in Imaging

Amit Kohli*
*University of California, Berkeley, USA

Abstract Image-to-image regression is an important learning task, used frequently in biological imaging. Current algorithms, however, do not generally offer statistical guarantees that protect against a model's mistakes and hallucinations. To address this, we develop uncertainty quantification techniques with rigorous statistical guarantees for image-to-image regression problems. In particular, we show how to derive uncertainty intervals around each pixel that are guaranteed to contain the true value with a user-specified confidence probability. Our methods work in conjunction with any base machine learning model, such as a neural network, and endow it with formal mathematical guarantees – regardless of the true unknown data distribution or choice of model. Furthermore, they are simple to implement and computationally inexpensive. We evaluate our procedure on three image-to-image regression tasks: quantitative phase microscopy, accelerated magnetic resonance imaging, and super-resolution transmission electron microscopy of a Drosophila melanogaster brain.

Intriguing Properties of Modern GANs

Roy Friedman* and Yair Weiss*
*Hebrew University of Jerusalem, Israel

Abstract Modern GANs have achieved remarkable success in generating high-quality, realistic images, leading many to believe that "GANs capture the training data manifold", a hypothesis backed by excellent performance in various evaluation metrics such as FID. Our findings reveal that GANs exhibit many concerning behaviors, despite their ability to produce convincing images. For example, GANs are better at reconstructing out-of-distribution images than those from the training set, displaying a bias toward images with less local variation. These biases lead to poor performance in tasks like classification, outlier detection, and compressed sensing. When more closely investigating the density over images implied by the generator of the GAN, we find that it is higher for images with less local variation and is surprisingly highly correlated with a model trained on small patches of the image. Our analysis suggests that while modern GANs are able to create incredibly realistic images, they do not faithfully approximate the training data manifold, calling into question their reliability for tasks that require precise modeling of data distributions. Underpinning preprint is available on arXiv: 2402.14098v1.

Uncertainty quantification for fast reconstruction methods using augmented equivariant bootstrap: Application to radio interferometry

Tobias Liaudat*
*CEA, France

Abstract The advent of next-generation radio interferometers like the Square Kilometer Array (SKA) promises to revolutionise our radio astronomy observational capabilities. The unprecedented volume of data these devices generate requires fast and accurate image reconstruction algorithms to solve the ill-posed radio interferometric imaging problem. Many image reconstruction methods are based on iterative algorithms, e.g. convex optimisation, and even if they achieve good reconstruction qualities, their long iterative nature makes them unsuitable for SKA-like scales. Recent developments based on deep learning have motivated new algorithmic structures with a reduced and fixed number of iterations, such as unrolled architectures or residual to residual network series, that can produce very fast, high-quality reconstructions. However, these state-of-the-art reconstruction methods lack trustworthy and scalable uncertainty quantification, which is critical for the rigorous scientific interpretation of radio observations. We propose an unsupervised technique based on a conformalized version of a radio-augmented equivariant bootstrapping method,

which allows us to quantify uncertainties for fast reconstruction methods. We propose a set of group transformations adapted to the radio problem that leads to reliable coverage probability estimations. Noticeably, we rely on reconstructions from ultra-fast unrolled algorithms, as the uncertainty quantification (UQ) is method-independent. The proposed UQ method brings more reliable uncertainty estimations to high-dimensional radio imaging problems than existing alternatives. Underpinning preprint is available on arXiv: 2410.23178.

Task-Driven Uncertainty Quantification in Inverse Problems via Conformal Prediction

Jeffrey Wen*, Rizwan Ahmad*, and **Philip Schniter***
*Ohio State University, USA

Abstract In imaging inverse problems, one seeks to recover an image from missing/corrupted measurements. Because such problems are ill-posed, there is great motivation to quantify the uncertainty induced by the measurement-and-recovery process. Motivated by applications where the end goal of imaging is determined by a downstream task like classification or the quantitative estimation, we propose a task-centred approach to uncertainty quantification. In particular, we use conformal prediction to construct an interval that is guaranteed to contain the task output from the true image up to a user-specified probability, and we use the width of that interval to quantify the uncertainty contributed by measurement-and-recovery. We then propose several methods that yield adaptive intervals that leverage posterior-sampling-based image recovery. Furthermore, we propose a multi-round measurement protocol, where measurements are collected until the task uncertainty falls below an acceptable level. We demonstrate our methodology on accelerated magnetic resonance imaging (MRI). Underpinning preprint is available on arXiv: 2501.17468.

Part III

Astronomical Imaging

Session 4

Scalable Interferometric Imaging in the Big Data Era

Organiser: **Stefan Wijnholds**

Summary *With the development and deployment of LOFAR2.0, SKA and ngVLA, radio interferometry is entering the Big Data Era. To keep up with the data deluge that will be generated by these instruments with an acceptable costing envelope for computing hardware and energy consumption, at least an order of magnitude improvement in computational performance needs to be realized. Also, data processing pipelines will need to scale well and use processing components optimized for energy efficiency. These challenges require us to rethink the way we process radio interferometric data, opening opportunities for exploration of new avenues. This session aims to present an overview of the challenges and provide a forum for discussion of new ways to process radio interferometric data.*

Scalability and energy efficiency challenges of radio interferometric imaging in the Big Data Era

Stefan Wijnholds*
*Netherlands Institute for Radio Astronomy (ASTRON), The Netherlands

Abstract With the development and deployment of LOFAR2.0, SKA and ngVLA, radio interferometry is entering the Big Data Era, in which a single observation will produce tens of TB of data. To manage this data deluge within an acceptable and sustainable costing envelope for computing hardware and energy consumption, at least an order of magnitude improvement in computational performance is needed. Data processing pipelines need to scale well across multiple nodes and use processing components optimised for energy efficiency. In this talk, I will introduce the magnitude of these challenges using experience from the International LOFAR telescope, showing why current practices are unsustainable. I will then discuss how these challenges are currently being addressed in the development of the iterative calibration pipeline for the SKA [1]. The performance of this pipeline is benchmarked regularly to identify processing bottlenecks and define steps to remedy them. These tests indicate that visibility predictions, calibration solvers and imaging are primary targets for performance improvements. These improvements include use of streaming visibility processing [1], baseline-dependent averaging (BDA) [2,3], a direction-dependent point-spread function [4], image-based prediction and calibration and imaging using direction-dependent effects described as continuous functions [5]. One of the key features of the SKA pipeline architecture is modularity, which should enable us to replace individual processing components by newly developed ones, thereby making these pipelines open to new processing approaches explored by the community. Related works are accessible at the DOIs: [1] 10.46620/URSIA-TRASC24/XDAS1340, [2] 10.23919/AT-AP-RASC54737.2022.9814290, [3] 10.1093/mnras/sty360, [4] 10.23919/URSIGASS57860.2023.10265341, [5] 10.1051/0004-6361/201832858.

R2D2: a deep neural network series for robust ultra-fast precision imaging in radio astronomy

Amir Aghabiglou*, Taylor Chu*, Arwa Dabbech*, and Yves Wiaux*
*Heriot-Watt University, UK

Abstract Mapping the radio sky with the new modern radio-interferometric (RI) arrays requires solving challenging inverse problems for the formation of high-resolution, high-dynamic-range images from large volumes of visibility data. A new generation of image reconstruction algorithms grounded in optimisation theory have demonstrated remarkable capability for imaging precision, well beyond the capability offered by CLEAN. These range from advanced proximal algorithms propelled by handcrafted regularisation operators, such as the SARA family, to hybrid plug-and-play (PnP) algorithms propelled by learned regularisation denoisers, such as AIRI. While already capable of scaling to large image and data dimensions, these techniques are highly iterative, which still hinders their ability to handle the extreme data volumes expected from future instruments. To address this scalability challenge, a novel deep learning approach was recently proposed, dubbed "Residual-to-Residual DNN series for high-Dynamic-range imaging". R2D2's reconstruction is formed as a series of residual images, iteratively estimated as outputs of DNNs, taking the image estimate from the previous iteration and the associated data residual as inputs. R2D2 thus features a hybrid structure between a PnP algorithm and a learned version of the matching pursuit algorithm underpinning CLEAN. We will dive into the R2D2 algorithmic structure and discuss its validation in simulation and on real data, showing that it opens the door to robust ultra-fast precision RI imaging. If time allows, we will discuss a transfer of technology for magnetic resonance imaging in medicine. Associated articles are available at the DOIs: 10.3847/1538-4365/ad46f5, 10.3847/2041-8213/ad41df, 10.23919/EUSIPCO63174.2024.10715010.

Radio-interferometric Imaging on the Sphere using Spherical Harmonics

Michael Kriele*
*University of Western Australia, Australia

Abstract In the international radio astronomy community, a key focus is the exploration of the early Universe through the detection of the 21cm hydrogen line from the Epoch of Reionisation (EoR). However, this faint signal is obscured by light from celestial bodies at present and low redshifts, overwhelming the 21cm line by as much as five orders of magnitude. Consequently, effectively removing these contaminating foregrounds is crucial for EoR detection. This can be achieved by generating sky maps that cover a broad range of frequencies and angular scales. In this talk, we will explore the use of spherical harmonics for computational-effective radio-interferometric imaging on the sphere whilst maintaining wide-field accuracy. Additionally, multi-system spherical harmonic interferometry will be introduced as a method to accurately capture the full celestial sphere using multiple interferometers simultaneously in a single 24-hour imaging sweep. Furthermore, we will discuss various statistical methods, including both stochastic and deterministic approaches; CLEANing techniques, with and without prior inference; and methods for noise and bias correction. Associated publications are accessible at the DOIs: 10.1017/pasa.2022.2, 10.23919/AT-AP-RASC54737.2022.9814200, 10.3847/1538-4357/ad2e9b, 10.5281/zenodo.7583243.

The Africanus Ecosystem: Scalable Radio Interferometric Data Processing, From On-Premises to Cloud

Oleg Smirnov*
*Rhodes Univerity, South Africa

Abstract The data volumes produced by new radio facilities such as MeerKAT, LOFAR and ASKAP has placed strain on existing software stacks, and with the advent of the SKA this problem will only become more acute. At the same time, algorithmic innovations (including those highlighted at this conference) create a pressing need for rapid, yet operationally scalable implementations, so that these algorithms can be tested and refined on real data. This testing also requires that these algorithms be deployed in scalable and distributed workflows, alongside existing software packages. The Africanus ecosystem is a new software stack designed to address some of these issues. By leveraging technologies developed in the PyData ecosystem (xarray, Dask, zarr, Apache Parquet), it allows for both rapid implementation and reasonable scalability of algorithms out of the box. This is combined with I/O layer libraries (dask-ms and xarray-ms) that provides API compatibility with existing data schemas (MSv2, FITS), while implementing parallel I/O in the backend, with support for cloud-based object store such as S3. Finally, the Stimela2 workflow management framework allows diverse packages to be deployed as complex, distributed and reproducible workflows, described by succinct and human-readable YaML recipes. Stimela2 supports various deployment scenarios, supporting both local software installs and Apptainer images, and scheduling jobs locally, via Slurm or Kubernetes. The latter option enables workflows to be executed on cloud providers such as AWS. As a real-life illustration, Africanus forms the foundation of a new high-cadence imaging and transient search pipeline named TRON, which is being deployed on archival MeerKAT data. This is routinely churning through Terabytes of data and finding new transient candidates, which I will illustrate in this talk.

DSA-2000: Novel Large-N Streaming Radio Camera

Joshua Albert*
*California Institute of Technology, USA

Abstract In this talk, I will present the DSA-2000, the first streaming large-N radio camera that offers new possibilities in computational imaging for astronomy. Traditionally, increasing the number of antennas N in radio interferometry has been challenging due to the $O(N^2)$ scaling of data volume and processing demands, along with the iterative nature of conventional inversion algorithms. However, hidden within large-N arrays is the possibility to eliminate the need for iterative imaging and enable real-time streaming processing. This advancement parallels the shift from analogue to "point-and-shoot" digital cameras. The DSA-2000 enables several significant capabilities: a streaming radio camera that achieves continuous, real-time imaging; ultra-fast survey speeds, accelerating survey

timelines; real-time follow-up of transient events; and enhanced sensitivity to nanohertz gravitational waves through pulsar timing arrays. To realise these capabilities, we have developed novel calibration and imaging techniques specifically designed for large-N systems, including a state-of-the-art forward modelling approach to validate and optimise the design against science requirements. In this presentation, I will discuss the signal processing innovations and computational strategies that underpin the DSA-2000's capabilities, addressing how we solve the inverse problem, both calibration and imaging, in real-time.

