Proceedings of the International BASP Frontiers Conference 2023

February 5-10, 2023

Villars-sur-Ollon, Switzerland

BASP Frontiers Conference 2023

Chair

Prof. Yves Wiaux, Heriot-Watt University, UK

Steering Committee

Dr. Katherine Bouman, California Institute of Technology, USA
Dr. Audrey Bowden, Vanderbilt School of Engineering, USA
Dr. Ulugbek Kamilov, Washington University in Saint-Louis, USA
Prof. Florian Knoll, Friedrich-Alexander University Erlangen Nuremberg, Germany
Prof. Anna Scaife, University of Manchester, UK
Prof. Julia Schnabel, Technical University of Munich, Germany
Prof. Philip Schniter, Ohio State University, USA
Prof. Jean-Luc Starck, CEA-Saclay, France
Prof. Yves Wiaux (Chair), Heriot-Watt University, UK

Local Organizing Committee

Mr. Taylor Chu, Heriot-Watt University, UK
Dr. Arwa Dabbech, Heriot-Watt University, UK
Mr. Chao Tang, University of Edinburgh, UK
Prof. Jean-Philippe Thiran, EPFL, Switzerland
Prof. Yves Wiaux (Chair), Heriot-Watt University, UK

Foreword

Astronomy and biomedical sciences find common roots in their need to process acquired data into interpretable signals or images. In these applications of signal processing, the complexity of the data to be acquired and processed is constantly increasing, thus challenging signal processing methodologies. 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 selected topics in astronomy and biomedical sciences around common challenges for signal processing, with a core focus on computational imaging.

12 years down the line ... Building on the success of the first five events in 2011, 2013, 2015, 2017, and 2019, (Covid-canceled in 2021), the 2023 conference will open its floor to synergetic discussions on state-the-of-art signal processing methodologies (from optimization to machine learning and sampling, from data-driven to physics-informed, from algorithm to hardware, etc.) and their application to data acquisition, image reconstruction, image analysis, and other scientific computing challenges. 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 in 2011. A variety of other imaging and more general sensing modalities have enriched the conference since then, and will also be represented, from astronomy to medicine, from low resource to extreme scale, etc.

Following our tradition, BASP Frontiers 2023 will gather no less than ninety participants in a very nice resort in the Swiss Alps named Villars-sur-Ollon, close to Lausanne and Lake Geneva. We believe that the most fruitful discussions often take place after the sessions themselves, during breakfast, coffee breaks, or dinner, on the terrace or on the ski slopes. We hope that the winter atmosphere will further promote creativity.

On Behalf of BASP Frontiers steering committee

Prof. Y. Wiaux

Conference Chair

Contents

Conference introductory talk

Multiscale Models for Data Generation and Inverse Problems with Deep Networks

Stéphane Mallat*
*Collège de France, France

Abstract Generating signals and images or solving inverse problems requires to build models of high-dimensional data distributions. Spectacular results have been obtained by deep neural networks, particularly with score diffusion models. A major issue is to understand how come such models can be estimated in high dimension. We show that many signals and images have local dependencies across scales, allowing to circumvent the curse of dimensionality. Similarly to Wilson renormalisation group in physics, we demonstrate that probability distributions can be factorised into conditional probabilities of wavelet coefficients across scales, which are local. It defines low-dimensional models of complex physical fields, but also of structured images. Applications are shown for image generation, noise removal and super-resolution recovery. We shall also relate these questions to classification and regression problems with deep neural networks.

Part I

Computational Imaging Theory

Session 1

Generative Inference and Calibration

Organizers: **Ivan Dokmanić** & **Philip Schniter**

Summary *Generative models are experiencing a second youth in imaging and scientific inference. New ideas include injective models for sampling high-dimensional posteriors, theoretical advances on statistical, approximation-theoretic, and topological questions, generating continuous functions which dovetail with the downstream PDE solvers, and creative uses of generative models to probe performance limits of inference systems. The session brings together prominent researchers spear-heading these exciting new directions.*

Gen-Alpha Generative Models for Imaging

Ivan Dokmanić*, **Amirehsan Khorashadizadeh***, **Konik Kothari**†, **Valentin Debarnot***, **Anadi Chaman**†, and **Maarten de Hoop**‡
*University of Basel, Switzerland, †University of Illinois at Urbana-Champaign, USA, ‡Rice University, USA

Abstract Generative models have become so good that they are a serious alternative to established "deep" approaches to imaging. In fact, on tasks like posterior sampling and uncertainty quantification, they often surpass them in flexibility, quality, and speed, with strong results even on highly ill-posed, nonlinear problems. In this preludial talk I will give a brief overview of the state of the generative art and then introduce our work on topological, approximation-theoretic, and analytic aspects of these emerging models. On the theory side, I will show that the so-called injective flows accurately model data distributions on almost arbitrary manifolds, subject to curious topological obstructions. On the practical side, I will show how we used these insights to build architectures that are particularly well-suited for regularization of ill-posed imaging problems. Finally, I will discuss how old ideas from signal processing inspired a data-driven interpolator which turns any old generative model into a differentiable, continuous functional generator at arbitrary scale. Applications abound in PDE-based problems and calibration.

Continuous Generative Neural Networks

Giovanni S. Alberti*, **Matteo Santacesaria***, and **Silvia Sciutto***
*University of Genoa Italy

Abstract In this talk, I will present and study Continuous Generative Neural Networks (CGNNs), namely, generative models in the continuous setting. The architecture is inspired by DCGAN, with one fully connected layer, several convolutional layers and nonlinear activation functions. In the continuous L^2 setting, the dimensions of the spaces of each layer are replaced by the scales of a multiresolution analysis of a compactly supported wavelet. We present conditions on the convolutional filters and on the nonlinearity that guarantee that a CGNN is injective. This theory finds applications to inverse problems, and allows for deriving Lipschitz stability estimates for (possibly nonlinear) infinite-dimensional inverse problems with unknowns belonging to the manifold generated by a CGNN. Preprint available on ArXiv (2205.14627).

Deep and Shallow Generative Models of Images

Yair Weiss*
*Hebrew University, Israel

Abstract Tremendous progress in the quality of generated image samples has been recently achieved by deep generative models such as Generative Adversarial Networks and Deep Diffusion Models. These models can now generate stunningly photorealistic images but when measured using the traditional goals of learning a prior over images for computer vision, their performance may still be worse than much simpler, classical methods. Furthermore, the conditions under which deep generative models will yield high-quality samples are poorly understood and small changes in the training protocol may dramatically decrease performance. In this talk, I will describe two shallow generative models that are better (in some senses) than the deep methods. The first [1] is the humble Gaussian Mixture Model which generates blurry images but allows efficient probabilistic inference. The second [2] is a direct method that attempts to generate images with with a desired distribution over patches at different scales. While these shallow methods are better in some senses than the deep methods, they are much worse in other aspects and I will also discuss what I think is needed for shallow methods to close the gap. Related work available on ArXiv ([1]1805.12462, and [2] 2203.11862).

Learning to Bound: A Generative Cramér-Rao Bound

Hai Victor Habi*, **Hagit Messer***, and **Yoram Bresler**†
*Tel Aviv University, Israel, †University of Illinois at Urbana-Champaign, USA

Abstract The Cramér-Rao bound (CRB) is a classical lower bound on the variance of any unbiased parameter estimator. The CRB enables to understand the fundamental limits in a parameter estimation problem, regardless of the algorithm used. It has been widely used, including in localization by sensor arrays, and in biomedical imaging. However, to compute the CRB, requires a precise and explicit statistical model for the data. In many applications, such a model is not available. Examples include device-specific noise statistics such as in image sensors, or source localization with unknown propagation channel characteristics or uncalibrated sensors. Instead, this work introduces a novel approach to approximate the CRB using data-driven methods. This approach is based on deep generative models, which have have shown state-of-the-art performance in modeling complex, high-dimensional distributions. We use a normalizing flow generative model to learn the measurement distribution from data. This is used, in turn, to generate samples of the gradient of the log-likelihood and obtain, as an empirical mean, an estimate of the Fisher Information Matrix (FIM). Inverting, we obtain the Generative Cramér-Rao Bound (GCRB). To assess the approximation quality by the GCRB, we provide theoretical bounds on its error relative to the true CRB in terms of two components: (i) imperfect learning, characterized by the total variation distance and the Fisher relative information between the true and learned distributions; and (ii) sampling error due use of an empirical mean to estimate the FIM using a finite number of samples generated by the normalizing flow model. We provide conditions under which both error components can be driven to zero. We validate the GCRB by simple examples of parameter estimation, and demonstrate its unique value by two examples from image processing: image denoising, and edge position detection, in the presence of realistic, camera-specific noise. Preprint available on ArXiv (2203.03695).

Interferometric Phase Image Estimation Using Importance Sampling

Mario Figueiredo*, and **Joshin P. Krishnan**†
*University of Lisbon, Portugal, †SimulaMet, Norway

Abstract Interferometric phase image estimation is a problem that arises in several imaging modalities, such as synthetic aperture radar and sonar, magnetic resonance imaging (MRI), optical interferometric imaging, and diffraction tomography. Formally, the problem can be formulated as that of denoising/estimating modulo-2π phase images from sinusoidal (2π-periodic) noisy observations, which makes it a challenging inverse problem. We describe a novel approach to tackle this problem, assuming the availability of an external data set, based on importance sampling to approximate a minimum risk Bayes estimate of the underlying complex image. The proposed method takes a class-specific data set of clean patches and clusters it using a mixture of circular symmetric Gaussian distributions. For each noisy patch, the closest cluster is identified and used to obtain a shift-invariant importance sampling estimator. Both the clustering mechanism and the estimation technique are developed for complex-valued signals by taking into account patch shift invariance, which is an important property for efficient phase denoising. The effectiveness of the proposed algorithm is shown via experiments on a semi-real interferometric data set constructed using human face images and on medical imaging applications involving real MRI data. It is shown that, in most of the experiments, the proposed estimator shows better results than the previous state-of-the-art methods, yielding a minimum improvement of 1 dB in peak signal-to-noise ratio (PSNR) for low to high noise levels. This talk is based on work by Joshin Krishnan, José B. Dias (1960-2020), and myself, accessible via the DOI 10.1109/ACCESS.2020.3021178.

Deep Invertible Approximation of Topologically Rich Maps between Manifolds

Maarten de Hoop[*]
[*]Rice University, USA

Abstract Neural networks based on topological data analysis (TDA) use tools such as persistent homology to learn topological signatures of data and stabilize training but may not be universal approximators or have stable inverses. Other architectures universally approximate data distributions on submanifolds but only when the latter are given by a single chart, making them unable to learn maps that change topology. By exploiting the topological parallels between locally biLipschitz maps, covering spaces, and local homeomorphisms, and by using universal approximation arguments from machine learning, we find that a novel network of the form $\mathcal{T} \circ p \circ \mathcal{E}$, where $\mathcal{E}$ is an injective network, p a fixed coordinate projection, and $\mathcal{T}$ a bijective network, is a universal approximator of local diffeomorphisms between compact smooth submanifolds embedded in $\mathbb{R}^n$. We emphasize the case when the target map changes topology. Further, we find that by constraining the projection p, multivalued inversions of our networks can be computed without sacrificing universality. As an application, we show that learning a group invariant function with unknown group action naturally reduces to the question of learning local diffeomorphisms for finite groups. Our theory permits us to recover orbits of the group action. Finally, our analysis informs the choice of topologically expressive starting spaces in generative problems. Joint research with M. Puthawala, M. Lassas, I. Dokmanić and P. Pankka. Preprint available on ArXiv (2210.00577).