Compressive radio-interferometric sensing with random beamforming as rank-one signal covariance projections

Laurent Jacques[*], Olivier Leblanc[*], Taylor Chu[†], and Yves Wiaux[†]
[*]Université Catholique de Louvain, Belgium, [†]Heriot-Watt University, UK

Abstract In Radio-interferometry (RI), an array of antennas probes cosmic signals coming from a given region of the sky. The covariance matrix of the vector gathering these antenna measurements offers, thanks to the Van Cittert-Zernike theorem, an incomplete and noisy Fourier sensing of the image of interest. The number of Fourier measurements—or visibilities—scales as the product of the square of the number of antennas and the number of short-time integration (STI) intervals. This work addresses the challenges posed by this vast volume of data, which will increase significantly with the advent of large antenna arrays, by proposing a compressive sensing technique applied directly at the level of the antenna measurements. First, we show how beamforming—a common technique of antenna signal dephasing usually used to focus on some region of the sky—is equivalent to sensing a rank-one projection (ROP) of the signal covariance matrix. We build upon a similar approach made in compressive lensless multiple core fibre imaging to propose a compressive sensing scheme relying on random beamforming, trading the dependence on the (square) number of antennas of the data size for a much smaller number of ROP measurements. We provide image recovery guarantees for sparse image reconstruction. Second, the data size is made independent of the number of STI intervals by modulating the ROP vectors with Bernoulli random weights. The resulting sample complexities, theoretically derived in a simpler case without modulations and numerically obtained in phase transition diagrams, are shown to scale linearly as the image sparsity level, at much lower sampling rates than the conventional approach. Underpinning preprint is available on arXiv: 2409.15031.

Democratizing Super-Resolution Imaging with the MeerKAT

Landman Bester[*]
[*]South African Radio Astronomy Observatory, South Africa

Abstract While CLEAN remains the de facto standard in radio interferometric imaging due to its maturity and speed, promising alternative algorithms with transformative potential have struggled to gain traction in the big data era, hindered by computational demands and complex parameter tuning. We present pfb-imaging: a flexible framework that implements the scaffolding required to accelerate general radio interferometric imaging algorithms in a way that scales with image size rather data volume. This breakthrough enables processing of terabyte-scale data using state-of-the-art algorithms on modest computing resources. As the imaging cornerstone of the Africanus software ecosystem (paper series in prep.), pfb-imaging aims to power MeerKAT and MeerKAT+ pipelines across platforms from high-performance computing clusters to commodity AWS instances. We demonstrate the framework's capabilities with breathtaking super-resolved MeerKAT images of the Milky way, opening new frontiers in radio astronomy.

MROP: Modulated Rank-One Projections for compressive radio-interferometric imaging

Taylor Chu[*], Olivier Leblanc[†], Laurent Jacques[†], and Yves Wiaux[*]
[*]Heriot-Watt University, UK, [†]Université Catholique de Louvain, Belgium

Abstract The emerging generation of radio-interferometric (RI) arrays are set to form images of the sky with a new regime of sensitivity and resolution. This in turn implies a significant increase in data volumes, *i.e.*, the number $\mathcal{O}(Q^2B)$ of noisy Fourier measurements–or visibilities–for Q antennas and B batches, calling for efficient

data dimensionality reduction techniques. Recently, we have proposed a compression scheme, coined modulated rank-one projection (MROP) model [1], that compresses the $Q \times Q$ covariance matrix associated with each batch into a smaller number P of ROPs and compresses across time by trading the B batches for a small number M of Bernoulli modulations of the ROP measurement vectors. We provide a dual perspective on MROP, which can either be understood as random beamforming at acquisition, or as a post-correlation compression. With the latter perspective, the MROP model is validated over extensive numerical experiments including multiple ground truth images, *uv*-coverages, dynamic ranges, and sampling regimes by comparison with the classical, subsampling, and baseline-dependent averaging (BDA) models, and using the uSARA optimisation algorithm for image formation. Related work is accessible on arXiv: [1] 2409.15031

Bayesian Imaging for Radio Interferometry with Score-Based Priors

Noé Dia*, M.J.Yantovski-Barth*, Alexandre Adam*, Micah Bowles†, Anna Scaife†, Yashar Hezaveh*‡, and Laurence Perreault-Levasseur*‡
*Université de Montréal, Canada, †The University of Manchester, UK, ‡MILA, Canada

Abstract In radio astronomy, interferometry is a powerful technique that enhances angular resolution beyond what single-dish telescopes can achieve. Since an interferometer measures a subsampled set of Fourier modes (called visibilities), reconstructing sky brightness images from incomplete Fourier information can be framed as an ill-posed inverse problem. CLEAN, the most popular imaging algorithm in radio astronomy, makes strong simplifying assumptions to address this imaging task. It assumes a point-source model of the sky, which breaks down for extended objects like galaxies or protoplanetary disks. Most importantly, CLEAN does not provide uncertainty estimates for reconstructed images. In this work, we leverage generative modeling in machine learning to apply Bayesian imaging to radio interferometry. Specifically, we use score-based models as plug-and-play priors for Bayesian inference in radio interferometric imaging. These models generate samples from a probability distribution by solving a Stochastic Differential Equation (SDE), enabling posterior sampling in high-dimensional spaces. Using a prior trained on optical galaxy images, we test our approach on ALMA data of protoplanetary disks from the DSHARP project. Our method achieves competitive resolution and dynamic range performance compared to CLEAN, despite inherent prior misspecification. This is consistent with previous works showing that, for a sufficiently flexible prior and an informative likelihood, the posterior distribution can deviate significantly from prior expectations. We further evaluate the impact of prior misalignment in statistical inversion by testing our approach across multiple settings, varying the type of score models used and the training dataset. Finally, we assess the calibration of our posteriors using a coverage test, TARP, and show the potential of our approach to achieve reliable posterior calibration for radio interferometry. Underpinning preprint is available on arXiv: 2311.18012.

HyperAIRI: physics-informed hyperspectral plug-and-play image reconstruction in radio interferometry

Chao Tang*, Arwa Dabbech†, Adrian Jackson*, and Yves Wiaux†,
*University of Edinburgh, UK, †Heriot-Watt University, UK

Abstract With the advent of a new generation of radio-interferometric (RI) telescopes, imaging algorithms face the challenge of efficiently forming high-resolution, high-dynamic-range images from large data volumes across wide frequency bands. Recently, AIRI, a plug-and-play algorithm that incorporates deep neural network (DNN) denoisers within a forward-backward proximal splitting framework, has demonstrated state-of-the-art performance in monochromatic RI imaging. In this work, we propose a hyperspectral extension of AIRI. The HyperAIRI denoiser is designed to reconstruct each spectral channel with side information from neighbour frequencies, also enforcing a physical power-law model whose spectral index can either be estimated during reconstruction or provided as an input. We train a shelf of denoisers, each tailored to a specific dynamic range, employing a Jacobian regularisation term for non-expansiveness to ensure algorithmic convergence. At each forward-backward iteration, the denoisers are applied to all channels in parallel, ensuring efficient computation and seamlessly adapting to the dynamic range of each channel and their total number. We demonstrate the performance of HyperAIRI, on both simulated data and real observations from the Australian SKA Pathfinder, in comparison with AIRI, the multi-frequency variant of CLEAN, and the hyperspectral variant of the state-of-the-art proximal splitting algorithm uSARA. Related works are accessible on arXiv: 2312.07137, 2302.14149

Session 5

High-dimensional Bayesian Astronomical Imaging and Inverse Problems with Machine Learning

Organiser: Laurence Perreault-Levasseur

Summary *Recent progress in machine learning and generative modelling has opened new avenues to tackle previously insoluble high-dimensional inverse problems in astronomy and astrophysics, particularly in Bayesian image reconstruction. While these methodologies show great promise in a range of applications from field-level cosmology to differentiable optics systems, multiple open problems stand in the way of groundbreaking discoveries. This session will explore recent applications and proposals addressing the development of computationally tractable methodologies to reconstruct posterior samples in imaging problems. We will also discuss assessing their accuracy in real-world settings and addressing the problem of robustness to distributional shifts, which remains a key focus of current research.*

Data-Driven High-Dimensional Inverse Problems: A Journey Through Strong Lensing Data Analysis and Other Imaging Inverse Problems

Laurence Perreault-Levasseur*
*Université de Montréal, Canada

Abstract Strong gravitational lensing offers a powerful tool to probe the universe, from mapping the distribution of dark matter to studying the formation of distant galaxies. However, due to the nonlinear, non-convex, and very high-dimensional nature of the inverse problem, traditional lensing analysis methods have struggled to convincingly extract all available information from existing observations. This talk explores recent advancements in Bayesian image reconstruction and high-dimensional inference, with a focus on applications to strong gravitational lensing. Using diffusion models and score-based generative approaches, I will present methods for reconstructing high-fidelity astronomical images while addressing key challenges such as out-of-distribution data robustness, uncertainty quantification, and joint inference of hierarchical properties. I will also present statistical tools to assess the accuracy of posterior samples obtained with machine learning. These advances in deep learning-based inference open new possibilities for studying complex lensing systems at scale, particularly in the era of large surveys like LSST, Euclid, and the Roman Space Telescope.

Multimodal and Foundational Priors for Astronomical Inverse Problems

François Lanusse*
*CNRS, France

Abstract It is by now well established that generative models can be used as data-driven priors for solving inverse problems in high-dimensions in a sound Bayesian way. However, when thinking about the broad applicability of such methods, a number of additional challenges appear. In particular, in observational sciences, high quality reference data which may be used for building such priors are either very rare, or very expensive to acquire. Furthermore, these reference samples may be systematically biased compared to the population of objects to be analysed at inference time. Finally, for the purpose of being able to easily reuse and share prior models, there is an interest to build model as generic and flexible as possible, potentially capable of modeling any type of astronomical data. In this talk, I will present our recent work on building methodologies that can address these challenges. In particular, a framework for building diffusion models from corrupted data using an Expectation-Maximization approach, which can also be used to adjust a data-driven prior to a particular target population. I will also present our work on building multimodal generative models, capable of modeling the joint and conditional distributions of various observational modalities (images, time series, spectra). Combined, these methodologies chart a path towards the creation of Foundational data-driven prior models, that have the flexibility to adapt to different desired populations at inference time, and can be used for many types of observations.

Complex galaxy models with conditional score models

Peter Melchior*
*Princeton University, USA

Abstract The morphology of galaxies is complex, and in many cases the noise level or resolution of astronomical images necessitates placing priors to solve inverse problems such as inpainting, superresolution, or deblending. In the last few years, score-matching models have emerged as powerful option for data-driven priors. We present an extension of this approach when multiple parameters, and thus their priors, are conditionally dependent. Specifically, we present the inference of spatially resolved dust morphologies, which are directly dependent on the light distribution because dust is created by stars. We will show the first results of applying this novel method on multi-band images of galaxies at cosmological distances and demonstrate how these results can be used to decontaminate galaxy photometry and study the interplay between stellar and dust distributions.

3D imaging in cosmology: from dark matter mapping to initial conditions reconstruction

Carolina Cuesta*
*Massachusetts Institute of Technology, USA

Abstract The upcoming era of cosmological surveys promises an unprecedented wealth of observational data that will transform our understanding of the universe. Surveys such as DESI, Euclid, and the Vera C. Rubin Observatory will provide extremely detailed maps of billions of galaxies out to high redshifts. Analysing these massive datasets poses exciting challenges that machine learning is uniquely poised to help overcome. In this talk, I will highlight recent examples from my work on probabilistic machine learning for cosmology. First, I will explain how a point cloud diffusion model can be used both as a generative model for 3D maps of galaxy clustering and as a likelihood model for such datasets. Moreover, I will present a generative model developed to reconstruct the initial conditions of the Universe from spectroscopic survey observations. Finally, I will introduce my work on continuous representations in cosmology and their use for data compression. When combined with the wealth of data from upcoming surveys, these machine learning techniques have the potential to provide new insights into fundamental questions about the nature of the Universe.

Autoregressive Modelling in Simulation-Based Inference for High-Dimensional Cosmological Problems

Noemi Anau Montel*
*University of Amsterdam, GRAPPA Institute, The Netherlands

Abstract High-dimensional inverse problems are central to modern cosmology, where inferring complex posterior distributions from observational data poses significant computational challenges. Recently, simulation-based inference (SBI) algorithms have advanced remarkably and are now applied across a wide range of cosmological analyses. We propose employing autoregressive modelling within SBI frameworks for efficient Bayesian inference in high-dimensional settings. We demonstrate the effectiveness of this approach in two frameworks for distinct astrophysical and cosmological analyses. First, we apply autoregressive modelling to neural ratio estimation, to accurately reconstruct posterior distributions over high-dimensional parameter spaces. We introduce a slice-based nested sampling algorithm to efficiently draw both posterior and constrained prior samples from ratio estimators—the latter being instrumental for sequential inference. We demonstrate the computational capabilities of this framework on proof-of-concept analyses of stellar streams and substructure searches in gravitational lensing images [1]. Second, we apply autoregressive modelling to field-level reconstruction of the universe's initial conditions from late-time density field observations. By modelling the posterior distribution over initial density fields auto-regressively and assuming a Gaussian likelihood, we efficiently sample from the high-dimensional posterior using Gibbs sampling with Exact Data Augmentation (GEDA). This allows us to reconstruct the primordial density field in a two-dimensional cosmological setting, showcasing the potential of autoregressive modelling for handling intrinsically nonlinear and high-dimensional cosmological data. The proposed technique is applicable to generic (non-differentiable) forward simulators and enables fast sampling from the posterior of the underlying field [2]. Our results highlight autoregressive modelling as a powerful tool for high-dimensional SBI in cosmology, capable of addressing complex inverse problems with improved computational efficiency. Associated articles are available on arXiv: [1] 2308.08597, [2] 2310.19910

Echoes in the Noise: Posterior Samples of Faint Galaxy Surface Brightness Profiles with Score-Based Likelihoods and Priors

Alexandre Adam*, Connor Stone*, Connor Bottrell†, Ronan Legin*, Yashar Hezaveh*, and Laurence Perreault-Levasseur*‡
*Université de Montréal, Canada, †The University of Western Australia, Australia, ‡Mila, Canada

Abstract Examining the detailed structure of galaxy populations provides valuable insights into their formation and evolution mechanisms. Significant barriers to such analysis are the non-trivial noise properties of real astronomical images and the point spread function (PSF) which blurs structure. Here we present a framework which combines recent advances in score-based likelihood characterization and diffusion model priors to perform a Bayesian analysis of image deconvolution. More specifically, we train a score-based model to learn the distribution of galaxy images as a data-driven informative prior. Furthermore, we train a score-based model on samples of noise from the Hubble Space Telescope to learn the distribution of noise, which is used to construct a likelihood function for posterior inference of idealized brightness profiles from noisy and blurred observation from the Hubble Space Telescope. The method, when applied to minimally processed Hubble Space Telescope data, recovers structures which have otherwise only become visible in next-generation James Webb Space Telescope imaging. Underpinning preprint is available on arXiv: 2311.18002.

Toward Attaining the Diffraction Limit and Beyond in Astronomical Imaging

Farzad Kamalabadi*
*University of Illinois, Urbana-Champaign, USA

Abstract It is well known that conventional spectral imaging systems for astronomical imaging are unable to achieve angular resolutions that are even close to the diffraction limit due to the inherent aberration effects in reflective optics. Furthermore, common approaches for aberration correction based on post-processing techniques are intrinsically limited in their effectiveness as they cannot compensate for the unknowable distortion effects in the image formation process. Recently, high-resolution imaging concepts and systems have been proposed based on the combination of diffractive optical elements and embedded computational imaging techniques to yield unprecedented resolution [1]. The approach requires the computational solution of an inverse problem to address the wavelength-dependent focusing nature of diffractive lenses, which motivates the development of statistical learning techniques for the solution of the multi-frame deblurring inverse problem resulting from multiple exposures at different focal lengths. The aim of the work presented here is to expand and enhance the capabilities of such diffractive imaging systems via computationally tractable methodologies for posterior sampling in the associated Bayesian image reconstruction problem via incorporating data-driven learned priors and to investigate their robustness to distribution shifts. Related work is accessible at the DOI: [1] 10.1109/TCI.2021.3075349.

Inverse-problem versus principal component analysis methods for angular differential imaging of circumstellar disk

Sandrine Juillard*, Valentin Christiaens*, and Olivier Absil*
*University of Liège, Belgium

Abstract The direct observation of faint planetary systems around bright stars is a promising method for studying planet formation. Nevertheless, raw images from high-contrast imaging instruments are often contaminated by residual stellar light, appearing as quasi-static speckles. An observing strategy, such as angular differential imaging, combined with image processing, is essential to mitigate these speckles. However, state-of-the-art algorithms have limitations, often introducing geometric distortions, particularly to extended structures. Recent observations have revealed a wide variety of structured dust disks around stars that host forming planetary systems (circumstellar disk). These disks could provide crucial insights into planet formation, but they may remain under-explored due to distortions introduced by image processing algorithms, which could also lead to false planet detections. Recently, two types of approaches have been proposed to recover more robust disk images: iterative principal component analysis (I-PCA) and inverse

problem (IP) approaches, with three distinct IP implementations. This raises the question of what differences exist between the three IP designs and what advantages they offer compared to I-PCA. An examination of the underlying assumptions of each algorithm was conducted to evaluate their reliability, accuracy, and practicality, as well as to understand their advantages and limitations. This is complemented by a testing pipeline to assess the performance of one IP approach (MUSTARD) in comparison to I-PCA. In this application, the IP approach is limited by our understanding of the dynamics and morphology of the speckle field, as well as nuisance phenomena that we cannot yet describe mathematically, which restrict the model's overall performance. In contrast, I-PCA is more robust and easier to parametrize. Underpinning article is available at the DOI: 10.1051/0004-6361/202347259

Posterior Sampling for Solar Snapshot Spectroscopy with Score-Based Priors

Ulas Kamaci[*] and Farzad Kamalabadi[*]
[*]University of Illinois, Urbana-Champaign, USA

Abstract Solar imaging spectroscopy is a fundamental technique for investigating the solar corona. Traditional imaging spectrographs require a time-consuming scanning process to build up the spatial-spectral data cube which renders them unsuitable to capture the quickly evolving transient events in the corona. Slitless spectrographs have been proposed to perform snapshot spectral imaging without requiring scanning; however, their measurements consist of overlapped spectra which needs to be disentangled to reconstruct the spatial spectral data cube. This configuration results in limited-angle tomography, which is a highly under-determined inverse problem. By using a Gaussian model for the spectra, we follow [1] to formulate the inverse problem as a parametric blur estimation problem with non-uniform blur where the goal is to estimate the parameters of the Gaussian blur at each spatial location, producing a nonlinear inverse problem. To solve the inverse problem we adapt a posterior sampling scheme with a score-based prior trained on Gaussian spectral fits of thousands of observations derived from the EIS spectrograph onboard the Hinode observatory. Comparison experiments with an end-to-end trained model demonstrate the effectiveness of the approach in terms of reconstruction quality and robustness against distribution shifts. As demonstrated in this work, such computational approaches to spectral imaging involving machine learning solutions to inverse problems show the promise of superior temporal and spatial resolution for observational astronomy. Related work is accessible at the DOI: [1] 10.1109/TIP.2014.2363903.

Towards a spherical extension of the R2D2 deep neural network series paradigm for wide field radio interferometric imaging

Ayoub Tajja[*], Emma Tolley[*], Jean-Philippe Thiran[*], and Yves Wiaux[†]
[*]EPFL, Switzerland, [†]Heriot-Watt University, UK

Abstract Recently, the R2D2 paradigm, standing for "Residual-to-Residual DNN series for high-Dynamic-range imaging" [1], was introduced for image formation in Radio Interferometry (RI) as a learned version of the traditional algorithm CLEAN. The first implementations of R2D2 are limited to planar imaging on small fields of view, failing to meet the spherical-imaging requirement of modern telescopes observing wide fields. To address this limitation, we propose the spherical-imaging extension S-R2D2. Firstly, we develop a fast implementation of the RI measurement model mapping a spherical image to visibility data. This model is built as the composition of an efficient Fourier-based interpolator mapping the spherical image onto the equatorial plane, with the standard RI operator mapping the equatorial-plane image to visibility data. However, the interpolator must inevitably operate at a lower-than-optimal resolution on the plane to achieve the required large resolution on the sphere, leading to a critical resolution-accuracy trade-off. Secondly, as R2D2, S-R2D2 encapsulates its minor cycles in existing deep neural network (DNN) architectures developed to process Euclidean images, but adapts its scheme to perform reconstructions on the sphere. More precisely, the final reconstruction is built as the sum of spherical estimations, which are the DNNs planar outputs back-projected onto the sphere with the adjoint of the interpolator. Importantly, S-R2D2 modifies R2D2's training loss to impose direct consistency with the spherical ground truth. Finally, while analysing the resolution-accuracy trade-off, we demonstrate in simulation S-R2D2's capability to perform fast and accurate reconstruction of monochromatic intensity spherical images, across a variety of high-resolution high-dynamic range settings. Related work is accessible at the DOI: [1] 10.3847/1538-4365/ad46f5.

Session 6

Computational Imaging for Precision Astrophysics

Organiser: Aviad Levis

Summary *The topic of this session is Computational Imaging for Precision Astrophysics, covering a variety of exciting new approaches for ultra high resolution astronomy. The session focuses on black holes and exoplanets both of which with their unique challenges. Horizon scale black hole science is incredibly challenging due the extreme resolution required (2̃0 micro arcseconds) to resolve the largest black holes on the sky: Sagittarius* $A^\star$ *and* $M87^\star$*. Nonetheless, innovations and advancements in interferometry over the last decade have opened the door to new horizon scale science: from the first images with the Event Horizon Telescope, to precise astrometry with GRAVITY. The talks and posters in this session cover a variety of topics from novel computational imaging that integrates interferometry with information theory and modern machine learning to new mission concepts for space interferometry. For exoplanets, beyond resolution, a key challenge is that of contrast: the signal from the planet could be up to a billion times fainter than that of its host star. Nonetheless, direct imaging holds the key to characterizing composition and habitability. In this session, we will hear about advancements in the field and how precise modelling of instrumental optics could be the key to pushing the detection limit towards the fundamental noise.*

Physics-constrained neural fields for 3D imaging in astronomy

Aviad Levis*
*University of Toronto, Canada

Abstract Three dimensional imaging in astronomy seems counter intuitive: we typically assume 3D information comes from multiple views (e.g. CT in medical imaging) and in most of astronomy (outside of our galactic neighborhood) we only have access to a single view point and cannot rely on parallax cues for 3D reconstructions. Nonetheless in many astronomical systems we find hidden 3D cues embedded in the time or spectral axes. While these signatures are not as strong as multiple views, in principle, they enable recovering some 3D information (or "2.5D"). Cosmological redshift is a classic example of getting 3D positions at the cosmic scale [1]. Doppler shits gives information about the projected velocities of disks. Fast temporal evolution of an orbiting system enables reconstructing the 3D structure through observations over time [2,3]. Modern telescope resolutions are now able to resolve many of those axes (spectral, temporal, polarimetric) to unprecedented accuracy which calls for new algorithms for 3D reconstructions to be developed. In the last few years, neural radiance fields (NeRFs) [4] have completely changed the landscape of 3D imaging within the computer vision and graphics communities. One of the reasons behind their success is the simplicity of constraining a neural representation through ray tracing. The focus of much of the research within the vision community is on developing better representations with efficiency and scalability for complex 3D scenes in mind. In astronomy (and more broadly in the natural sciences), the scenes are not very complex; nonetheless, the radiative transfer physics are often a bottleneck in terms of complexity and runtime, preventing the solution of inverse problems. In this work, I will show how neural fields constrained by ray tracing and complex physics enable 3D imaging in astronomical environments: black holes, protoplanetary disks, and dark matter. Related works are available on arXiv: [1] 2309.04437, [2] 2310.07687, [3] 2204.03715, [4] 2003.08934.