Diffusion-based Models meet Image Priors

Erich Kobler[*]
[*]University Hospital Bonn, Germany

Abstract Recently, diffusion-based models have gained increasing interest from various research fields due to their superb performance on numerous tasks such as text, audio, or image synthesis. These models build upon the principle of score matching proposed by Hyvörinen 2005, which relies on learning a model,Äôs parameters by matching its gradient w.r.t. the input in a logarithmic domain to the counterpart of the data distribution. To extend this approach to unknown data distributions, Vincent 2011 introduced denoising score matching, which used to revert stochastic differential equations if conditioned on diffusion time. To deepen the understanding of this learning paradigm, we study denoising score matching applied to learning statistics of natural images; therefore, image patches are perturbed by additive isotropic Gaussian noise with variance t. This results in a one-parameter family of corresponding smooth probability densities on the space of image patches. In the negative log domain, this family of densities becomes gradually more convex. This observation suggests that denoising-score-matching-based models implement a graduated non-convexity scheme (Blake and Zisserman 2003), enabling efficient inference or sampling. Taking this effect into consideration, we introduce a novel prior modeling low-level statistic of natural images, which is an important pillar for various image reconstruction problems. Our prior operates on the space of image patches and generalizes the Fields-of-Experts prior (Roth and Black 2005). To account for arbitrary scale-space parameters t, the prior implements smooth potential and activation functions using 2d splines. After learning this prior using denoising score matching, it can be applied to various reconstruction problems. It turns out that our formulation generalizes several recently proposed approaches including variational networks and score-based generative models within a unified formulation. More details available at bit.ly/3FddR1b.

Validation Diagnostics for SBI algorithms based on Normalizing Flows

Julia Linhart[*], **Alexandre Gramfort**[*], and **Pedro Rodrigues**[*]
[*]INRIA, France

Abstract Simulation based inference (SBI) has the potential to revolutionize experimental sciences, as it allows to invert arbitrary complex mechanistic models with intractable likelihoods. Building on the recent trend of new deep generative models known as Normalizing Flows (NF), new SBI-algorithms can now efficiently accommodate arbitrary complex and high-dimensional data distributions. The development of appropriate validation methods however has fallen behind. Indeed, most of the existing metrics either require access to the true posterior distribution, or fail to provide theoretical guarantees on the consistency of the inferred approximation beyond the one-dimensional setting. This work proposes easy to interpret validation diagnostics for multi-dimensional conditional (posterior) density estimators based on NF. It also offers theoretical guarantees based on results of local consistency. The proposed workflow can be used to check, analyze and guarantee consistent behavior of the estimator. The method is illustrated with a challenging example that involves tightly coupled parameters in the context of computational neuroscience. This work should help the design of better specified models or drive the development of novel SBI-algorithms, hence allowing to build up trust on their ability to address important questions in experimental science. These results were presented at the 2022 NeurIPS workshop "Machine Learning and the Physical Science".

Stable deep MRI reconstruction using Generative Priors

Martin Zach[*], **Florian Knoll**[†], and **Thomas Pock**[*]
[*]Graz University of Technology, Austria, [†]Friedrich-Alexander-Universität, Germany

Abstract Data-driven approaches recently achieved remarkable success in medical image reconstruction, but integration into clinical routine remains challenging due to a lack of generalizability and interpretability. Existing approaches usually require high-quality data-image pairs for training, but such data is not easily available for any imaging protocol and the reconstruction quality can quickly degrade even if only minor changes are made to the protocol. In addition, data-driven methods may create artificial features that can influence the clinicians decision-making. This is unacceptable if the clinician is unaware of the uncertainty associated with the reconstruction. In this paper, we address these challenges in a unified framework based on generative image priors. We propose a novel deep neural network based regularizer which is trained in an unsupervised setting on reference images without requiring any data-image pairs. After training, the regularizer can be used as part of a classical variational approach in combination with any acquisition protocols and shows stable behaviour even if the test data deviates significantly from the training data. Furthermore, our probabilistic interpretation provides a distribution of reconstructions and hence allows uncertainty quantification. We demonstrate our approach on magnetic resonance imaging, where results show competitive performance with state-of-the-art end-to-end deep learning methods, while preserving the flexibility of the acquisition protocol and allowing for uncertainty quantification. To solve parallel imaging problems, we propose a novel algorithm to jointly estimate the image and the spatially varying coil sensitivities. Preprint available on ArXiv (2210.13834).

A Regularized Conditional GAN for Posterior Sampling in Inverse Problems

Matthew Bendel[*], **Rizwan Ahmad**[*], and **Phillip Schniter**[*]
[*]Ohio State University, USA

Abstract In inverse problems, one seeks to reconstruct an image from incomplete and/or degraded measurements. Such problems arise in magnetic resonance imaging (MRI), computed tomography, deblurring, superresolution, inpainting, and other applications. It is often the case that many image hypotheses are consistent with both the measurements and prior information, and so the goal is not to recover a single "best" hypothesis but rather to explore the space of probable hypotheses, i.e., to sample from the posterior distribution. In this work, we propose a regularized conditional Wasserstein GAN that can generate dozens of high-quality posterior samples per second. Using quantitative evaluation metrics like conditional Frechet inception distance, we demonstrate that our method produces state-of-the-art posterior samples in both multicoil MRI and inpainting applications. Preprint available on ArXiv (2210.13389).

Session 2

Large-scale Optimization for Computational Imaging

Organizers: **Gitta Kutyniok** & **Ulugbek Kamilov**

Summary *Large-scale optimization problems arise in a variety of imaging tasks. Examples include dictionary learning, low-rank matrix recovery, blind deconvolution, and phase retrieval. Conventional approaches for solving many of these optimization problems involve designing algorithms that can effectively leverage a wide-variety of structural constraints. The session provides an excellent opportunity for the wider signal processing and imaging community to come together and share recent developments, open challenges, and future directions in large-scale optimization methods suitable for analyzing imaging data.*

The Future of Optimization for Imaging

Gitta Kutyniok*
*Ludwig Maximilian University of Munich, Germany

Abstract In this introductory talk to the session on "Large-scale Optimization for Computational imaging", we will first survey current optimization methods, also related to the subsequent presentations. We will then discuss fundamental limitations of optimization approaches, focusing in particular on the fact that current algorithms are predominantly run on digital hardware. As we will show, this can cause serious problems from a computability viewpoint, which can be provably resolved by using analogue hardware such as neuromorphic computing or quantum computing.

Efficient Lip-1 spline networks for convergent PnP image reconstruction

Michael Unser*, **Stanislas Ducotterd***, and **Pakshal Bohra***
*École Polytechnique Féderale de Lausanne, Switzerland

Abstract Neural networks are constructed via the composition of simple modules: affine transformations (e.g. convolutions) and pointwise non-linearities. PnP schemes for image reconstruction make use of such convolutional neural networks for the so-called denoising step. The condition for convergence is that the denoising CNN should be non-expansive. It is achieved by imposing a tight Lipchitz-1 constraint on each module; for instance, by using ReLU non-linearities and by spectrally normalizing each linear layer. The downside of this simple normalization approach is that the resulting network generally loses expressivity. Concretely, this means that PnP with current Lip-1 CNN fall short of reaching the best possible (state-of-the-art) image quality, which is price to pay for the consistency and stability of the reconstruction. The approach that is investigated in this work is to replace the ReLU activations by more expressive trainable non-linearities, subject to the Lip-1 constraint. The foundation for our approach is a representer theorem for the design of deep neural networks with Lipchitz-constrained free-form activations subject to TV^2 regularization to favor "simple" neuronal responses. It states that the global optimum of this constrained functional optimization problem can be achieved with a configuration where each neuron is a linear spline with a small number of adaptive knots (break points). We then describe a B-spline framework for the efficient implementation and training of such Lip-1 spline networks. Finally, we apply our framework to the reconstruction of magnetic resonance images and show that it compares favorably with other existing Lip-1 neuronal architectures. Related work available on ArXiv (2210.16222, 1802.09210)

Optimization of deep neural networks under privacy constraints

Georgios Kaissis*, **Alexander Ziller***, and **Daniel Rueckert***
*Technische Universität München, Germany

Abstract The utilization of deep neural networks in domains with sensitive data represents a challenging trade-off between data utilization and data protection. Modern privacy-enhancing technologies like differential privacy are able to bridge that gap and alleviate privacy concerns. However, they introduce undesirable privacy-utility trade-offs, both from a computational and an accuracy point-of-view. To tackle these trade-offs, innovative techniques for optimization must be developed. This talk will discuss such techniques by focusing on two distinct areas: On one hand, computational techniques which mitigate the decreased efficiency of private optimization. On the other hand, architecture design techniques which facilitate convergence in the presence of biased and noisy gradients necessary for differential privacy. Moreover, it will provide a concise introduction to the theory of differential privacy from a signal processing point of view. The works which will be discussed here are accessible at nature.com/articles/s41598-021-93030-0, and on ArXiv (2210.00053, 2209.04338, 2205.04095, 2203.00324).

Hybrid Learning to Sense and Solve for Computational Imaging

M. Salman Asif*
*University of California - Riverside, USA

Abstract Deep learning-based methods for solving inverse problems have gained considerable attention in the last few years. Deep learning principles have also been used to represent and optimize sensing and reconstruction modules as differentiable networks. In principle, these networks can be learned in an end-to-end manner using data alone, but performance of such systems degrades with small changes in data. A hybrid approach is often more stable in which we can fix some parts of the network using prior knowledge of the forward model or physics priors and learn some parts of the network using training data. In this talk, I will present some of our recent work on the hybrid approaches learning to sense and solve. I will first discuss 3D lensless imaging using measurements captured with a programmable mask. We formulate the 3D recovery problem as a parallel set of simple least-squares problems. With part of the recovery algorithm fixed, we learn the mask pattern to optimize the 3D reconstruction and learn a refinement network to identify and remove artifacts in the reconstruction. Then I will discuss an affine phase retrieval problem in which we represent the iterative phase retrieval algorithm as an unrolled (differentiable) network and optimize over the additive reference signal. The learned reference provides significant performance enhancement compared to standard Fourier phase retrieval. Preprints available on ArXiv (2007.14621, 2108.07966).