Uncovering Images of Exoplanets with Differentiable Optical Models for Coronagraphs

Jason Wang*, Brandon Feng†, Rodrigo Ferrer-Chavez*, Katherine Bouman‡, Aviad Levis⋆, and William Freeman†
*Northwestern University, USA, †Massachusetts Institute of Technology, USA, ‡Caltech, USA, ⋆University of Toronto, Canada

Abstract Faint exoplanets are tens of thousands to billions of times fainter than the host stars they orbit, making it extremely challenging to distinguish the signal of a planet from the bright glare of their stars. A coronagraph, a series of optics in an exoplanet imaging instrument, is used to suppress and mitigate the glare of the star. Because optical systems are not perfectly stable, small perturbations (changes in the optics of just a fraction of the wavelength of light due to drifts, thermal expansion, turbulence, etc.) cause the residual glare of the star to change over time, and our imperfect removal of the stellar glare limits our sensitivity to finding planets. However, the physics of diffractive optics is known, and one should be able to forward model the changes in the glare of the star due to optical perturbations and subtract them off. I will present a differentiable optical model for coronagraphs that we have applied to simulated and real data from the James Webb Space Telescope. I will explain the optical perturbations we have to model, our optimization procedure, and the improved exoplanet imaging sensitivity that we are achieving with the technique.

Black Hole and Exoplanet Imaging with Information Field Theory

Vishal Johnson* and Torsten Enßlin*
*Max-Planck Institute for Astrophysics, Germany

Abstract The observation of black holes and exoplanets is challenging due to the small angular scales to be resolved through a time dependent ionosphere and atmosphere, the incomplete data, and the noise. Information field theory (IFT) addresses the image reconstruction problem in a probabilistic way. Algorithms designed with IFT exploit sky brightness correlations in space and time, while learning those from the observed data, and are able to calibrate iono- and atmospheric effects. This talk will introduce into IFT and show its application to astrophysics, including exoplanet and black hole imaging.

Kine: Interferometric Dynamic Imaging with Neural Networks

Marianna Foschi*
*Instituto de Astrofisica de Andalucia, Spain

Abstract The talk will present a new approach to dynamic imaging of radio interferometric data, based on a Multi-Layer Perceptron neural network. The method, named Kinetic NeRF (or Kine), was developed to provide time-regularized video reconstructions of sparse observations of variable radio sources, like black holes, relativistic jets and blazars. In contrast to pre-existing methods, Kine does compensate for the sparse available data by incorporating explicit, human-defined priors, which are prone to bias. Instead, it enforces implicit spatial and temporal regularization by parametrizing the output image or video via a neural network. The method relies on a coordinate-based neural representation, an architecture which has recently gained popularity in the computational imaging community as an alternative to discrete array-based representations of functions defined on low-dimensional domains. Kine is especially effective for imaging extremely sparse very long baseline interferometry observations of dynamic sources, such as observations of the Sgr A black hole with the Event Horizon Telescope (EHT), but also for imaging videos of multi-epoch observations of slowly evolving sources, such as blazar observations by monitoring programs. This talk will cover an explanation of the algorithm, show example reconstructions of synthetic data from black hole simulations, and present reconstructions of real multi-epoch observations of relativistic jets.

Extracting movement information from radio interferometry with the Optimal Transport regularizer

Shiro Ikeda*
*Institute of Statistical Mathematics, Japan

Abstract In a set of EHT observations, the number of data points in the Fourier domain is fixed, and dynamic imaging requires dividing the observation data into frames. One straightforward approach to performing dynamic imaging is to extend the Regularized Maximum Likelihood (RML) methods of static imaging, where we now solve for a series of images. Each of these images corresponds to a small data segment. Given that the data points are extremely limited, it is necessary to include additional terms that regularize the images in time. These temporal regularizers control characteristics of the movie as a whole, such as the smooth evolution of the target over the entire observation. The Optimal Transport (OT) distance, which defines a distance metric between two frames, is thus a strong candidate for such a temporal regularizer. The OT distance contains information about the underlying domain of the image and thus can inherently encode information about the physical motion between images, making it particularly useful when we expect coherent motion. By incorporating efficient transport based on the optimal transport problem, we aim to detect dynamics near black holes and assess its potential capabilities as a dynamic regularizer. We will show our results of dynamic imaging with the OT regularizer.

Reliable Deep Generative Priors for Astronomical Imaging

Berthy Feng* and Katherine Bouman*
*California Institute of Technology, Caltec

Abstract Deep generative priors promise to help push the limits of imaging extreme astronomical phenomena. Although deep generative models offer sophisticated image priors, it is unclear how to leverage them in a principled manner. In this poster, I will present two lines of work that make imaging with deep generative priors more reliable. The first proposes a variational approach to Bayesian inference with a diffusion-model prior [1], which allows us to obtain more accurate posteriors than possible with popular guidance-based diffusion approaches. We apply this approach to the real M87* data that the Event Horizon Telescope used in 2019 for the first-ever image of a black hole. We re-imagine M87* under a variety of diffusion-model priors [2], each trained on a different image dataset. By testing the image reconstruction under different assumptions, we can assess which image features are robust to imaging assumptions (e.g., the presence of a ring and location of the bright spot) and which are hallucinations beyond the data. This type of analysis offers a way to leverage AI hallucinations responsibly for astronomical imaging. The second proposes a way to enforce physics constraints onto deep generative models, thereby making them more reliable priors. We propose neural approximate mirror maps [3] to constrain generative models by construction, where a forward mirror map transforms constrained data into an unconstrained space and has a corresponding inverse map.

We can train a generative model in the unconstrained space and map its samples back to the constrained space via the inverse map. Whereas analytical mirror maps only exist for convex constraints and are difficult to derive, our approach for learning the mirror map works for arbitrary constraints. We demonstrate applying it to solve constrained inverse problems, such as data assimilation with PDE constraints. Associated articles are available on arXiv: [1] 2309.01949, [2] 2406.02785, [3] 2406.12816.

From the low ressource to the high ressource setting: Solving imaging challenges in the ngVLA era with lessons learned from VLBI

Hendrik Müller*
*National Radio Astronomy Observatory, USA

Abstract For decades, VLBI imaging relied on CLEAN, which introduced human bias. Recently, RML, Bayesian, and AI methods have emerged, enabling high-quality, automated, super-resolved reconstructions. These methods excel in low-resource settings, where traditional techniques struggle to account for spatial and temporal correlations. The low-resource data regime is ideal for such development: data sizes are manageable enough for computationally intensive methods, and imaging performance significantly improves with advanced spatial correlation processing. Key innovations include dynamic source movies in full polarimetry [1,2], super-resolution imaging with sparse arrays [3,4], joint Bayesian inference of images and calibration quantities [5,6], and calibration-robust polarimetric imaging from closure quantities [7,8]. Next-generation radio interferometers like ngVLA however will bring a new data challenge due to large data volumes, potentially requiring a return to classical matching pursuit approaches within a major/minor loop framework. This contribution reviews recent low-resource developments, examines their potential for arrays like ngVLA, and outlines strategies for effective utilization. Related works are available at the DOIs: [1] 10.3390/galaxies11010012, [2] 10.1051/0004-6361/202449325, [3] 10.3847/1538-4357/aab6a8, [4] 10.1051/0004-6361/202243244, [5] 10.21105/joss.04457, [6] 10.1051/0004-6361/202449663, [7] 10.1051/0004-6361/202346207,[8]10.1051/0004-6361/202450437.

Generative modelling for mass-mapping with fast uncertainty quantification

Jessica Whitney*, Tobias Liaudat*, Matthew Price†, Matthijs Mars†, and Jason McEwen†
*University College London, UK, †CEA, France, ‡University of Leiden, The Netherlands

Abstract Understanding the nature of dark matter in the Universe is an important goal of modern cosmology. A key method for probing this distribution is via weak gravitational lensing mass-mapping - a challenging ill-posed inverse problem where one infers the convergence field from observed shear measurements. Upcoming stage IV surveys, such as those made by the Vera C. Rubin Observatory and Euclid satellite, will provide a greater quantity and precision of data for lensing analyses, necessitating high-fidelity mass-mapping methods that are computationally efficient and that also provide uncertainties for integration into downstream cosmological analyses. In this work we introduce MMGAN, a novel mass-mapping method based on a regularised conditional generative adversarial network (GAN) framework, which generates approximate posterior samples of the convergence field given shear data. We adopt Wasserstein GANs to improve training stability and apply regularisation techniques to overcome mode collapse, issues that otherwise are particularly acute for conditional GANs. We train and validate our model on a mock COSMOS-style dataset before applying it to true COSMOS survey data. Our approach significantly outperforms the Kaiser-Squires technique and achieves similar reconstruction fidelity as alternative state-of-the-art deep learning approaches. Notably, while alternative approaches for generating samples from a learned posterior are slow (e.g. requiring $\tilde{1}0$ GPU minutes per posterior sample), MMGAN can produce a high-quality convergence sample in less than a second. This contribution originated from discussions at BASP Frontiers 2023. Underpinning preprint is available on arXiv: 2410.24197.

Searching for Protoplanets with ALMA Using Regularized Maximum Likelihood Imaging Techniques

Brianna Zawadzki* and Ian Czekala†
*Wesleyan University, USA, †University of St Andrews, UK

Abstract In recent years, protoplanetary disks have been observed at high enough resolutions to detect and characterize a variety of disk substructures, some of which are indicative of active planet formation. The planet-hunting ALMA large program exoALMA has observed 15 protoplanetary disks at high angular and spectral resolution, characterizing disk structures and kinematics in enough detail to detect non-Keplerian features (NKFs) caused by forming protoplanets. As these signatures of young planets can be faint and tiny, robust imaging procedures are critical for detecting embedded protoplanets at high confidence. The exoALMA collaboration has employed two different imaging procedures to ensure that we consistently detect the NKFs: CLEAN, the traditional iterative deconvolution process, and Regularized Maximum Likelihood (RML), an imaging process recently popularized by the Event Horizon Telescope Collaboration. RML constructs an image by maximizing the likelihood of visibility data given predicted model visibility values and regularizers. When applied to sub-mm interferometric observations, RML can achieve higher angular resolution and image fidelity compared to CLEAN, while also serving as an independent verification of marginal features observed in CLEAN images. We used the open source Python package MPoL to obtain full RML image cubes from the exoALMA observations of all protoplanetary disks with localized NKFs, as well as several disks exhibiting large-scale non-Keplerian arcs. We find that RML imaging methods consistently reproduce the NKFs seen in the fiducial CLEAN images of the 12CO emission around several sources, suggesting that the NKFs are real disk features rather than artifacts from a specific imaging procedure.

EHT Julia Organization: Julia Meets Black Hole Imaging and Beyond

Kazu Akiyama*
*MIT Haystack Observatory, USA

Abstract The Event Horizon Telescope (EHT) is a global millimeter-wavelength very long baseline interferometry (VLBI) array. With its planet-sized synthesized aperture and short observing wavelength, the EHT has the unique capability to capture the finest-scale structures in the universe, resolving features on scales of a few tens of microarcseconds. This capability ultimately led to the first-ever images of black holes, including those of M87* and Sgr A*. As a computational telescope, the EHT relies heavily on data processing, which is a fundamental and integral part of the instrument. Currently, the landscape of EHT data processing software is quite heterogeneous – different aspects of data processing are handled by various software packages written in multiple programming languages, with a majority in Python. However, these packages often lack scalability. This poses challenges in meeting the growing demands for more coherent inference schemes that can jointly handle physical, image, and calibration parameters, especially given the significantly larger and more complex datasets expected from the next-generation EHT array. This future array will observe multiple frequencies, at larger bandwidths, and with a substantial increase in the number of antennas. To address these software challenges, we have initiated the EHT Julia Organization to develop and provide a coherent software infrastructure, encompassing calibration, imaging, and modeling, using the high-performance computing language Julia. Julia offers a unique capability to facilitate the development of fast, composable, and extendable software toolkits that can be easily integrated into scalable parallel computing environments with CPU/GPU clusters. Additionally, it supports large-scale inference of complex models through robust auto-differentiation capabilities. In this poster, we present an overview of the EHT Julia packages and discuss future prospects.

Part IV

Medical Imaging

Session 7

"R4" Reconstruction, Resolution, Regularization and Representation – Overcoming Medical Imaging Challenges

Organiser: Julia Schnabel

Summary *The aim of this session is to guide the audience through the medical imaging pipeline – from acquisition and reconstruction to analysis - with a little help from AI. The talks will cover different image modalities, and focus on inverse problems in medical image reconstruction, motion estimation and motion correction. As image quality and accurate diagnostic assessment are key for the applicability of AI-based solutions in the medical context, the talks cover these aspects and provide insights into the uncertainty of AI-based algorithms in image reconstruction and analysis.*

Deep Learning for Motion Correction in MRI

Julia Schnabel[*†‡]
[*]Helmholtz Munich, Germany, [†]Technical University of Munich, Germany, [‡]King's College London, UK

Abstract Medical imaging can be regarded as a composite of "R4: Reconstruction, Resolution, Regularization and Representation" - all leading to different image quality challenges that need to be overcome in a clinical setting. As an example, here we focus on motion artefacts in magnetic resonance imaging, and how these can be corrected through motion reconstruction. Motion of the patient during scan time, including respiratory or other physiological motion patterns, can severely affect further clinical downstream tasks, and may even render the acquired images unsuitable for subsequent diagnostic reporting. This greatly affects clinical workflow in hospitals, leading to undesirable patient recall, delayed timely diagnosis or patient treatment. Identifying and resolving motion corruption either retrospectively or already at scan time, can help alleviate these problems, paving the way for automated scan quality control and active scanning. We first review the state-of-the-art in motion reconstruction [1], after which we develop a number of deep learning based methods that aim to detect and correct motion during MRI reconstruction, for a range of clinical applications in cine cardiac MRI [2,3], liver MRI [4] and quantitative brain MRI [5]. Notably, we operate in the raw MRI k-space domain to identify and remove motion corruption at the source. Using a range of supervised and unsupervised techniques, efficiently combined with undersampled image reconstruction and embedding of different clinical downstream tasks, we demonstrate a family of deep learning methods that can be readily deployed and extended to active scanning and online quality control. We will close by providing a brief overview of some of the current hot topics in the wider medical imaging domain, which aim to overcome these "R4" challenges across different medical imaging modalities, with newly emerging methodological advances designed for different, dedicated clinical applications. Related works are accessible via the DOIs: [1] 10.1109/TMI.2023.3323215, [2] 10.1109/TMI.2020.3008930, [3] 10.1016/j.media.2019.04.009. Underpinning articles are available at the DOIs: [4] 10.1007/978-3-031-72104-5_59, [5] 10.1007/978-3-031-72104-5_54.

Longitudinal MRI with Interleaved High- and Low-Field Scans and Feature-Fusion Transformers

Efrat Shimron[*]
[*]Technion - Israel Institute of Technology, Israel

Abstract High-field MRI systems offer superb-quality images but face high costs and require advanced infrastructure. Emerging low-field MRI systems are affordable and portable, hence they offer exciting potential, yet they are limited by inferior image quality and long scan duration. In future healthcare workflows, these systems may hence be used as complementary modalities. We envision a longitudinal imaging workflow integrating high- and low-field scans, particularly suited for patients requiring ongoing monitoring or for healthy individuals who need routine cancer screening. We hence introduce a computational framework to boost image quality and speed in follow-up low-field scans by harnessing data from baseline high-field scans. We introduce ViT-Fuser, a novel vision transformer architecture that learns and merges features from both scans through a unique feature-fusion module, along with a hybrid loss function that facilitates sharper, more detailed low-field reconstructions. Experiments on longitudinal glioblastoma MRI data demonstrate ViT-Fuser's advantage in image quality and reconstruction metrics. Our method achieves state-of-the-art results, producing cleaner images than techniques that do not utilize prior information. Our strategy can hence enhance and boost future clinical workflows and contribute to increasing MRI accessibility worldwide.

Model-based quantitative MRI meets machine learning

Gary Zhang[*]
[*]University College London, UK

Abstract Magnetic Resonance Imaging (MRI) has established itself as one of the most versatile medical imaging tools in modern healthcare systems. However, conventional MRI is qualitative and produces image contrast that conflates a multitude of underlying properties of the imaged sample. Model-based quantitative MRI offers fundamental advantages over its conventional counterpart but suffers from its own limitations, due to longer acquisition time and processing time. Recent advances in machine learning present major opportunities to address these limitations. In this talk, I will highlight one recent development in this area from my group concerning the challenges

associated with degeneracy in model-based quantitative MRI. Model-based quantitative MRI aims to quantify a set of tissue properties of some sample of interest from a set of MRI signals collected according to some chosen acquisition protocol. The relationship between the measured signals and the underlying tissue properties (and the protocol) is typically described by some mathematical model which underpins the success of this approach. Degeneracy refers to the situations under which similar measured signals arise from very different tissue properties, which presents challenges to standard machine learning approaches. Chemical-shift-encoded MRI (CSE-MRI) is one notable model-based quantitative MRI example where the presence of degeneracy is its hallmark. This talk will use this as an example to illustrate the consequence of the degeneracy on standard machine learning approaches and present a novel solution we have developed to overcome this challenge. Underpinning preprint is available on arXiv: 2403.01178.

Machine Learning in PET Image Reconstruction – Opportunities, Challenges, and Common Pitfalls

Georg Schramm*
*KU Leuven, Belgium

Abstract Positron Emission Tomography (PET) is a medical imaging technique that reconstructs three-dimensional images of the body's metabolic and physiological processes using coincidence measurements of paired gamma-ray emissions, which result from positron decay of a radiotracer attached to a molecule of interest. Among the primary challenges in PET imaging are the extremely high and highly-varying noise levels in the acquired data, limited spatial resolution, high dynamic range, varying image contrasts, and the computational complexity of the forward model, compounded by the large size of raw data. The integration of Machine Learning (ML) into the PET data processing and image reconstruction pipeline has opened up numerous opportunities to improve both the quality of raw data and the resulting images, potentially leading to enhanced diagnostic accuracy. ML techniques in PET are applied at various stages of the imaging pipeline in research, including more accurate position, energy, and time-of-flight estimation from raw detector signals, as well as noise suppression and resolution enhancement during or after image reconstruction. However, implementing ML in PET imaging presents distinct challenges and potential pitfalls, which differ significantly from other modalities like MRI, where ML research is probably most mature. This talk provides an overview of recent and promising ML applications in PET imaging, while also addressing common pitfalls and challenges that may not be immediately apparent to those unfamiliar with PET-specific constraints. Related work is accessible via the DOI:10.2967/jnumed.121.262303.

Towards Accurate Quantitative Photoacoustic Tomography: Combining Model and Learning Based Algorithms

Andreas Hauptmann*
*University of Oulu, Finland

Abstract In photoacoustic tomography biological tissue is illuminated with a short pulse of near infrared light. The absorbed energy creates a local pressure increase that propagates through the tissue, governed by the acoustic wave equation, and we can measure the pressure wave on the boundary. From this measured time-series we first aim to reconstruct the initial pressure in the tissue, providing valuable information on local structures, such as microvasculature. Subsequently, it is possible to recover quantitative absorption and scattering values. Correct recovery of the optical parameters would provide valuable functional and biological information for medical purposes. In practice, solving both the acoustic and optical inverse problem comes with challenges. Starting from an often encountered limited-view geometry, restricting the measurement surface and resulting in a severely ill-posed linear inverse problem for the acoustic inversion. Reconstruction errors from the acoustic problem will naturally propagate when attempting to solve the optical problem. In this talk, we discuss the combination of model-based and data-driven reconstructions for both, the acoustic and optical inverse problem. We discuss conceptual differences between learning a reconstruction for linear and nonlinear inverse problems. In both cases, we consider the possibility to use approximate models to reduce computational complexity.

AI-guided maternal and fetal low field MRI

Jana Hutter*
*FAU Erlangen-Nüremberg, Germany

Abstract Complex motion patterns are a hallmark of all abdominal and pelvic imaging. These are amplified further in fetal MRI, where unpredictable fetal motion and fast fetal cardiac activity interact with the maternal motion and the needs of MR imaging. AI-guided MRI techniques allow to react and correct to motion. Rethinking the entire fetal examination by including real-time processing and planning has led to a self-driving fetal MRI scan allowing comprehensive assessment of fetal and maternal health. Combined with recently re-emerging low field MRI - carrying the essential benefits of lower costs, reduced distortion artifacts and longer T2 - it allows to democratize access to fetal MRI.

Neural Implicit k-space Representations

Daniel Rueckert* and Wenqi Huang*
*Technical University of Munich, Germany

Abstract Neural implicit k-space representations offer significant potential for reconstructing dynamic MRI at high spatial and temporal resolution. While existing methods bin acquired data from neighbouring time points to reconstruct individual time frames, neural implicit k-space representations can learn continuous and subject-specific k-space representations without the need for external training data. In this work we demonstrate the potential of this approach for the reconstruction of cardiac MR images and motion-compensated reconstruction. We evaluate different regularisation strategies for neural implicit k-space representations. Furthermore, we evaluate the effect of different k-space sampling schemes as well as different undersampling rates on learning neural implicit k-space representations that are suitable for image reconstruction. Associated articles are accessible via the DOIs: 10.1007/978-3-031-34048-2_42, 10.1007/978-3-031-72104-5_59.

Physics-Informed Deep Learning for Motion-Corrected Reconstruction of Quantitative Brain MRI

Hannah Eichhorn*†, Veronika Spieker*†, Kerstin Hammernik†, Elisa Saks†, Lina Felsner†, Kilian Weiss‡, Christine Preibisch†, and Julia Schnabel*†
*Helmholtz Munich, Germany, †Technical University of Munich, Germany, ‡Philips Healthcare, Germany

Abstract Quantitative MRI offers the advantage to estimate physical tissue properties from a series of qualitative images with varying imaging parameters. However, patient motion during image acquisition is a challenge for brain MRI in general, and quantitative MRI is particularly sensitive to motion due to its intrinsically long imaging times. To address this issue, we have developed PHIMO, a physics-informed learning-based motion correction method tailored to quantitative MRI. PHIMO leverages information from the quantitative parameter estimation process to detect motion-corrupted k-space lines in a self-supervised fashion for each subject individually. To subsequently perform motion correction, the detected motion-corrupted lines are excluded from a data-consistent learning-based reconstruction. We demonstrate the potential of PHIMO for the application of $T2^{\star}$ quantification from gradient echo MRI. This sequence is particularly sensitive to motion due to its sensitivity to magnetic field inhomogeneities, which are also affected by patient motion. A state-of-the-art technique for motion correction of $T2^{\star}$ maps requires redundant acquisition of the k-space center, which significantly increases the acquisition time. We show that PHIMO can detect and exclude intra-scan motion events and, thus, correct for severe motion artifacts. PHIMO outperforms a previously proposed motion correction method based on outlier-rejecting bootstrap aggregation, which confirms the importance of incorporating information from the $T2^{\star}$ decay in our physics loss. Moreover, PHIMO approaches the performance of the state-of-the-art motion correction method, while substantially reducing the acquisition time by over 40%, facilitating clinical applicability. Our code is available on GitHub. Associated article is available at the DOI: 10.1007/978-3-031-72104-5_54.