Deep Model-Based Architectures for Inverse Problems under Mismatched Priors

Shirin Shoushtari*, **Jiaming Liu***, **Yuyang Hu***, and **Ulugbek S. Kamilov***
*Washington University in St. Louis, USA

Abstract There is a growing interest in deep model-based architectures (DMBAs) for solving imaging inverse problems by combining physical measurement models and learned image priors specified using convolutional neural nets (CNNs). For example, well-known frameworks for designing DMBAs include plug-and-play priors (PnP), deep unfolding (DU), and deep equilibrium models (DEQ). While the empirical performance and theoretical properties of DMBAs have been widely investigated, the existing work has primarily focused on their performance when the desired image prior is known exactly. This work addresses this gap by providing new theoretical and numerical insights into DMBAs under mismatched CNN priors. Mismatched priors arise naturally when there is a distribution shift between training and testing data, for example, due to test images being from a different distribution than images used for training the CNN prior. They also arise when the CNN prior used for inference is an approximation of some desired statistical estimator (MAP or MMSE). Our theoretical analysis provides explicit error bounds on the solution due to the mismatched CNN priors under a set of clearly specified assumptions. Our numerical results compare the empirical performance of DMBAs under realistic distribution shifts and approximate statistical estimators. Additional technical details can be found in the preprint available at wustl-cig.github.io/mismatchwww.

Signal processing with rank-one projections and optical quadratic random sketches

Laurent Jacques[*], **Remi Delogne**[*], **Vincent Schellekens**[†], and **Laurent Daudet**[‡]
[*]Université catholique de Louvain, Belgium, [†]CEA List - LFIC, France, [‡]LightOn, France

Abstract Rank-one projections (ROP) of matrices and quadratic random sketching of signals are at the heart of recent processing and imaging applications, such as covariance matrix estimation, phase retrieval, and lensless interferometry by speckle illumination. In this context, it is often desirable to perform (approximate) processing techniques (such as pattern detection or classification) directly in this sketched domain without accessing the original data. While such a possibility was previously demonstrated for linear random sketching methods, as those used in standard compressive sensing applications, we show how to estimate simple signal processing tasks (such as deducing frequencies and local variations in signals or images) directly from the quadratic random projections provided by an optical processing unit (or OPU), a device allowing for superfast computing of random data sketching in the optical domain. Mathematically, this is made possible by showing that the ROP measurement operator approximately preserves, with a controlled distortion, specific comparisons between two signals in the sketched domain provided one projects one signal sketch onto the sign of the other signal sketch. In other words, under certain conditions, the ROP operator satisfies a generalized sign product embedding (SPE) property already considered in the field of one-bit compressive sensing. Leveraging our approach and techniques, we finally report several conclusive experiments of image classification directly operated in the quadratic sketched domain recorded by an actual OPU.

Block Delayed Majorize-Minimize Subspace Algorithm for Large Scale Image Restoration

Jean-Baptiste Fest[*] and **Emilie Chouzenoux**[†‡]
[*]INRIA, France, [†]CentraleSupélec, France, [‡]Université Paris-Est Marne la Vallée, France

Abstract The ever-increasing complexity and precision of modern acquisition devices as those met in large scale image processing often implies the resolution of large scale optimization problems. In such a context, the use of distributed algorithms has been developed as a response to alleviate the high-dependency of computation performances regarding the dimension of the initial problem. When the cost function is differentiable but non-necessarily convex function, an efficient family of optimization methods relies on the concept of majorization-minimization (MM). In this work, we introduce the Block Delayed MM Memory Gradient (BD3MG) scheme which provides a distributed implementation of the MM method from Chouzenoux et al. SIIMS 2013. We show the convergence of the method under mild assumptions, using recent tools from non-smooth analysis. The performance of the algorithm is illustrated on the restoration of synthetic and real 3D images degraded by depth-variant 3D blur, arising in multiphoton microscopy. More details on our contribution can be found at hal.science/hal-03764677.

A constrained optimization-based approach for multiphoton microscopy restoration

Ségoléne Martin*, **Jean-Christophe Pesquet***, **Emilie Chouzenoux***†, and **Claire Lefort**‡
*CentraleSupelec, France, †Université Paris-Est Marne la Vallée, France, ‡XLIM Limoges, France

Abstract In the last three decades, multiphoton microscopy (MPM) has received much interest in the field of biomedical visualization. This optical solution is a first choice for observing cells in action, providing contactless subcellular resolution deep in living tissues. Nevertheless, the volumes acquired are usually strongly corrupted by multiplicative noise and instrumental image distortion. Sophisticated approaches based on adaptive optics can mitigate these problems, but such physical strategies are expensive and technically complex. An alternative strategy is to rely on computational approaches from image processing to improve the visual quality of MPM images. However, most computational strategies require the adjustment of some hyper-parameters, especially a regularization parameter, which is tedious and time-consuming given the large volume of data one must cope with. In this talk, we describe a new hyperparameter-free restoration method for MPM imaging. Our optimization-based approach avoids tuning a regularization parameter by using both a constraint on the data-fit term and statistical properties of the noise. In addition, the method can deal with heteroscedastic Gaussian noise (i.e., noise variance dependent on the pixel's intensity), which is a reasonable approximation of Poisson noise. Based on our recent work [1], we provide an efficient and reliable algorithm to solve the constrained minimization problem. Finally, we display some restoration results obtained on real multiphoton volumes of a mouse's muscle. Related work accessible via the DOI [1] 10.1007/s10851-022-01112-z.

Unfolded proximal denoising network for versatile plug-and-play algorithm

Audrey Repetti*
*Heriot-Watt University, UK

Abstract Since early 2000's, proximal methods have become state-of-the-art to solve minimization problems. In the context of imaging inverse problems, they are usually used to find maximum a posteriori (MAP) estimates of unknown images from degraded measurements. During the last decade, hybrid proximal algorithms involving neural networks (NNs) have emerged. Two main frameworks can be distinguished: unfolded NNs, where the architecture of the NN mimics a finite number of iterations of an optimization algorithm, and plug-and-play (PnP) algorithms, where some non-linear steps (e.g., proximity operators) of iterative optimization algorithms are replaced by NNs. Unfolded NNs usually incorporate the measurement operator in the learning process, which can be prohibitive for applications with non-fixed measurement operators. PnP do not have this drawback but involved NNs still depend on the underlying statistical models (e.g., higher noise level on the measurements requires stronger denoisers). In addition, although both frameworks are grounded in optimization theory, they do not directly benefit from the same theoretical guarantees. Often, either the NN structures must be constrained, or some regularizations must be incorporated during the training process, making it more difficult in practice. In this work, we propose a PnP algorithm based on forward- backward (FB) iterations, where the learned denoiser is an unfolded NN based on proximal splitting techniques. In particular, this NN is built to mimic a Gaussian denoiser from a MAP viewpoint. This allows us to introduce a regularization parameter in the model to tune the regularization strength, similarly to standard variational approaches. This has the advantage of making the learned NN more adaptive to a variety of inverse problem statistical models, without need for training the NN for different noise levels. Preliminary publication (EUSIPCO 2022) available via the DOI 10.23919/EUSIPCO55093.2022.9909564.

Session 3

Modern Regularization

Organizers: Thomas Pock & Philip Schniter

Summary *The availability of expressive regularizers is a very important component in solving ill-posed inverse problems in imaging. In recent years, hand-designed regularizers have been gradually replaced by data-driven ones. Provided sufficient training data is available, it is nowadays possible to learn tailored regularizers for a certain problem class. This usually leads to a huge increase in reconstruction quality, but the learned regularizers are usually much harder to analyze and it is much harder to give guarantees on convergence behavior, generalization ability, or reconstruction error. In this session, we will present and discuss the latest methods, techniques and applications in this cutting-edge field of research.*

An Introduction to Modern Regularization

Thomas Pock[*]
[*]Graz University of Technology, Austria

Abstract. The availability of expressive regularizers is a very important brick in solving ill-posed inverse problems in imaging. In recent years, hand-designed regularizers have been gradually replaced by data-driven ones. Provided sufficient training data is available, it is nowadays possible to learn tailored regularizers for a certain problem class. This usually leads to a huge increase in reconstruction quality, but the learned regularizers are usually much harder to analyze and it is much harder to give guarantees on convergence behavior, generalization ability, or reconstruction error. In this introductory talk, I will present and compare different regularization techniques, from classical regularization methods based on convex energies to fully learned regularizations based on deep learning techniques. A special focus will be on generative models, i.e. models that allow to generate samples from their distribution. These models have the advantage that they can be learned in an unsupervised setting, are more stable to changes in data acquisition, and allow for some quantification of uncertainty. We show applications for image reconstruction mainly in the medical context, i.e. computed tomography and magnetic resonance imaging.

Solving inverse problems in imaging via learned regularizers and learned image features

Stacey E. Levine[*]
[*]Duquesne University, USA

Abstract Effective data driven approaches for learning regularizers for inverse problems in imaging have greatly increased in recent years. This is due to exploding computational power which has enabled models to learn salient image features through the use of deep neural networks that depend on a large number of learned parameters. Still, powerful regularizers for solving inverse problems in imaging based on physics and functional semi-norms have provided theoretically sound models for solving a large range of imaging problems, leading to predictable and interpretable solution properties. In this work we consider a hybrid approach, in that regularizers are learned for both image data as well as for structured but learned image features, both of which are inspired by model based image geometry. This approach shows improvement over comparable model-based deep learning image denoising algorithms both with respect to image quality metrics as well as the preservation of fine features. We also study how these deep neural network architectures for both learned regularizers as well as the learned features affect the accuracy of the image reconstruction, ultimately decreasing reconstruction error.

Risk Control for Online Learning Models

Yaniv Romano[*]
[*]Technion, Israel

Abstract Modern machine learning algorithms have achieved remarkable performance in a myriad of problems, and are increasingly used to make impactful decisions in high-stakes applications such as autonomous driving and healthcare diagnostics. The use of data-driven algorithms in such applications is exciting yet alarming: these methods are extremely complex, often brittle, and notoriously hard to analyze or interpret. Naturally, concerns have been raised about the reliability of the output of such machines. This talk focuses on making reliable predictions in an online setting, in which the underlying data distribution can drastically–and even adversarially–shift over time. We will introduce statistical tools that can be wrapped around any online "black-box" machine learning model to provide valid and informative uncertainty estimates.

Generative regularization approaches in imaging - a function space perspective

Martin Holler[*], and **Andreas Habring**[*]
[*]University of Graz, Austria

Abstract Classical regularization approaches for inverse problems in imaging often comprise a natural modeling of images in function space. The most famous example here is probably the Total Variation functional and the corresponding space of functions of Bounded Variation. This perspective on images has the advantage that questions such as well-posedness of the regularization approach or regularity of recovered images can be answered independently of particular discretizations. With the rise of neural-network- and learning-based regularization approaches for imaging, classical methods are becoming increasingly replaced by new techniques. This also raises new challenges, as the mathematical understanding of neural-network- or learning-based regularization approaches, in particular in view of function space modeling, still lacks significantly behind what is known in the classical theory. The purpose of this talk is to address these challenges from two perspectives. First, we introduce a variational model for image regularization that follows the same principles as the famous deep image prior ‚Äì but within a simpler, function-space modeling. The latter not only yields results comparable to those of the deep image prior with a reduced number of parameters, but also enables us to obtain precise statements about the regularity of images that can be recovered. In a second part of the talk, we will then discuss a general regularity analysis for images generated by convolutional-neural-network architectures. As we will show, depending on the the type of weight-penalization using during training, and depending on the network depth, the regularity of images obtained with such approaches does not match the classical modeling of images as piece-wise smooth functions. As practical consequence, this suggest to refrain from basic L2 regularization of network weights in case of images being the network output. Relevant work accessible via the DOIs: 10.48550/arXiv.2204.10588, and 10.1137/21M1414978.

Critical point regularization with non-convex regularizers

Markus Haltmeier[*] and **Daniel Obmann**[*]
[*]University of Innsbruck, Austria

Abstract Regularization is one of the key elements for the stable, efficient and accurate solution of inverse and imaging problems. The central tool in this context is variational minimization, also known as Tikhonov regularization. In recent years, many such methods have been developed, for example using regularizers defined by neural networks as penalty terms in variational methods. One of the basic assumptions in the stability and convergence analysis of these methods is the ability to find global minimizers. However, such an assumption is often not achievable when the regularizer is a black box or non-convex, which makes finding global minimizers of the involved Tikhonov functional a difficult task. Instead, standard minimization methods are used, which generally only guarantee that a critical point of the Tikhonov functional is found. To address this issue, in this contribution we study the stability and convergence properties of critical points of Tikhonov functionals with possibly nonconvex regularizers. To this end, we introduce a new concept of relative subdifferentiability and study its fundamental properties. Based on this concept, we develop a convergence analysis assuming relative subdifferentiability of the regularizer. For the case where the noise level tends to zero, we derive a limiting problem representing first-order optimality conditions for a related constrained optimization problem. Finally, we present numerical simulations that support our theoretical findings and highlight the need for the type of analysis presented in this work. Preprint available on ArXiv (2207.14480).

Adaptive Sparse Phase Recovery for Efficient Reconstruction of Holographic Lens-Free Images

René Vidal*, **Carolina Pacheco***, and **Benjamin D. Haeffele***
*Johns Hopkins University, USA

Abstract In-line lensless digital holography has great potential in multiple applications; however, reconstructing high-quality images from a single recorded hologram is challenging due to the loss of phase information. Typical reconstruction methods are based on solving a regularized inverse problem and work well under suitable image priors, but they are extremely sensitive to mismatches between the forward model and the actual imaging system. This paper aims to improve the robustness of such algorithms by introducing the adaptive sparse reconstruction method, ASR, which learns a properly constrained point spread function (PSF) directly from data, as opposed to solely relying on physics-based approximations of it. ASR jointly performs holographic reconstruction, PSF estimation, and phase retrieval in an unsupervised way by maximizing the sparsity of the reconstructed images. Like traditional methods, ASR uses the image formation model along with a sparsity prior, which, unlike recent deep learning approaches, allows for unsupervised reconstruction with as little as one sample. Experimental results in synthetic and real data show the advantages of ASR over traditional reconstruction methods, especially in cases where the theoretical PSF does not match that of the actual system. Related work accessible at the DOIs: 10.1007/978-3-319-66185-8_13, and 10.1364/OE.458360.

Statistical Sparsity, Non-Convexity, and Smoothness, All in One Place: The Cauchy Penalty

Alin Achim*, **Luca Calatroni** †, **Serena Morigi**‡, and **Gabriele Scrivanti**⋆
*University of Bristol, UK, †CNRS, UCA, INRIA, France, ‡University of Bologna, Italy, ⋆Centrale Supélec, France

Abstract We describe a smooth, non-convex and content-adaptive regularization model for the reconstruction of high-resolution Optical Coherence Tomography (OCT) images from low-resolution measurements. We follow a sparse-representation approach, whereby sparsity is modeled with respect to a suitable representation learned from high-resolution OCT data. To do so, we employ a pre-learned representation tailored to model α-stable statistics in the non-Gaussian case, i.e. $\alpha < 2$. The image reconstruction problem renders here particularly challenging due to the high level of noise degradation and to the heterogeneity of the data at hand. Hence, as a regulariser we propose a separable Cauchy-type penalty. To favor adaptivity to image contents, we introduce a space-variant modeling, by which the local degree of non-convexity given by the local Cauchy shape parameter is estimated via maximum likelihood. This strategy is shown to produce an adaptive map of parameters, which self-tunes the regularization strength depending on the type of local information observed. In order to exploit the overall smoothness of the model considered, we opted for a variant of the well-known Broy-den-Fletcher-Goldfarb-Shanno (BFGS) algorithm where a cautious update of the descent direction is used to adapt to the different convex vs. non-convex regimes. Several numerical tests show gains in computational efficiency. We show the advantages of having an adaptive regularization model with respect to a fixed one and compare our results with the ones obtained by means of popular non-smooth convex or non-convex regularizers frequently employed to enforce sparsity. We also show how the reconstructed data can be successfully employed for OCT image segmentation. Finally, we investigate the use of Cauchy-based loss functions as components within variational autoencoders and illustrate their benefits when techniques relying on deep unfolding are employed. Related publication available at the DOI 10.2139/ssrn.4047898.

Modified Loss Functions for Improved Plug-and-Play Methods

Danica B. Fliss[*], **Robert Nowak**[*], and **Willem Marais**[*]
[*]University of Wisconsin - Madison, USA

Abstract Plug-and-play (PnP) methods are a popular approach for image reconstruction. These methods mimic proximal gradient algorithms, but replace the formal proximal step with an off-the-shelf denoiser (e.g., BM3D or a deep neural denoiser). PnP methods are attractive because they leverage powerful, general-purpose denoisers, but this can also be a weakness in certain cases. For instance, striping gain factors occur with some satellite data due to pixel aggregation differences and is corrected by multiplying different sections of the image by different gain values. To correct for striping would distort the noise and cause poor results using black-box denoisers, but to denoise before correcting the stripes may give the impression of underlying signal structures that do not exist and lead to reconstruction artifacts. Plug-and-play methods help to combat this issue by iteratively correcting for stripes and denoising incrementally. Nonetheless, our analysis shows that ideally we would like to locally scale the level of denoising in proportion to each of striping gain, but it is difficult or impossible to incorporate this modification into black-box denoisers. To address this challenge, we show that instead the loss function can be appropriately modified to achieve the desired effect, with no modification needed to the black-box denoiser. We demonstrate our new PnP approach with the novel modified loss function on real satellite imagery using BM3D and achieve promising results. Empirically we compare our method to the traditional PnP method, using wavelet shrinkage and BM3D as our denoisers, and see our method outperforms the traditional with gain correcting reconstruction problems like that in our satellite data. For deblurring tasks, we see similar improvements in accuracy. We also observe that our new method is more stable with respect to tuning parameters than traditional PnP methods.

Deep regularization via bi-level optimization with adversarial regularization

Srijith S. Nair[*] and **Phillip Schniter**[*]
[*]Ohio State University, USA

Abstract In the variational approach to solving imaging inverse problems, one recovers the unknown image from noisy measurements by minimizing the sum of a data-fidelity term and a regularization term. Selecting the data-fidelity term is usually straightforward, but designing the regularization term is a long-standing, challenging problem. Ideally, we want that the regularization i) guarantees feasible optimization and ii) promotes recovered images with high accuracy. Like several recent works, we propose to construct the regularization term using a deep neural network and leverage auto-differentiation when minimizing the variational cost, a framework we refer to as "deep regularization". Inspired by the classical work of Tappen et al., we propose to train our deep regularizer using a bi-level approach, which directly rewards accuracy in the solution of the variation optimization problem. Our primary contribution is to use adversarial regularization during training to ensure the convexity and/or smoothness of the deep regularizer. Image recovery experiments demonstrate the advantages of our approach.

Deep network series for large-scale high-dynamic range imaging

Amir Aghabiglou[*], **Matthieu Terris**[*], **Adrian Jackson**[†], and **Yves Wiaux**[*]
[*]Heriot-Watt University, UK, [†]University of Edinburgh, UK

Abstract We propose a new approach for large-scale high-dynamic range computational imaging. Deep Neural Networks (DNNs) trained end-to-end can solve linear inverse imaging problems almost instantaneously. While unfolded architectures provide necessary robustness to variations of the measurement setting, embedding large-scale measurement operators in DNN architectures is impractical. Alternative Plug-and-Play (PnP) approaches, where the denoising DNNs are blind to the measurement setting, have proven effective to address scalability and high-dynamic range challenges, but rely on highly iterative algorithms. We propose a residual DNN series approach, where the reconstructed image is built as a sum of residual images progressively increasing the dynamic range, and estimated iteratively by DNNs taking the back-projected data residual of the previous iteration as input. We demonstrate on simulations for radio-astronomical imaging that a series of only few terms provides a high-dynamic range reconstruction of similar quality to state-of-the-art (optimization and) PnP approaches, at a fraction of the cost. Preprint accessible on ArXiv (2210.16060).

Part II

Astronomical Imaging

Session 4

Astronomical Imaging with Optical Telescopes: From Thirty-meter Observatories to Future Space Missions

Organizers: **He Sun** & **Katie Bouman**

Summary *Large optical telescopes have a rich history of pushing the boundaries of imaging technology and astronomical science. With a series of revolutionary optical telescopes launched recently or planned to be launched soon, in the next decade we expect to discover astronomical phenomena that will redefine our view of the universe. These telescopes will open windows to studying galaxies in the early universe, Earth-like planets around distant stars, and unknown discoveries beyond our imagination. This session reviews recent development of computational imaging techniques in large space- or ground-based optical observatories. Topics include adaptive optics for future space telescopes (e.g. RST) and 30m level ground telescopes (e.g. E-ELT), data processing and early results from JWST, and machine learning methods for gravitational lensing and exoplanet study.*

Measuring Properties of Imaged Exoplanets

Jason Wang*
*Northwestern University, USA

Abstract Imaging exoplanets, planets orbiting other stars, requires removing the bright glare of the star to reveal faint exoplanets. Often the starlight subtraction is done with principal component analysis (PCA) or similar algorithms. The process of removing starlight often involves distorting the planet signal and leaving behind spatially and spectrally correlated noise. Distortions in the planet signal make measurements of the brightness and flux biased. To overcome this, we can forward model the impact of PCA-based algorithms on the planet light. To extract out positions and fluxes, we also fit for the spatial covariance in the residual noise by modeling it as a correlated Gaussian process. This had led to the most precise measurements of imaged exoplanet properties from single-dish telescopes [1]. To better measure the atmospheric spectra of exoplanets, we have developed new instrumentation to feed the light of the planet into a spectrograph with sufficient spectral resolution to resolve individual absorption lines. This allows us to spectrally disentangle planet light from residual starlight, with the potential to reach the precision set by photon noise. We demonstrate this technique on-sky, detecting multiple exoplanets and reaching down to detector thermal noise limits, although the thermal noise was notably high [2]. Related work accessible via the DOIs: [1] 10.3847/0004-6256/152/4/97, and [2] 10.3847/1538-3881/ac1349.