Optimizing learned unrolled networks for low-latency image reconstruction

Marc Vornehm*, Jens Wetzl†, Florian Fürnrohr*, Daniel Giese†, Rizwan Ahmad‡, and **Florian Knoll***,
*FAU Erlangen-Nüremberg, Germany, †Siemens Healthineers, Germany, ‡Ohio State University, USA

Abstract Learned unrolled neural networks for MR image reconstruction [1] have demonstrated high image quality on open benchmark datasets [2,3,4]. A key advantage is that the computationally expensive training occurs in advance of the actual patient imaging. At inference, running trained models is computationally lean, which them well-suited for interactive real-time MRI applications. The goal of this work is to investigate learned unrolled networks for MRI-guided cardiac interventions, where a maximum latency of 200 ms is required [5], and to analyze network design choices in the context of low-latency imaging. Our proposed architecture is based on our recent CineVN [6]. We trained all networks on fully sampled cine data from the OCMR dataset [7] using retrospective undersampling. For testing, we used 14 prospectively undersampled real-time cine series. To enable low-latency, we used a sliding window where the most recent frame was inferred using data from the latest 7 frames. For reference, we obtained CS reconstructions using BART [8]. We used NVIDIA A100 GPUs for all numerical experiments and measured the reconstruction time per frame, including data transfer to and from the GPU. Our results show that there is an inherent tradeoff between image quality and inference time. This is to be expected, since the computation time at inference is dominated by the size of the deep neural network. Nevertheless, we achieved an inference latency of only 69 ms per timeframe, which is acceptable for real-time interventions, while achieving an SSIM of 0.97 with respect to the reference CS reconstructions [9].
This research was supported by NIH/NIBIB grant R01EB029957 and NHR HPC resources under project b143dc. Related work are accessible via the DOIs: [1] 10.1002/mrm.26977, [2] 10.1148/ryai.2020190007, [3] 10.1002/mrm.28338, [4] 10.1109/TMI.2021.3075856, [5] 10.1002/jmri.25749, [6] 10.1002/mrm.30260, [7] 10.48550/arXiv.2008.03410, [8] 10.5281/zenodo.592960, [9] 10.58530/2024/0017.

The R2D2 Deep Neural Network Series Paradigm for Scalable Non-Cartesian Magnetic Resonance Imaging

Yiwei Chen*, Amir Aghabiglou*, Motahare Turki*, Shijie Chen*, Chao Tang*†, Ruud B. van Heeswijk‡, and **Yves Wiaux***,
*Heriot-Watt University, UK, †University of Edinburgh, UK, ‡Lausanne University Hospital, Switzerland

Abstract We introduce the R2D2 Deep Neural Network (DNN) series paradigm for fast scalable image reconstruction from highly-accelerated non-Cartesian k-space acquisition sequences in Magnetic Resonance Imaging (MRI). While unrolled architectures provide a robust image formation approach via data-consistency layers, embedding non-uniform FFT operators in a DNN can become impractical to train at large scale, e.g. in 2D MRI with a large number of coils, or for higher-dimensional imaging. Alternative Plug-and-Play approaches, alternating a learned denoiser blind to the measurement setting with a data-fidelity step, are not affected by this limitation but their highly-iterative nature implies slow reconstruction. To address this scalability challenge, we leverage the R2D2 paradigm recently introduced to enable ultra-fast reconstruction for large-scale Fourier imaging in radio astronomy [1]. R2D2's reconstruction is formed as a series of residual images iteratively estimated as outputs of DNNs taking the previous iteration's data residual as input. The method can be interpreted as a learned version of the Matching Pursuit algorithm. A series of R2D2 DNN modules were sequentially trained in a supervised manner on the fastMRI dataset and validated for 2D multi-coil MRI in simulation and on real data [2], targeting highly under-sampled radial k-space sampling. Results suggest that a series with only few DNNs achieves comparable image reconstruction quality to its unrolled incarnation R2D2-Net [1], whose training is however much less scalable. R2D2 also offers superior speed, precision, and robustness, over the state-of-the-art diffusion-based approach "Decomposed Diffusion Sampler" (DDS, [3]) specifically developed to address large-scale inverse imaging problems.
Related works are accessible via the DOIs: [1] 10.3847/1538-4365/ad46f5, [2] 10.23919/eusipco63174.2024.10714974, [3] 10.48550/arXiv.2303.05754.

Session 8

Medical Imaging in a Low-resource Setting

Organisers: **Andrew Webb** & **Florian Knoll**

Summary *This session will explore different challenges associated with medical imaging in low resource settings. MRI, X-ray/CT and ultrasound image requirements will be compared with one another, as well as issues of data transmission and interpretation. The different potential roles of AI in image acquisition, processing and diagnosis will be a common theme throughout.*

What practical role can deep learning play in sustainable low field MRI in low resource settings?

Andrew Webb*
*Leiden University Medical Center, The Netherlands

Abstract Recent advances in (mobile) low field MRI technology have shown immense promise in terms of producing diagnostic quality images for neuroimaging and whole-body applications at a fraction of the purchase, siting and maintenance costs of standard commercial MRI systems operating at much higher field strengths [1–3]. Some of these systems rely heavily on deep learning (DL) algorithms to improve image quality, whereas others are meticulously designed to produce high quality images with minimal signal processing. Nevertheless, exploitation of such technology is slow, and it is an open question as to what role DL can play in decreasing the barriers to expansion. Should it be used to make image contrast more familiar to what radiologists are used to, should it be used for image interpretation given the lack of radiologists in many low-income settings, can it be used to help scanning by general medical staff (task-shifting) or should it be used to design even lower cost and more forgiving systems which can be used for specific tasks in more challenging environments? This talk will review current state-of-the-art low field MRI technology and the image quality that is currently possible, and discuss how DL might be employed in each of the use cases outlined above. Related works are accessible at the DOIs: [1] 10.1126/science.adm7168, [2] 10.1101/2024.04.01.24305081, [3] 10.1038/s44222-023-00086-w.

Addressing Ecosystem Needs for Advancing MRI Access in Low and Middle-Income Settings

Sean Deoni*
*Bill & Melinda Gates Foundation, USA

Abstract Low and ultra-low field MRI (with magnetic field strengths less than 100mT) offers the potential to reshape health equity - enabling advanced medical imaging in settings that have traditionally lacked access [1]. This includes low- and middle-income settings where conventional high-field systems are cost-prohibitive and the basic infrastructure to support such systems is lacking. However, an inexpensive low-field MRI system will, by itself, only improve access with concurrent innovation in personnel training (including maintenance technicians, radiographers, and radiologists), data transmission, internet connectivity, access to stable electricity, and integration into clinical workflows. This talk will discuss the challenges of deploying low-field MRI systems across various global health settings in sub-Saharan Africa, India, Pakistan, and Bangladesh, including rural primary care centers and higher-resource urban tertiary care hospitals, and approaches to addressing training and limited infrastructure. In addition, we will discuss the development and optimization of imaging methods tailored explicitly to health conditions that are more prevalent in low and middle-income countries. Related work is accessible via the DOI: [1] 10.11604/pamj.2018.30.240.14000.

Computational Brain Imaging in Low-Resource Settings: From Automated Diagnostics to AI Hallucinations

Steven Schiff*
*Yale University, USA

Abstract Most people who need neurosurgical care have no access to brain imaging. The advent of low-field MRI opens a rapidly expanding frontier of image availability into health care settings where clinicians have had no prior experience in interpreting MRI, the signal-to-noise levels of the imaging is low, speciality radiologists are not available, and AI image enhancement techniques can generate misleading structural information. I will explore a range of efforts at computational imaging enhancement using learning libraries [1,2], super-resolution using structural priors [3], childhood brain growth [4,5], clinician acceptance of low-quality imagery [6], infection diagnostics [3], and the inherent nature of AI hallucinations from image libraries of structural pathologies [8]. These studies are supported by NIH 7R01AI145057, 5R01HD085853 and 1U01NS107486. Related works are accessible via the DOIs: [1] 10.1109/TBME.2017.2783305, [2] 10.3171/2021.7.PEDS21209, [3] 10.1109/TIP.2019.2942510, [4]

10.3171/2021.2.PEDS201006, [5] 10.1056/NEJMoa1707568, [6] 10.1016/j.nicl.2021.102896, [7] 10.1088/1741-2552/acd9ee, [8] 10.1016/j.nec.2024.05.010.

The Imaging Revolution: Convergence of Point-of-Care Testing, AI and Digital Health

Kristen DeStigter*
*University of Vermont, USA

Abstract The integration of point-of-care testing (POCT), artificial intelligence (AI), and digital health is poised to revolutionize access to imaging in low-resource settings. Drawing on 16 years of continuous experience as an NGO in sub-Saharan Africa (Imaging the World, Corp.), we have demonstrated that distributed clinical ultrasound diagnostics can be seamlessly incorporated into an existing tier-based healthcare system at the primary care level. Enhancing access to imaging care is achievable through a comprehensive package that includes sustainable technology, training, supervision, and quality management, addressing a broad spectrum of conditions via distributed, democratized radiology. For successful implementation and scalability, it is crucial to address all elements of technology, process, and particularly, culture. Transformative implementation must be context-specific and involve collaboration with policymakers and key stakeholders. The proven model of distributed, affordable, and reliable ultrasound services provides a blueprint for the democratization of other imaging technologies, representing a paradigm shift in diagnostic access for populations in resource-limited settings. Related works are accessible via the DOIs: [1] 10.1109/GHTC.2011.39, [2] 10.1371/journal.pone.0078450, [3] 10.1186/s12884-014-0424-9, [4] 10.1136/heartjnl-2018-313810, [5] 10.1186/s13244-021-01073-8, [6] 10.1371/journal.pone.0255918, [7] 10.1016/S0140-6736(21)00673-5.

Advances in Image Quality Transfer and application in low-resource settings

Daniel Alexander*
*University College London, UK

Abstract I will present recent advances in Image Quality Transfer [1,2,3] and related ideas such as SynthSR [4,5]. I will focus particularly on their potential in low-resource settings, their utility to enhance the utility of low-field MRI scanners, and more broadly their role in international efforts to democratize MRI expertise and capability. These techniques use machine learning to estimate high quality images, e.g., from a powerful experimental scanner, from lower quality images, e.g., acquired on a standard hospital scanner or low-field systems. I will cover the history and basic principles of these ideas, show some of the latest results, speculate about future opportunities, and describe some challenges and observations of implementing these ideas in LMIC scenarios. Related works are accessible via the DOIs: [1] 10.1016/j.neuroimage.2017.02.089, [2] 10.1016/j.neuroimage.2020.117366, [3] 10.1016/j.media.2023.102807, [4] 10.1016/j.neuroimage.2021.118206, [5] 10.1126/sciadv.add3607. Underpinning preprint is available on arXiv: 2404.05980.

Nonlinear tomographic reconstruction via nonsmooth optimization

Vasileios Charisopoulos* and Rebecca Willett*
*University of Chicago, USA

Abstract We study iterative signal reconstruction in computed tomography (CT), wherein measurements are produced by a linear transformation of the unknown signal followed by an exponential map. Approaches based on pre-processing the data with a log transform and then solving the resulting linear inverse problem are tempting since they are amenable to convex optimization methods; however, such methods perform poorly when the underlying image has high dynamic range, as in X-ray imaging of tissue with embedded metal. In this work, we focus on methods that solve the nonlinear inverse problem directly, thus avoiding the aforementioned pathology. In particular, we show that a suitably initialized subgradient method applied to a natural nonsmooth, nonconvex loss function produces iterates that converge to the unknown signal of interest at a geometric rate under the statistical model proposed by [1]. In particular, the computational and sample complexities of our method scale polynomially with the intensity of the

unknown image, while past methods exhibit exponential scaling. Our recovery program enjoys improved conditioning compared to the formulation proposed by the latter work, enabling faster iterative reconstruction from substantially fewer samples. Related work is accessible via the DOI: [1] 10.48550/arXiv.2310.03956. Underpinning preprint is available on arXiv:2407.12984.