Machine learning for direct exoplanet imaging

Olivier Absil*, **Carles Cantero Mitjans***, and **Marc Van Droogenbroeck** *
*University of Liège, Belgium

Abstract Directly imaging faint exoplanets around bright stars is a formidable challenge, due to the large flux ratio and small angular separation between the two objects. From the ground, this task is furthermore complicated by the effects of atmospheric turbulence, which create a bright field of speckles around the star in short-exposure images. In order to reduce the brightness of the speckle field, high-contrast imaging (HCI) relies on advanced instrumentation, including adaptive optics and stellar coronagraphy. This is however generally not enough to enable exoplanet detection, for which an appropriate combination of observing strategies and image processing techniques is key. In this talk, I will briefly review the state of the art in image processing for HCI, and the main lessons learned so far. I will then present our recent attempt to reformulate the detection of faint point sources in HCI data sets as a supervised machine learning problem (Gonzalez et al. 2018), which led to the inception of the SODINN algorithm, based on a combination of convolutional neural networks and long-short term memory networks. The rest of the talk will follow our ride along the hype cycle, exploring successively (i) the peak of inflated expectation, when we obtained a first set of very promising ROC curves, (ii) the trough of disillusionment, when our SODINN entry to the exoplanet imaging data challenge showed a very large number of false positives, and (iii) the slope of enlightenment, which we are currently climbing thanks to a more in-depth comprehension of the noise distribution in HCI data sets. I will conclude this talk by showing the latest results obtained with an upgraded version of SODINN. Related work available on ArXiv (1712.02841).

Toward Milli-Arcsecond Imaging: Distributed Diffractive Optics and Learned Reconstruction

Farzad Kamalabadi[*]
[*]University of Illinois at Urbana-Champaign, USA

Abstract Conventional spectral imaging systems for space remote sensing rely on reflective optics which are not naturally scalable, as the desired increase in angular resolution requires larger apertures which cannot be manufactured with sufficient accuracy to attain the diffraction limit. In this work, I describe a scalable ultra-high-resolution imaging concept based on diffractive optics and computational image formation methods. The wavelength-dependent focusing nature of diffractive lenses requires long baselines demanding distributed telescope configurations attained with precision formation flying of multiple spacecraft, and sophisticated computational imaging and machine learning algorithms to solve the multi-frame deblurring inverse problem resulting from multiple exposures at different focal lengths. I will describe the fundamental theory and principles as well as the necessary technological breakthroughs for such imaging systems and the ongoing efforts toward demonstrating such experiments from space by a multi-university consortium in collaboration with NASA and under sponsorship by US NSF. Related work available at the DOI 10.1109/TCI.2021.3075349.

Astronomical Data Fusion

Peter M. Melchior[*]
[*]Princeton University, USA

Abstract Astronomical image analysis has traditionally been done piecemeal, using one image from one telescope in one filter band. While apparently successful, this approach leads to shortcomings whenever the analysis of that single image is incomplete. And it misses the opportunities provided by large multi-band surveys and the existence of complementary observing programs with substantial overlap on the sky. I will present a framework that unifies data analysis across multiple different data sources, from space-borne imagers to ground-based spectrographs, and discuss the challenges that arise in astronomical data fusion. I will present several ongoing projects to tackle these problems through constrained optimization as well as generative modeling with autoencoders, normalizing flows, and diffusion models.

Space Starlight Suppression Technology Demonstration: The Nancy Grace Roman Space Telescope Coronagraph

Marie Ygouf[*]
[*]NASA Jet Propulsion Laboratory, USA

Abstract After the James Webb Space Telescope (JWST), NASA's next flagship astrophysics mission is the ambitious Nancy Grace Roman Space Telescope currently on track for a 2027 launch. The Roman Space Telescope Coronagraph Instrument will be the first high-performance stellar coronagraph using active wavefront control for deep starlight suppression in space, providing unprecedented levels of contrast, spatial resolution, and sensitivity for astronomical observations in the optical. During its Technology Demonstration phase, the Roman Coronagraph will resolve the signal of an exoplanet via photometry and spectroscopy and directly image and measure the polarization of disks. Future flagship mission concepts aim to characterize Earth analogues with visible light flux ratios of 10-10, and the Roman Coronagraph is a critical intermediate step toward that goal, with predicted capability of 10-9. Here, we introduce the ideas of adaptive optics and coronagraph design, present the coronagraph's capability as well as some anticipated results from its technology demonstration.

PolyCLEAN: A Polyatomic Frank-Wolfe Algorithm for Scalable Interferometric Imaging

Adrian Jarret[*], **Julien Fageot**[*], and **Matthieu Simeoni**[*]
[*] École Polytechnique Féderale de Lausanne, Switzerland

Abstract Penalized optimization methods allow to take advantage of powerful image priors and thus are perceived as strong contenders to the standard CLEAN-based pipeline for the next generation of radio interferometric imaging. The tremendous amount of data being produced by modern and future radio telescopes represents however a real challenge for such techniques, which are not inherently as scalable as CLEAN and its various embodiments. In this work, we present PolyCLEAN, a new optimization algorithm for radio interferometric imaging relying on sparse priors and a polyatomic variant of the Frank-Wolfe algorithm. The resulting algorithm is inherently scalable and admits a very intuitive iterative architecture, similar to that of CLEAN. We believe that this architectural proximity will strongly facilitate the adoption of the method by astronomers, as well as its integration in pre-existing CLEAN-based interferometric imaging pipelines. Thanks to its solid grounding in convex optimization theory, PolyCLEAN comes moreover with much more theoretical guarantees than CLEAN: it is provably convergent (with known convergence rate), comes with a natural and well-defined stopping criterion, and is amenable to a Bayesian interpretation allowing for uncertainty quantification. Finally, it was shown experimentally to be competitive with state-of-the-art optimal proximal algorithms. We demonstrate the superiority of PolyCLEAN over CLEAN on realistic synthetic datasets simulated with RASCIL. Reconstruction quality is drastically improved while maintaining similar computation times. Our algorithm is written in Python using the optimization package Pycsou, that comes with an efficient NUFFT-based implementation of the (de-)gridding operations. Related work available on ArXiv (112.02890).

Instantaneous Spectral Imaging with Learned Reconstruction

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

Abstract Spectral imaging, simultaneous spectroscopy and photography of a scene, is a fundamental diagnostic technique in physical sciences. Due to the limitations of two dimensional detectors to capture the three dimensional spatio-spectral data, conventional spectral imagers require a scanning process via taking multiple exposures, which renders them unsuitable for dynamic scenes. A computational technique for non-scanning spectral imaging has been proposed in [1] and its effectiveness has been demonstrated in the context of spectral imaging of stellar coronae [2], where a nonlinear parametric forward model is developed, and the resulting semi-blind deblurring problem is solved to estimate the underlying unknown parameters from multiplexed measurements. However, the proposed algorithm is computationally expensive, and has limited performance in noisy settings. Nevertheless, the characteristics of the image formation model, in particular the non-linearity and non-convexity of the optimization problem, motivate a machine learning approach to the solution of the inverse problem as recently proposed for a related computational imaging modality [3]. In this work, we develop a deep learning based approach to solve the ill-posed spectral image reconstruction problem based on a convolutional neural network using supervised learning with previous observational data. We analyze the effect of network architecture, training data, and loss function on the reconstruction performance, and evaluate the method on various observational settings. Related work accessible via the DOIs: [1] 10.1109/TIP.2014.2363903, [2] 10.3847/1538-4357/ab372a, and [3] 10.1109/TCI.2021.3075349.

Towards Unsupervised VLBI Imaging

Hendrik Müller*
*Max-Planck Institute for Radioastronomy, Germany

Abstract For decades VLBI imaging was done with CLEAN introducing a significant human bias. More recently RML and Bayesian approaches were developed and provide excellent, super-resolving reconstruction results. However, the reconstruction still relies on the finetuning of a larger number of non-trivial hyperparameters, which makes extended parameter surveys and an intensive simulation work necessary. I will review several current, complementary developments towards a completely unsupervised data analysis. Firstly, in multiscalar imaging the image is modelled by a multiscalar dictionary. The basis functions of the dictionary define filters in the Fourier domain that allow for an ideal separation between observed features and gaps in the uv-coverage, hence providing reconstruction robustness. In a recent work we demonstrated that this completely data-driven approach recovers images of compatible quality as state-of-the-art RML imaging software, but without the need of a tedious parameter selection procedure [1,2]. The approach has been extended successfully to polarimetry, and dynamic reconstructions where it is demonstrated to reliably recover dynamics at the event horizon on minute-timescales [3]. Secondly, as a natural extension we currently study AI-based reconstruction algorithms in which a series of image filters (encoder/decoder) is learned to compress the sidelobe pattern of partially removed beams [4]. Thirdly, as a complementary approach towards unsupervised imaging, we study multi-objective optimization with parallel genetic evolution [5]. I will present our current developments in these directions, future roadmaps and first applications to observations. Related publication and preprints: [1] Müller, H. & Lobanov, A., A&A, 2022, [2] Müller, H. & Lobanov, A., subm. to A&A, [3] Müller. H. & Lobanov, A., subm. to A&A, [4] Müller, H. & Lobanov, A., in prep., and [5] Mus, A. & Müller, H., in prep.

Learned Interferometric Imaging for the SPIDER Instrument

Matthijs Mars*, **Jason McEwen***, and **Marta Betcke***
*University College London, UK

Abstract The Segmented Planar Imaging Detector for Electro-Optical Reconnaissance (SPIDER) is an optical interferometric imaging device that aims to offer an alternative to the large space telescope designs of today with reduced size, weight and power consumption. This is achieved through interferometric imaging. State-of-the-art methods for reconstructing images from interferometric measurements adopt proximal optimization techniques, which are computationally expensive and require handcrafted priors. In this work we present two data-driven approaches for reconstructing images from measurements made by the SPIDER instrument. These approaches use deep learning to learn prior information from training data, increasing the reconstruction quality, and significantly reducing the computation time required to recover images by orders of magnitude. Reconstruction time is reduced to $\sim$10 milliseconds, opening up the possibility of real-time imaging with SPIDER for the first time. Furthermore, we show that these methods can also be applied in domains where training data is scarce, such as astronomical imaging, by leveraging transfer learning from domains where plenty of training data are available. Preprint available on ArXiv (2301.10260).

Session 5

Physics Informed Machine Learning in Astronomy

Organizers: **François Lanusse** & **Jean-Luc Starck**

Summary *Machine Learning had a significant success in Astronomy in recent years, and it becomes obvious that useful applications of ML require a tight connection to physical modeling. In this session, we explore several aspects of imbuing physics as part of a ML model, from building hybrid models that merge both deep learning and physical models, to using known physical symmetries and equivariances to design dedicated neural architectures.*

Combining Physics and Machine Learning in Astronomy

François Lanusse*
*CEA - Saclay, France

Abstract As we move towards the next generation of astronomical surveys, our field is facing new and outstanding challenges at all levels of scientific analysis, from pixel-level data reduction to cosmological inference. As powerful as deep learning has proven to be in recent years, in most cases a deep learning approach alone proves to be insufficient to meet these challenges, and is typically plagued by issues including robustness to covariate shifts, interpretability, and proper uncertainty quantification, impeding their exploitation in scientific analysis. In this review talk, I will survey the different approaches that have emerged so far in our field to combine machine learning models with our physical knowledge and insight as to overcome these limitations. These include the use of deep generative models as part of a physically motivated causal structure, the design of neural networks architectures preserving exact symmetries, or the use of automatic differentiation to enable fast inference over physical models.

Data-driven analysis of strong gravitational lenses: opportunities and challenges for machine learning.

Laurence Perreault-Levasseur*
*Université de Montréal, Canada

Abstract Despite the remarkable success of the standard model of cosmology, the lambda CDM model, at predicting the observed structure of the universe over many scales, very little is known about the fundamental nature of its principal constituents: dark matter and dark energy. In the coming years, new surveys and telescopes will provide an opportunity to probe these unknown components. Strong gravitational lensing is emerging as one of the most promising probes of the nature of dark matter, as it can, in principle, measure its clustering properties on sub-galactic scales. The unprecedented volumes of data that will be produced by upcoming surveys like LSST, however, will render traditional analysis methods entirely impractical. In the recent years, machine learning has been transforming many aspects of the computational methods we use in astrophysics and cosmology. I will share our recent work in developing machine learning tools for the analysis of strongly lensed systems. I will describe how, using physics-informed methodologies such as simulation-based inference or score-based diffusion models, not only can we achieve greater analysis speeds by 5 to 6 orders of magnitude, but we can also out perform state-of-the-art traditional methods in terms of accuracy. I will also discuss the issue of proper uncertainty quantification in these analyses.

Multiscale modeling of galaxy-scale strong lenses

Galan Aymeric*
*École Polytechnique Fédérale de Lausanne, Switzerland

Abstract The striking observational phenomenon called strong gravitational lensing is a unique tool to study the evolution of the Universe and its content. In particular, it allows to resolve the distant source galaxies that would otherwise be too faint for detailed analyses, and constrain the dark and luminous mass distribution of the lens galaxies. Current and future high-resolution imaging data imply solving challenging image processing problems in order to capture all the complexity of the lens and source galaxies. In this talk, I will discuss about recent developments in the field of galaxy-scale strong lens modeling, with a particular emphasis on multiscale pixelated techniques and deep learning-based approaches. I will talk about regularization techniques based on wavelet transforms and continuous neural fields, embedded in the efficient framework of differentiable programming. I will also present specific examples of lens modeling applications to imaging data from the Hubble Space Telescope.

Rethinking data-driven point spread function modeling with a differentiable optical model

Tobias Ignacio Liaudat[*]
[*]CEA - Saclay, France,

Abstract In astronomy, upcoming space telescopes with wide-field optical instruments have a spatially varying point spread function (PSF). Specific scientific goals require a high-fidelity estimation of the PSF at target positions where no direct measurement of the PSF is provided. Even though observations of the PSF are available at some positions of the field of view (FOV), they are undersampled, noisy, and integrated in wavelength in the instrument's passband. PSF modeling represents a challenging ill-posed problem, as it requires building a model from these observations that can infer a super-resolved PSF at any wavelength and position in the FOV. I will present a novel framework for data-driven point spread function modeling by changing the PSF modeling space from the pixels to the wavefront. Our approach is based on a differentiable optical forward model that allows propagating the wavefront from the modeling plane to the pixel representation at the focal plane. The proposed change allows the transfer of a great deal of complexity from the instrumental response into the forward model. The proposed PSF model relies on efficient automatic differentiation technology and modern stochastic first-order optimization techniques developed by the thriving machine-learning community. I will then present an application for the Euclid space mission where the PSF is required with a super-resolution of factor three for every wavelength and FOV position. Our method represents a breakthrough in performance with respect to previous data-driven PSF models. Finally, I will introduce an extension of the model where we tackle wavefront recovery, which belongs to the family of phase retrieval problems. The proposed extension uses the differentiable forward modeling approach to exploit the diversity from the PSF spatial variations to estimate the underlying wavefront only from degraded in-focus observations. Preprint available on ArXiv (2203.04908).

Learning the Galaxy-Halo Connection by Imposing Exact Physical Symmetries

Kate Storey-Fisher[*], **David Hogg**[*], **Shy Genel**[†], and **Soledad Villar**[‡]
[*]New York University, USA, [*]Flatiron Institute, USA, [‡]Johns Hopkins University, USA

Abstract Astronomical data and simulations obey exact physical symmetries, such as translation, rotation, and permutation; it is crucial that machine learning approaches applied to this data preserve these symmetry rules. A primary current challenge in cosmological inference is learning the relationship between dark matter halos and the properties of the galaxies they host. Machine learning has shown success at this task, but most approaches either do not preserve the symmetries of cosmological simulations or do not utilize the full halo phase-space information. I present a new approach for enforcing physical symmetries for machine learning models: the core idea is that geometric features can be contracted into "invariant scalars" that remain the same under transformations. These can then be used as input features of the model and guarantee invariance (or equivariance) of the outputs. I apply this approach to produce an invariant description of dark matter halos, and show that a neural network trained on these features makes highly precise predictions of central galaxy properties. The approach achieves increased precision compared to non-invariant features containing the same information, and to standard halo properties. Our results are interpretable and give insight into the halo properties that govern galaxy formation, relevant to the open question of galaxy assembly bias. As the approach operates at the feature level, it can be integrated into a wide range of ML frameworks and spatial problems in astrophysics and other domains.

Sampling high-dimensional inverse problem posteriors with neural score estimation

Benjamin Remy[*], **François Lanusse**[*], **Jean-Luc Starck**[*], **Niall Jeffrey**[†], **Jia Liu**[‡], and **Tim Schrabback**[⋆]
[*]CEA - Saclay, France, [†]École Normale Supérieure Ulm, France, [‡]Kavli IPMU, University of Tokyo, Japan, [⋆]University of Bonn, Germany

Abstract I will present a novel methodology to address many ill-posed inverse problems, by providing a description of the posterior distribution, which enables us to get point estimate solutions and to quantify their associated uncertainties. Our approach combines Neural Score Matching for learning a prior distribution from physical simulations, and a novel posterior sampling method based on an annealed HMC algorithm to sample the full high-dimensional posterior of our problem. In the astrophysical problem we address, by measuring the lensing effect on a large number of galaxies, it is possible to reconstruct maps of the Dark Matter distribution on the sky. But because of missing data and noise dominated measurement, the recovery of dark matter maps constitutes a challenging ill-posed inverse problem. We propose to reformulate the problem in a Bayesian framework, where the target becomes the posterior distribution of mass given the galaxies shape observations. The likelihood factor, describing how light-rays are bent by gravity, how measurements are affected by noise, and accounting for missing observational data, is fully described by a physical model. Besides, the prior factor is learned over cosmological simulations using a recent class of Deep Generative Models based on Neural Score Matching and takes into account theoretical knowledge. We are thus able to obtain samples from the full Bayesian posterior and can provide Dark Matter map reconstruction alongside uncertainty quantification. We present an application of this methodology to the reconstruction of the COSMOS field and yield the highest quality convergence map of this field to date. We find the proposed approach to be superior to previous algorithms, scalable, providing uncertainties, using a fully non-Gaussian prior and promising for future weak lensing surveys. We also apply this framework to MRI reconstruction, illustrating the diversity of inverse problems we can solve. Preprint available on ArXiv (2201.05561).

Deep Learning-based galaxy image deconvolution

Utsav Akhaury[*], **Jean-Luc Starck**[*], **Pascale Jablonka**[*], **Frédéric Courbin**[*], and **Kevin Michalewicz**[*]
[*]École Polytechnique Féderale de Lausanne, Switzerland

Abstract With the onset of large-scale astronomical surveys capturing millions of images in multiple frequency bands, there is an increasing need to develop fast and accurate deconvolution algorithms that generalize well to different images. A powerful and accessible deconvolution method would allow for the reconstruction of a cleaner estimation of the sky. The deconvolved images would be helpful to perform photometric measurements to help make progress in the fields of galaxy formation and evolution. While for the vast majority of astrophysical cases, the purpose of the deconvolution is to reconstruct galaxy images at high spatial resolution in each photometric band independently, there are cases, essentially at low signal-to-noise ratios, for which a joint multi-channel deconvolution can improve the detection and characterization of the systems. As such, we proceeded to tackle the astrophysical deconvolution problem in two steps: 1. We investigated the single channel deconvolution problem and proposed a new deconvolution method based on the Learnlet decomposition. We have compared the performance of different Unet architectures and Learnlet for image deconvolution in the astrophysical domain by following a two-step approach: a Tikhonov deconvolution with a closed-form solution, followed by post- processing with a neural network. Our numerical results based on the simulated dataset show a detailed comparison between the considered methods for different noise levels. 2. We then extended the idea to multi-channel deconvolution across three different frequency bands. Visually, our naive joint deconvolution algorithm already gives us very encouraging results. The ultimate goal is to come up with a robust multi-channel deconvolution method that could efficiently be applied to real data in the future. Underlying work accessible via the DOI 10.3389/fspas.2022.1001043.

Scalable and equivariant spherical CNNs by discrete-continuous (DISCO) convolutions

Jason McEwen*
*University College London, UK

Abstract No existing spherical convolutional neural network (CNN) framework is both computationally scalable and rotationally equivariant. Continuous approaches capture rotational equivariance but are often prohibitively computationally demanding. Discrete approaches offer more favorable computational performance but at the cost of equivariance. We develop a hybrid discrete-continuous (DISCO) group convolution that is simultaneously equivariant and computationally scalable to high-resolution. While our framework can be applied to any compact group, we specialize to the sphere. Our DISCO spherical convolutions not only exhibit SO(3) rotational equivariance but also a form of asymptotic SO(3)/SO(2) rotational equivariance, which is more desirable for many applications (where SO(n) is the special orthogonal group representing rotations in n-dimensions). Through a sparse tensor implementation we achieve linear scaling in number of pixels on the sphere for both computational cost and memory usage. For 4k spherical images we realize a saving of 10^9 in computational cost and 10^4 in memory usage when compared to the most efficient alternative equivariant spherical convolution. We apply the DISCO spherical CNN framework to a number of benchmark dense-prediction problems on the sphere, such as semantic segmentation and depth estimation, on all of which we achieve the state-of-the-art performance. For further details see the preprint on ArXiv (2209.13603).

Delensing Gravitational Lensing Effects with Physics-Informed Neural Networks

Ayoub Tajja*, and **Emma Tolley***
*École Polytechnique Fédérale de Lausanne, Switzerland

Abstract Gravitational lensing is caused by massive amounts of matter which create gravitational fields that distort light from distant galaxies. In a strongly lensed system, the distortion is strong enough to create multiple images or Einstein rings of the distant light source. For an observed lensed source, reconstructing both the original source and the lensing potential of the strong lens system enables us to have a better understanding of the mass distribution and composition of the Universe. The new-generation optical surveys (Webb, NGRT) produce massive high-resolution images of the sky and gravitational lenses. These data require faster and more efficient lens modeling techniques. Image decomposition methods in mathematical bases and Convolutional Neural Networks have been increasingly used in the field of astronomy to accelerate or automate data analysis. However, these models may not be able to reconstruct properly the original image as well as the lensing potential mapping. In order to help the network with these accuracy issues, physics-based constraints can be used. This kind of network is a Physics-Informed Neural Network (PINN). We will discuss the two pipelines developed to reconstruct both original image and lensing potential, which use a physics-based constrained network and a preprocessing image decomposition analysis.