Locality enables end-to-end supervised learning for Cryo-ET

Vinith Kishore*, Valentin Debarnot*, AmirEhsan Khorashadizadeh*, and Ivan Dokmanić*,
*University of Basel, Switzerland

Abstract Cryo-electron tomography (CryoET) is a powerful tool in structural biology for 3D visualization of cells and biological systems at resolutions sufficient to identify individual proteins in situ. The measurements are collected similar to standard computed tomography; frozen samples are tilted and exposed to an electron beam at various angles. However, due to the nature of the samples and the effects of electron exposure on it, we can only obtain limited and highly noisy measurements, making standard tomographic reconstruction methods inadequate. Current methods address this issue by applying post-processing techniques to denoise and recover the missing information using the noisy reconstructions. Among these, self-supervised methods are popular, using deep neural networks trained independently on each dataset either by splitting the measurements or enforcing rotational equivariances. In this work, we introduce the first supervised learning method that directly reconstructs the tomogram from the noisy measurements. The primary drawback of supervised learning is the need for paired datasets of measurements and reference volumes that accurately represent real data, which are lacking in cryoET. We address this issue by collecting a set of paired datasets of raw cryo-ET data with reconstructions from leading self-supervised methods. Secondly, we use local reconstruction networks, which require very less data and generalize well on different volumes in practice. The proposed approach has a low memory footprint which allows us to efficiently recover large volumes, enabling high-resolution 3D reconstructions. We also release the paired dataset for developing and benchmarking new CryoET reconstruction methods. Underpinning preprint is available on arXiv:2401.00816.

Merging the worlds of Cartesian and Non-Cartesian sampling for highly accelerated MRI

Chaithya Radhakrishna*†, Maxime Bertrait*†, Aurelien Massire‡, Alexandre Vignaud†, and Philippe Ciuciu*†
*Université Paris-Saclay, Inria, CEA, France, †Université Paris-Saclay, NeuroSpin, CEA, France, ‡Siemens Healthcare SAS, France

Abstract Accelerating MRI acquisitions is possible through compressed sensing with variable density sampling, which is most efficiently achieved using non-Cartesian imaging. In this regard, SPARKLING algorithm was introduced to optimize scanner-compliant sampling patterns that match a variable density, over-sampling the k-space center and tapering toward the periphery. However, conventional SPARKLING trajectories often oversample the center of k-space beyond the Nyquist rate, which is suboptimal as these extra samples can be dispatched to higher frequencies for improved image reconstruction. Furthermore, non-Cartesian sampling limits the direct use of established and clinically used parallel imaging acceleration techniques. To address these challenges, we propose a novel acquisition trajectory using affine constraint and a modified sampling density to end up with hybrid sampling pattern consisting of: Cartesian sampling at center of k-space at optimal Nyquist rate and non-Cartesian curves at the periphery to maximize k-space coverage. Each trajectory passes through the center of k-space as a Cartesian line, enabling compatibility with established Cartesian acceleration techniques in the center like GRAPPA and CAIPIRINHA. Further, this design also allows for additional acceleration through undersampling at higher frequencies using non-Cartesian and flexible portions of the same sampling trajectory. Finally, we propose a novel dual-reconstruction scheme: the k-space center is reconstructed with GRAPPA, while the fully reconstructed center and periphery data are processed using a conventional FISTA-based iterative reconstruction. Applied to anatomical MRI, our method achieves $10-15\times$ acceleration with minimal image quality loss, enabling 1mm isotropic whole-brain 3D imaging in under a minute. Future work will integrate echo-planar imaging, expanding these gains to functional and diffusion MRI. The approach extends on our recent work accessible via the DOI: 10.1002/mrm.29702.

Robust Plug-and-Play Methods for Highly Accelerated Non-Cartesian MRI Reconstruction

Pierre-Antoine Comby*, Benjamin Lapostolle*, **Matthieu Terris***, and Philippe Ciuciu*
*Université Paris-Saclay, Inria, CEA, France

Abstract Magnetic Resonance Imaging (MRI) reconstruction at high acceleration factors remains challenging because of the ill-posed nature of the inverse problem. Although compressed sensing (CS) methods offer robustness across acquisition settings, they struggle with reconstruction quality at high acceleration factors. Deep learning approaches have improved reconstruction quality, but are prone to overfitting and hallucination effects when acquisition parameters vary between training and test times. Plug-and-Play (PnP) methods, which replace CS priors with deep neural network denoisers, offer a promising middle ground, but face challenges in MRI reconstruction due to proximal gradient descent instabilities and the lack of clean training data. This work presents two key contributions to address these challenges. First, we introduce an unsupervised preprocessing pipeline that generates clean, noiseless complex MRI signals from multicoil k-space data. The pipeline combines virtual coil combination with a Neighbor2Neighbor denoising network trained without ground truth data. Second, we propose an annealed Half-Quadratic Splitting (HQS) algorithm incorporating preconditioning techniques to address stability issues in PnP methods. We evaluated our approach in the fastMRI dataset using spiral sampling patterns at various acceleration factors ($4\times$ to $16\times$). Results demonstrate that our method achieves superior reconstruction quality compared to existing PnP approaches and maintains robust performance at high acceleration factors where other algorithms struggle. The preconditioning significantly improves the image quality, while the HQS scheme enhances stability compared to standard proximal gradient descent methods. Notably, our approach generalizes well across acceleration factors without requiring retraining, outperforming supervised deep learning methods at high acceleration factors. Underpinning preprint is accessible on arXiv:2411.01955.

Session 9

Medical Imaging in a High-resource Setting

Organiser: Daniel Sodickson

Summary *If we had all the hardware power, compute resources and data analysis techniques we could dream of, what could we achieve in medical imaging? What fundamental scientific or clinical questions could we answer? This session will explore advanced imaging capabilities which are currently being enabled or may soon be enabled by emerging hardware platforms, software tools, and datasets. Speakers will share a focus on new information that can be gleaned from MRI, PET, CT, ultrasound, and other imaging or sensing modalities. They will also comment on which developments may eventually translate to, or otherwise influence, the low-resource settings to be explored in more detail in a companion session.*

Putting medical imaging in context: connected scanners of the future

Daniel Sodickson*
*New York University, USA

Abstract If we had all the hardware power, computer resources, and data analysis techniques we could dream of, what would we do differently in medical imaging? We would certainly aim to image sharper, faster, and better. One feature of medical imaging which has changed little during its remarkable advancements to date, however, is the way we generally treat each imaging session or acquisition sequence in isolation. Though we have mastered dynamic imaging with high frame rates, we still tend to consider each distinct interaction with a medical imaging device as a single snapshot. Health, on the other hand, is dynamic. In many cases, what physicians are most interested in is change. This lecture will describe some early attempts to put medical imaging in context. I will explore two specific examples, drawn from MRI but applicable to multiple medical imaging modalities as well as to astronomical imaging. The first involves use of side information for image reconstruction. Side information can include other contrasts, other modalities, prior scans, or even medical record text. It can also include representations learned from tasks other than image reconstruction, or from foundation models. I will introduce new ways to disambiguate image content in ill-posed or ill-conditioned cases, while protecting against hallucination. The second example involves use of longitudinal imaging for risk assessment. I will consider intelligent imaging protocols that meter the extent of data acquisition based on disease risk computed in real time from undersampled data together with prior imaging data. I will also explore the use of incomplete or low-quality imaging data for longitudinal monitoring of health. Such uses of context require extensive computational power and a high degree of coordination. They are therefore well suited, at least at first, to high-resource settings. That said, they may also enable new, high-impact uses of accessible imaging devices in low-resource settings.

Revealing microstructure with MRI: new information from advanced gradient technology

Ileana Jelescu*
*Lausanne University Hospital, Switzerland

Abstract Like a kaleidoscope, the MRI acquisition can be adjusted to yield a stunning variety of information. Among other features, the MRI scanner can be turned into an in vivo microscope, able to quantify tissue features at the cellular level, while providing such detailed information over an entire organ, in living humans; in other words: a microscope uniquely suited for longitudinal studies of healthy and patient populations. The microstructure information is encoded using diffusion. In a diffusion MRI experiment, the mean squared displacement of water molecules in biological tissue is a few microns – a distance characteristic of the cellular scale – and their random walk is largely governed by interactions with the cellular architecture. In other words, diffusion MRI encodes information about micron-scale features of the tissue that we otherwise cannot access directly given the coarse MRI spatial resolution (0.1 – 1 mm). Developing quantitative techniques to decode this information and map specific biophysical parameters of tissue microstructure ('Microstructural MRI') has been the driving force of our research field [1]. The sensitization of the MRI signal to diffusion is achieved by applying magnetic field gradients. The stronger and faster the gradients, the shorter length scales we can probe. Thus, advanced gradient technology (e.g., on small-bore preclinical MRI systems, or human 'Connectom' scanners) has truly opened new avenues for Microstructural MRI. In this talk, I will show how advanced gradient technology, combined with innovative MRI sequence design and analytical tissue modeling, have enabled the reliable quantification of microstructure features in brain white and gray matter [2,3]. More recently, 'dynamic' cell microstructure fluctuations associated with brain activity are also being explored as a novel and specific functional MRI contrast [4]. Finally, I will discuss what further advancements in gradient technology could bring to our field, and how the information that is gleaned from high resource settings can help or augment the interpretation of data acquired in lower resource settings such as hospital scanners. Associated articles are available at the DOIs: [1] 10.1016/j.jneumeth.2020.108861, [2] 10.7554/elife.49855, [3] 10.1162/imag_a_00104, [4] 10.1101/2024.10.01.615823.

Emerging horizons in ultrahigh field MRI

Lucio Frydman*
*Weizmann Institute of Science, Israel

Abstract Magnets are THE key to nuclear magnetic resonance imaging (MRI). Sufficiently strong magnets could enable imaging not just the water in living organisms, but also the molecules that support the chemistry and biology underlying Life. Even if constrained to water observations, stronger magnets could improve the spatial resolution and introduce new kinds of contrasts that escape lower field observations, thereby enabling–as stronger telescopes do–new discoveries. This talk will rely on my experience as part of the "High Mag" panel that, convened in 2023/4 by the US National Academies, reported on the kind of opportunities that would open up by stronger fields. Starting with a review of ongoing achievements in NMR and MRI magnet and conductor developments, I will describe realistic recommendations suggested that should be pursued by US funding agencies in collaboration with industry; these include an all-superconducting 40 T NMR system; a 28 T magnet for small animal imaging; a 14+ T large bore magnet with applications to human MRI. I will then focus on the applications-and foremost on the supra-linear benefits–that would arise from these ultra-strong magnets. As will be shown, even these ultra-high fields would still leave frontiers in the understanding of human biology and the human brain –as well as valuable potential capabilities for the early detection of diseases–beyond the realm of the possible.

Beat Pilot Tone: Wireless Radio-Frequency Motion Sensing in MRI at Arbitrary Frequencies

Michael Lustig*
*University of California, Berkeley, USA

Abstract Motion in Magnetic Resonance Imaging (MRI) scans causes image corruption and remains a barrier to clinical imaging. We propose Beat Pilot Tone (BPT), a simple serendipitous system exploitation that turns any MRI receiver chain into a radio frequency (RF) motion sensing system that can operate at arbitrary frequencies (up to several GHz). Our contact-free system can be implemented on any MRI scanner regardless of field strength. Through electromagnetic field simulations and experiments, we explain BPT's novel mechanism: two or more transmitted RF tones form motion-modulated standing wave patterns that are sensed by the same receiver coil arrays used for MR imaging. These waves are incidentally mixed by intermodulation and digitized simultaneously to the MRI data. BPT achieves an order of magnitude greater sensitivity to motion than other methods in detecting and separating common motion types (respiratory, bulk, cardiac, and head motion) in volunteers. Moreover, BPT offers tunable sensitivity to motion based on the transmit frequencies; at microwave frequencies, BPT can detect millimeter-scale vibrations (ballistocardiograms). With multiple antennas and frequency-multiplexing, BPT can operate as a multiple-input multiple-output (MIMO) system. BPT significantly expands the capability of any MRI system, paving the way toward multi-modality, motion-robust, and simultaneous RF and MR imaging. Underpinning article is available at the DOI: 10.1002/mrm.30150.