Session 6

Frontier of Interferometric Imaging from SKA to ngEHT

Organizers: Kazunori Akiyama & Yves Wiaux

Summary *Computational imaging is a key process in radio and optical/near-infrared interferometry to reveal the fine views of the universe from observational data taken in Fourier space. Over the last decade, significant progress has been made in the development of new computational imaging techniques to address and overcome various challenges brought by the advent of the new instruments including the Event Horizon Telescope (EHT), Low Frequency Array (LOFAR), MeerKAT and Very Large Telescope Interferometer (VLTI), as well as upgrades of existing facilities such as Very Large Array (VLA). Many algorithmic and data processing challenges arise in our quest to endow these instruments, set to observe the sky at new regimes of sensitivity and resolution, with their expected acute vision. In this new era, imaging encompasses not only forming 2D spatial maps of observed fields of view, but also the reconstruction of the spectrum, polarization, and dynamics of the sources of interest, not to mention the mapping of underlying physical quantities. This focused session will gather interferometric imaging experts to review the exciting frontiers in the field.*

Broadening horizons of black hole imaging with the EHT and ngEHT

Kazu Akiyama*
*MIT Haystack Observatory, USA

Abstract Earlier this year, the Event Horizon Telescope (EHT) Collaboration released the first images of the Milky Way supermassive black hole Sgr A⋆. These results are followed by the first-ever black hole imaging made for the supermassive black hole M87⋆ in the M87 galaxy. The images of Sgr A⋆ and M87⋆ show bright rings of light closing the central dark shadows of the 4 million and 6.5 billion solar mass black holes, respectively. These black hole images were obtained from interferometric data obtained in April 2017 at 230 GHz (or 1.3 mm wavelength) using the technique of very long baseline interferometry with an array of eight radio telescopes at six geographic sites spread across the globe. A key integral part allowing these breakthroughs is the innovative algorithms of interferometric imaging specialized for the EHT. In radio interferometry, the imaging problem is formulated as an under-determined problem that reconstructs an image from measurements in its Fourier space often with imperfect calibrations. EHT has two unique challenges: its sparse Fourier coverage and severe limitations in data calibration. The limitations of the traditional techniques to deal with these two issues have driven the development of new dedicated forward imaging techniques including the regularized maximum likelihood methods and the Bayesian imaging methods. In these methods, images are directly solved from data through the measurement equation, offering flexibility to the imaging such as the usage of closure quantities immune to calibration uncertainties. These new dedicated methods are shown to outperform the traditional imaging techniques, not only for the EHT but for other interferometric arrays such as VLA, ALMA, and ngVLA. This review talk will give an overview of the exciting frontier of black hole imaging with the EHT, and new technical challenges with the next generation EHT (ngEHT) for the next decade.

Next Generation Black-Hole Imaging: Beyond a 2D Image

Aviad Levis*
*California Institute of Technology, USA

Abstract The Event Horizon Telescope (EHT) is a unique computational camera with a goal of imaging the glowing gas surrounding supermassive black holes. In 2022, the EHT collaboration revealed the first imaging results of the black hole at the center of our galaxy: Sagittarius A⋆ (Sgr A⋆). The image of Sgr A⋆ was computationally reconstructed from measurements of synchronized telescopes around the globe. Imaging Sgr A⋆ was challenging due to its evolution, which is dynamic on the timescale of acquisition. While an image certainly offers interesting insights, looking toward the future, we are developing new computational algorithms that aim to go beyond a 2D image. For example, could we use EHT observations to recover the dynamic evolution or even 3D structure? In this talk, I will highlight new and exciting prospects for the future of black-hole imaging. By peeling away different layers of the underlying physics, we show how algorithms could shed light on complex dynamic processes. Our hope is that in the not-too-distant future, they will enable scientific discovery and even provide a glimpse into the very nature of space-time itself in our galaxy's most extreme environment.

The state of the art in classical radio astronomical deconvolution algorithms

André Offringa*
*ASTRON, The Netherlands

Abstract Högbom CLEAN, a classical algorithm for the deconvolution of radio astronomical images, is a greedy method to perform deconvolution. It is similar to the matching pursuit sparse approximation algorithm. Various variants of CLEAN exist, such as multi-scale (MS), multi-frequency (MF), multi-polarization and auto-masking algorithms. Almost all scientific products from interferometric imaging observations are, as of the moment of writing, made using variants of the classical CLEAN algorithm. WSClean is an interferometric imager, and I will discuss the latest deconvolution algorithms that are implemented in this imager. It implements a MSMF algorithm described in Offringa & Smirnov (2014), with various optimizations to allow accurate deconvolution of large images (100K x 100K) with acceptable computational requirements. The basis functions used are delta-components and circular Gaussians. These typically need more components compared to other dictionaries (shapelets, wavelets), but a practical advantage is that these can be analytically inverted, which is an advantage when the produced models are to be used in calibration. Recently, the MSMF deconvolution algorithm was extended with a forced-spectrum option (Ceccotti et al., in prep; see also [1]). We show this method produces improved results compared to per-component spectral fitting, and is particularly useful for producing calibrator models. Another extension is the support for a direction-dependent PSF during deconvolution, without added computational costs (Van der Tol et al., in prep; see also [2]). These algorithms are packaged in a library, Radler, which provides a C++ and Python interface. Various new gridders have also become available. Related research available at [1] https://wsclean.readthedocs.io/en/latest/wideband_deconvolution.html#forced-spectral-indices, and [2] https://wsclean.readthedocs.io/en/latest/direction_dependent_psfs.html.

Scalable wide-field imaging in Radio interferometry: from sparsity-based to learned regularization

Arwa Dabbech*, **Matthieu Terris***, **Amanda Wilber***, **Adrian Jackson**†, and **Yves Wiaux***
*Heriot-Watt University, UK, †University of Edinburgh, UK

Abstract Aperture synthesis by radio interferometry probes the sky through incomplete noisy measurements in the Fourier domain. Modern radio telescope arrays are expected to form high-resolution high-dynamic range images over wide fields of view and from sheer data volumes, in which context image formation has never been more challenging. Over the last decade, the SARA family of algorithms has been developed to address this challenge, leveraging state-of-the-art algorithmic structures borrowed from optimization theory. We discuss here its latest incarnations, built on the forward-backward algorithm, which simply alternate a handcrafted (uSARA) or a learned (AIRI) regularization denoiser with a gradient-descent data-fidelity step in a provably convergent scheme. This new framework is equipped with automated functionalities for the faceting of its denoisers and the block-decomposition of the wide-field measurement operator, providing seamless scalability to extreme image and data dimensions. The imaging precision and computational performance of uSARA and AIRI, as well as the suitability of the images formed for subsequent scientific analysis is thoroughly validated from MeerKAT and ASKAP observations of complex structure with diffuse and faint emission, and in comparison with the widely used WSClean imager. Preprints accessible on ArXiv (2207.11336), and on the Heriot-Watt University Research Portal (52551784, 54234902).

Designing calibration and imaging pipelines for portability and reproducibility

Oleg Smirnov*
*SKA SA, Rhodes University, South Africa

Abstract Radio interferometric imaging pipelines are complex beasts with many moving components (over and above the imaging algorithms per se); radio data reduction is often described as "death by a million papercuts". This has made it very difficult to make pipelines that are (a) robust; (b) portable across computing architectures; (c) yield reproducible results and therefore (d) can provide a basis for fair comparisons of novel and traditional calibration and imaging approaches. I will describe the dask-ms and codex africanus ecosystem, as well as the Stimela2 workflow management framework, which goes a long way towards addressing this problem. The code ecosystem has enabled a new generation of calibration and imaging packages that are system-agnostic, in the sense they can be deployed on single nodes, on HPC systems, and in cloud environments. In the meantime, Stimela2 provides a framework for tying these packages (as well as legacy software) into concise and modular pipeline "recipes". I will demonstrate a full self-calibration recipe that can be applied to publicly available (and cloud-hosted) MeerKAT datasets, and can form the basis for a new style of open and reproducible "radio imaging contest".

Resolve: Multi-domain Imaging of Radio-interferometric Data through Bayesian Inference

Jakob Knollmüller*
*Technical University of Munich, Germany

Abstract Radio-interferometric observations are inherently incomplete, as they only sparsely probe the Fourier components of a source on the sky. In order to recover the original source structure from noisy and incomplete data it is necessary to solve an inverse problem. Unfortunately there is an infinite number of images that are compatible with the data. The Resolve algorithm approaches the problem from a probabilistic perspective by imposing a prior model over plausible source configurations and combines this information with the likelihood of the data in order to obtain a probability distribution of possible source configurations that are a priori plausible and compatible with the observation. As the result is given in form of a probability distribution, an uncertainty quantification of image features is inherently provided. The method is based on Gaussian processes with adaptive kernels along different domains and it can deal with multi-frequency, dynamic and polarization imaging tasks, and combinations thereof. Efficient methods for approximate Bayesian inference allow its application in high-resolution settings with millions of parameters. It can make use of regular visibility data, as well as closure quantities and can also accommodate for self-calibration. We will showcase several examples, including dynamic reconstructions of black holes from simulated, as well as EHT data. One example application can be found at nature.com/articles/s41550-021-01548-0.

PnP algorithms for large-scale high-dynamic range imaging in radio astronomy

Matthieu Terris*, **Chao Tang***, **Chung San Chu***, **Arwa Dabbech***, **Adrian Jackson**†, and **Yves Wiaux***
*Heriot-Watt University, UK, †University of Edinburgh, UK

Abstract Plug-and-Play (PnP) algorithms are attractive alternatives to proximal algorithms when solving inverse imaging problems. By learning a Deep Neural Network (DNN) behaving as a proximal regularization operator, one waives the computational complexity induced by sophisticated image priors, and the sub-optimality of handcrafted priors. In this work, we propose AIRI, a PnP framework tailored to aperture synthesis by Radio Interferometry (RI) in astronomy, where reconstruction algorithms are to form images with extreme dynamic range from sheer volumes of Fourier data. Our approach relies on three steps. Firstly, given the absence of appropriate high-dynamic RI database, we borrow and curate low-dynamic range training databases from alternative imaging modalities. Secondly, we train a shelf of DNN denoisers, each tailored to a specific dynamic range. Our training procedures rely on two important elements. On the one hand, the noise level is adjusted to the target dynamic range, and the database dynamic range is enhanced on the fly via image exponentiation. On the other hand, the training loss is enhanced with a Jacobian regularization term endowing the denoisers with a nonexpansiveness property ensuring PnP convergence. Thirdly, the denoisers are plugged into proximal splitting structures, yielding dedicated PnP algorithms for high dynamic range imaging. We performed an extensive validation of AIRI, using databases borrowed from either optical astronomy or magnetic resonance imaging, and plugging denoisers into either Forward-Backward or Primal-Dual optimization structures. All AIRI incarnations were demonstrated on RI simulations to be competitive with the state-of-the-art optimization technique SARA, whilst providing significant acceleration. AIRI's precision and scalability was also demonstrated on real large-scale data from the MeerKAT telescope. Part of this work accessible on ArXiv (2202.12959, 2207.11336).