Pushing the Limits of CT: New Clinical Information from Multi-Energy and Photon Counting Systems

Aaron Sodickson*
*Brigham and Women's Hospital, USA

Abstract CT technology has undergone exciting advances in recent years, with proven clinical benefits throughout the body. Conventional CT creates high resolution cross-sectional images displaying the relative X-ray absorption of the contained tissues. Multi-Energy CT (also called Dual-Energy CT) goes a step further by processing x-ray absorption data from distinct x-ray energy ranges in order to differentiate materials based on their varied energy-dependent x-ray absorption behaviors, adding color to previously grey-scale images. These techniques have created new opportunities to improve detection and characterization of pathology, with associated opportunities to reduce radiation exposure to patients, and downstream imaging utilization. Photon-Counting CT is the latest new major step forward in CT technology, accomplished by translating detector technology designed for high energy particle physics into the medical arena. These detector systems replace the traditional scintillator material /photodiode detector

array with a semiconductor layer capable of individually counting and quantifying the energy of each incident X-ray. These scanners have unlocked further advances in high resolution imaging, and promise to further advance multi-energy material characterization applications. This talk will highlight key features of these advanced CT scanners, and demonstrate some of the established and developing clinical benefits enabled by these technology breakthroughs. Underpinning article is available at the DOI: 10.1007/s10140-020-01785-2.

MRI for Assessing Phase Aberration Corrections for Transcranial Ultrasound

Kim Butts Pauly*, Morteza Mohammadjavadi*, Kasra Naftchi-Ardebili*, and Gary Glover*
*Stanford University, USA

Abstract Transcranial ultrasound has many compelling applications, including diagnostic imaging, thermal ablation, selective drug delivery, and brain stimulation. These applications, however, are challenged by the heterogeneity of the human skull. Variations in thickness and composition give rise to phase aberrations, attenuation, and absorption. Physics-based beam simulations in combination with MRI or CT can be used to find phase aberration correction (PAC) terms. In the simplest implementation, short TE MRI provides an input to these simulations, and time reversal or ray tracing are used to find the PAC terms. Because physics-based simulations are time consuming, these may be done only in advance of the study. During the study, MRI can be used to assess the focal position, with MR thermometry or MR acoustic radiation force imaging (MR-ARFI) with and without a temperature rise, respectively. The focal position can be easily modified with phased array transducers. In addition, MR-ARFI can potentially be used to assess the intensity reaching the focus. In a large animal model, we demonstrated that the MR-ARFI displacement was correlated to the neuromodulatory effect we were trying to achieve, a reduction in visual evoked potential. But, to date, there is no means to update the PAC terms in real time. We had a proof of concept study in 2014 that demonstrated that beam simulations could be used to find the PAC terms that were most consistent with the acquired MR-ARFI image. Since then, our research group has developed an approach to beam simulations and PAC based on machine learning, reducing simulation time by a factor of 1000 while providing the PAC terms simultaneously. With the speed afforded by the ML approach, we can now imagine using the MR-ARFI images to provide feedback on the PAC in real-time or used as inputs to the simulation, alongside images of the skull. Underpinning article is available at the DOI:10.1118/1.4865778.

Kinetic Model-Informed Generative AI for Multi-Tracer 4D PET Reconstruction

Andrew Reader*, George Webber*, and Bolin Pan*
*King's College London, UK

Abstract Generative AI is the most recent frontier in PET image reconstruction, and we have recently proposed Poisson-likelihood scheduling for sampling from the posterior distribution via diffusion models [1]. This provides a principled approach to integrating generative AI within PET image reconstruction with fewer hyper-parameters than other methods, even if computationally demanding. We have also extended our diffusion model based prior to harness higher signal to noise ratio training data from multiple co-registered subject scans when training the diffusion model [2]. In tandem with these recent developments we have proposed kinetic-model informed deep learning for separating single-tracer images from multiplexed PET acquisitions [3]. The latter development allows dynamic PET images of two or more tracers scanned within the same imaging session to be separated with greater fidelity and with fewer training examples compared to purely AI data-driven approaches. We envision the integration of a multi-subject fully 3D diffusion model based image reconstruction into our multi-tracer kinetic-model based deep learning pipeline for multi-tracer separation. The resulting framework would be able to provide single-tracer 4D PET reconstruction and kinetic parametric maps with credibility intervals, whether the acquired PET data are single or multi-tracer. Associated articles are accessible via the DOIs: [1] 10.1109/NSS/MIC/RTSD57108.2024.10657861, [2] 10.1109/NSS/MIC/RTSD57108.2024.10657446. Underpinning article is available at the DOI: [3] 10.1186/s40658-024-00660-0.

Evaluating the Quantum Fourier Transform for Magnetic Resonance Image Reconstructions

Martin Blaimer*, Felix Breuer*, Kilian Dremel*, Anastasia Papadaki*, Thomas Lang*, Dimitri Prjamkov*, Markus Firsching*, Anja Heim*, Stefan Kasperl*, and Theobald Fuchs*,
*Fraunhofer IIS, Germany

Abstract Magnetic resonance imaging (MRI) systems perform data acquisition with up to 128 receiver channels thereby allowing time-resolved measurements with high spatial resolution. Alternative data processing solutions including quantum computing (QC) are being explored to handle the ever-increasing amounts of raw data [1]. QC makes use of principles of quantum mechanics and exposes unique properties such as quantum entanglement to potentially enable a massive parallel processing of data encoded in a quantum state. Our purpose is to explore the potential of the quantum Fourier-transform (QFT) for MRI reconstructions. Inspired by previous work in radio astronomy [2], we demonstrate an efficient representation of multi-channel MRI data by the quantum probability image encoding (QPIE) [3] and present initial QFT reconstructions using QASM simulations. Finally, we explore the performance of the QFT on a 156-qubit device from IBM. Our results indicate an efficient data encoding using QPIE. The processing of an MRI dataset of size $16 \times 256 \times 256$ (channels $\times$ lines $\times$ rows) requires 20 entangled qubits. Due to the quantum nature, the qubit states must be prepared and measured multiple times to obtain statistically robust results. We found empirically that, $\mathcal{O}(N^2)$ measurements are necessary for sufficient quality of $N \times N$ images. Therefore, quantum advantage may be difficult to achieve even though the QFT requires only $\mathcal{O}(\log_2 N \times \log_2 N)$ operations when QPIE is employed [2]. Noisy hardware and limitations in the number of entangled qubits limit the application to small images. However, the rapidly evolving hardware with reduced noise and error rates show promise for future applications, especially when additional processing steps are performed in the quantum domain. This work is supported by BayQS. Related works are accessible via the DOIs: [1] 10.48550/arXiv.2004.02036, [2] 10.1016/j.ascom.2024.100796, [3] 10.1103/PhysRevX.7.031041.

(Self-)supervised reconstruction for high spatial and temporal resolution in DCE-MRI for breast cancer

Erik Gösche*, Kathrin Flaβkamp†, Sungheon Gene Kim‡, and Florian Knoll*,
*FAU Erlangen-Nüremberg, Germany, †Saarland University, Germany, ‡Weill Cornell Medicine, USA

Abstract Dynamic Contrast Enhanced (DCE) MRI is the most accurate tool for diagnosing breast cancer [1,2]. It works by injecting a contrast agent that shortens T1 time in surrounding tissue, enhancing the T1-weighted MRI signal. The increased uptake of contrast agent in tumor regions, caused by angiogenic leaky blood vessels, makes DCE-MRI highly effective for cancer detection. However, the ultimate potential of DCE-MRI has not been fully utilized in clinical practice due to limitations in data acquisition methods, which cannot provide both high 3D spatial resolution ($< 1mm^3$) and high temporal resolution ($< 10s$/frame) simultaneously. In this work, we propose novel reconstruction models to preserve high 3D spatial and temporal resolution. In the absence of a ground truth, this task becomes especially challenging. Therefore, we are pursuing two different strategies: supervised training of reconstruction models using simulated data, and self-supervised training approaches that do not rely on labeled datasets. We perform this analysis on simulated data generated by a digital reference object (DRO) toolkit [3] and on real in-vivo breast cancer MRI exams which are part of the fastMRI Breast dataset [4]. Our results show that supervised models trained on simulated data can significantly improve the spatial and temporal resolution trade-offs compared to conventional reconstruction methods. Our model achieves near-isotropic 3D spatial resolution of $1mm^3$ and a temporal resolution of approximately $8s$/frame, without significant loss of image quality. The self-supervised approach shows promise in adapting to real clinical data, maintaining image fidelity without labeled training data. This study builds on our groundwork published at [5]. This research was supported by HPC resources by the Erlangen National High Performance Computing Center under project b143dc. Related works are accessible via the DOIs: [1] 10.1002/jmri.26654, [2] 10.1097/RLI.0b013e3181b4c127, [3] 10.1002/mrm.30152, [4] 10.48550/arXiv.2406.05270, [5] 10.1002/mrm.30260.

Ultra-High-Field Diffusion-weighted MRI Reconstruction with Spatial-angular Low-rank Regularization for Rigid Motion Correction

Annika Hofmann[*], Zhengguo Tan[†], Patrick Liebig[‡], Frederik Bernd Laun[*], and Florian Knoll[*]
[*]FAU Erlangen-Nüremberg, Germany, [†]University of Michigan, USA, [‡]Siemens Healthineers, Germany

Abstract Diffusion MRI is essential for examining brain microstructures with intricate fiber arrangements, though averaging over large voxels can introduce partial volume effects [1]. While increasing spatial resolution could address this, it also reduces the signal-to-noise ratio (SNR). Ultra-high field MRI, such as 7T, enhances SNR. But the shorter echo times at 7T can lead to blurring artifacts in single-shot EPI acquisitions. To mitigate these artifacts, multishot EPI is often used. Motion during scanning induces spatially varying phase variations in the signal between the individual shots. A common approach to counteract these variations involves using low-resolution navigator scans to estimate and correct phase shifts [2]. However, significant rigid motion between navigators and imaging scans poses challenge, which even state-of-the-art reconstruction methods, such as MUSE [3], struggle to compensate effectively, resulting in substantial image degradation. In this work, we present how the joint reconstruction for shift-encoded navigator-based interleaved echo planar imaging method [4], can enable the reconstruction of diffusion-weighted images even in cases where navigator phase corrections fail. This technique leverages local low-rank regularization by organizing spatial-angular matrices from the acquired diffusion-weighted images and applying Singular Value Thresholding to enforce low-rank conditions [5]. We show that JETS-NAViEPI can successfully recover diffusion-weighted images despite severe rigid-body motion and fully corrupted navigator data; situations in which conventional methods would yield nondiagnostic images. Related works are accessible via the DOIs: [1] 10.1038/s41597-021-00904-z, [2] 10.1002/mrm.22024, [3] 10.1016/j.neuroimage.2013.01.047, [4] 10.1162/imag_a_00085, [5] 10.1002/mrm.28025.

Exact Parameter Identification in PET Pharmacokinetic Modeling

Martin Holler[*], Erion Morina[*], and Georg Schramm[†]
[*]University of Graz, Austria, [†]KU Leuven, Belgium

Abstract In quantitative dynamic positron emission tomography (PET), one aims to reconstruct kinetic tissue parameters based on a time series of images that provide the concentration of the PET tracer over time. This reconstruction is challenging, as it corresponds to a non-linear parameter identification problem. Further, it is usually based on an estimate of the arterial input function obtained via arterial blood sampling and analysis, which is a costly and time consuming process. In our work, we analyze the identifiability of kinetic tissue parameters from multi-region measurement data in quantitative PET imaging solely based on image data. We show that, for the frequently used two-tissue compartment model and under reasonable assumptions, it is possible to uniquely identify these parameters without the need of additional concentration measurements from blood samples. Our identifiability result, which holds in the idealized, noiseless scenario, indicates that costly concentration measurements from blood samples in quantitative PET imaging can be avoided in principle. The connection to noisy measurement data is made via a consistency result, showing that exact reconstruction is maintained in the vanishing noise limit. Numerical experiments with a regularization approach are further carried out to support these analytic results in an application example. This contribution presents results of the work accessible via the DOI:10.1088/1361-6560/ad539e, and its extensions.