The importance of compressed sensing and regularization: An application to Faraday depth imaging

Miguel A.Carcamo *, **Anna Scaife***, **Emma Alexander***, and **Patrick Leahy***
*University of Manchester, UK

Abstract Studies of astrophysical Faraday rotation have led to a number of improvements in our understanding of astrophysical systems, including our own Galaxy, external galaxies and on larger scales the intra- and inter-cluster medium. Moreover, the advent of broadband polarization measurements with modern radio telescopes have enabled the widespread use of the RM Synthesis technique. This technique is based on the fact that there is Fourier relationship between polarized intensity (corrected by the spectral dependency) as a function of wavelength-squared and the Faraday dispersion function. However, the reconstruction of Faraday depth structures from incomplete spectral polarization radio measurements is an under-constrained problem that requires additional regularization. In this talk I will present a novel object-oriented compressed sensing framework for Faraday depth reconstruction: CS-ROMER. I will explain how we are able to identify the optimal wavelet basis for a given observational configuration in order to reconstruct Faraday depth signals from incomplete measurements, and I will demonstrate the application of this technique in three different scenarios using real data from the JVLA, eMERLIN and MeerKAT radio-telescopes. Related work accessible via the DOI 10.1093/mnras/stac3031.

PRIISM: Imaging tool for ALMA telescope

Shiro Ikeda[*], **Takeshi Nakazato**[†], **Takashi Tsukagoshi**[‡], **Akio Taniguchi**[⋆], **Masayuki Yamaguchi**[⊤], **George Kosugi**[▷], **Mareki Honma**[▷], **Ryohei Kawabe**[▷], **Kazu Akiyama**[◇], **Toshinobu Takagi**[○], **Yoshihiko Yamada**[○], and **Atsushi Miyazaki**[○]
[*]Institute of Statistical Mathematics, Japan, [†]National Astronomical Observatory of Japan, Japan, [‡]Ashikaga University, Japan, [⋆]Nagoya University, Japan, [⊤]Institute of Astronomy and Astrophysics, Japan, [▷]National Astronomical Observatory of Japan, Japan, [◇]MIT Haystack Observatory, USA, [○]Japan Space Forum, Japan

Abstract Radio interferometry is a method to obtain images by combining radio telescopes separated by distance. Recently revealed black hole shadow images, published by the Event Horizon Telescope collaboration (EHTc), were taken by a Very Long Baseline Interferometer (VLBI), which highlighted the advantage of radio interferometry, The "imaging" process, which finalizes the image from the data, is essential for the radio interferometer. In addition to the traditional CLEAN, the EHTc developed and used new imaging methods which include Regularized Maximum Likelihood (RML) methods. The RML method estimates an image by balancing the fidelity term and the image constraint terms, such as sparsity and smoothness. Since the EHT collaboration has demonstrated the validity of RML methods, we are now developing the Python module for Radio Interferometry Imaging with Sparse Modelling (PRIISM), which is the RML method designed for the imaging of the Atacama Large Millimeter-wave and Submillimeter-wave Interferometer (ALMA) telescope data. Our goal is to provide an imaging tool which estimates good-quality images automatically from ALMA data. Because the ALMA has around 50 telescopes, which is much larger than the EHT, observed data are dense and estimation of high-resolution images is required. This makes the imaging process computationally expensive. It is important to tune newly introduced RML parameters efficiently. We propose a new criterion to evaluate the imaging parameters and use Bayesian optimization for parameter selection. Due to atmospheric effects, the gains of ground-based telescopes fluctuate. In radio interferometry, the fluctuations have a significant impact on images, so it is necessary to reduce them. This operation is called self-calibration. We have extended the method of Repetti et al. 2017 and implemented it for PRIISM. With this new self-calibration, the final images are automatically obtained. We will show some imaging results.

Scalable Analysis and Synthesis in Radio Interferometric Imaging via a Chunked NUFFT

Matthieu Simeoni[*], **Joan Rue Queralt**[*], and **Sepand Kashani**[*]
[*]École Polytechnique Fédérale de Lausanne, Switzerland

Abstract Analysis and synthesis are key steps of the radio-interferometric imaging process, serving as a bridge between the visibility and sky domain. They can be expressed as partial Fourier transforms involving a large number of non-uniform frequencies and spherically-constrained spatial coordinates. Due to the data non-uniformity, these partial Fourier transforms are computationally expensive and represent a serious computational bottleneck in the image reconstruction process. The W-gridding algorithm achieves log-linear complexity for both steps by applying a series of 2D non-uniform FFTs (NUFFT) to the data sliced along the so-called W frequency coordinate. A major drawback of this method however is its restriction to direction-cosine meshes, which are fundamentally ill-suited for large field of views. (Heavy pixel distortion as the radial distance increases.) In this presentation we introduce a novel algorithm for analysis/synthesis based on a 3D-NUFFT. Unlike W-gridding, the latter is compatible with arbitrary spherical meshes such as the popular HEALPix scheme for spherical data processing. Relying on the 3D-NUFFT allows one to optimally select the size of the inner FFTs, in particular the number of W-planes. This results in a better performing and user-friendly algorithm. We moreover show how to leverage chunking to cope with the challenging scale of next-generation of radio telescopes such as the Square Kilometer Array (SKA). Our chunking strategy partitions both the Fourier and image domains using a quad-tree algorithm to balance the load amongst workers. It then creates and distributes a 3D-NUFFT computational graph over multiple threads or processes with reduced memory requirements. We conclude with RASCIL-based benchmark results comparing our solution to W-gridding in terms of accuracy, computation time, and memory usage for various realistic setups.

Part III

Medical Imaging

Session 7

Medical Imaging in Low-Resource Settings

Organizers: **Nicholas Durr** & **Audrey Bowden**

Summary *Despite significant progress in modern healthcare, state-of-the-art healthcare in first-world environments differs significantly from that in low-resource settings. Constraints on cost, size, usability and environmental stability pose interesting challenges to engineer suitable solutions in these spaces. This session brings together impactful research that showcases innovative strategies for imaging that are suitable for implementing in low-resources settings such as rural areas, primary care centers and low- and middle-income countries.*

Medical imaging in low-resource settings with an example in eye health

Nicholas Durr*
*Johns Hopkins University, USA

Abstract Large financial markets have spurned technology innovation to develop solutions for many important healthcare problems in high-resource settings. However, there is perhaps a greater opportunity for societal impact in creating accessible healthcare technologies for the low-resource world. The first part of this presentation will provide context for the session and introduce the technologies and applications being explored. Nanofabrication technologies will be described that enable the detection of parasites and bacteria in challenging samples at the point of care. Computer vision technologies in mobile phones can task-shift vector surveillance and ultimately reduce the prevalence of mosquito-borne diseases. Computational spectroscopy technologies are being developed to enable compact and low-cost optical coherence tomography systems for improved ocular diagnostics. Lastly, hardware and software innovations are paving the way for portable magnetic resonance imaging systems that are practical for use in low resource settings. The second part of this presentation will cover the idea-to-market cycle of a computational optics product for to eye care. Over a billion people have poor vision that could be corrected with a simple pair of prescription eyeglasses. These uncorrected refractive errors have a major impact on productivity, quality of life, and access to education. I will present our work to address this problem via problem framing, technology development, device design, and commercial product deployment. We created a new wavefront sensing technology that allows eyeglass prescriptions to be accurately and easily measured without any moving components. We tested several autorefractor prototypes in large international clinical studies in rural and urban, high- and low-resource settings. The resulting autorefractor has been used to prescribe over 4 million pairs of eyeglasses over the last five years.

Portable Technologies for Medical Diagnostics and Imaging in Low Resource Settings

Hatice Ceylan Koydemir*
*Texas A&M University, USA

Abstract Developments in micro- and nano-fabrication technologies have substantially advanced the powerful and smart portable biomedical devices for point-of-care diagnostics, especially for low-resource settings. Such portable devices offer great potential to transform the future of healthcare through personalized telemedicine. They can improve people's life quality through rapid and accurate detection and classification of various analytes using disposable tests with integrated computational learning approaches. This talk will present 14 portable optical platforms engineered for detecting and/or imaging tissue, parasite, cell, bacteria, and virus plaques in complex environments [1-13]. Miniaturization of a variety of imaging techniques, including fluorescence imaging, speckle imaging, brightfield imaging, and lens-free holographic imaging, will be introduced together with the design details of each platform and their assay components. The advantages and limitations of each miniaturized and portable platform will be discussed compared to the standard imaging methods. Procedures followed in the evaluation of their design robustness will be provided by highlighting the major considerations in future designs. With the interdisciplinary work, these inspiring and engineered medical diagnostic and imaging technologies could be helpful in building the next generations' more advanced, accessible, and affordable imaging and sensing devices. Related work accessible via the DOIs: [1] 10.48550/arXiv.2207.00089, [2] 10.1038/s41377-020-00358-9, [3] 10.1038/s41746-020-0282-y, [4] 10.1039/C9AN00532C, [5] 10.1039/c4lc01358a, [6] 10.1016/j.tmaid.2018.12.001, [7] 10.1038/s41598-019-56474-z, [8] 10.1039/c9lc00652d, [9] 10.1038/s41377-019-0203-5, [10] 10.1039/c8ra06473c, [11] 10.1021/acsphotonics.8b00146, [12] 10.4269/ajtmh.16-0912, [13] 10.1515/nanoph-2017-0001, and [14] 10.1038/s41377-018-0110-1.

Sustainable magnetic resonance imaging for low and middle income countries

Andrew Webb*
*Leiden University Medical Center, Netherland

Abstract Commercial magnetic resonance imaging (MRI) systems cost millions of pounds to purchase, require large electromagnetically shielded spaces to house, are extremely expensive to maintain and require highly trained technicians to operate. These factors together means that their distribution is confined to centrally-located medical centers in large towns and cities. Globally over 70% of the world's population has absolutely no access to MRI, and clinical conditions which could benefit from even very simple scans cannot be treated [1]. The magnetic fields typically used are very high, which means that there are severe contraindications so that, for example, MRI cannot currently be used in the emergency room. From the considerations above it is clear that if low-field MRI could be made more portable, accessible and sustainable then it would open up new opportunities in both developed and developing countries [2-4]. In order to achieve portability, we design systems that use thousands of very small low-cost permanent magnets, arranged in designs which have no fringe field and therefore very easy siting requirements. The low magnetic fields allow scanning of patients with implants, and the scanner could potentially be transported on an ambulance for differentiation of hemorrhagic or ischemic stroke, for example. This talk will cover aspects of magnet, gradient and RF coil design for low fields (∼50 mT), as well as corrections for gradient- and B0-distortions, and present the latest in vivo results as well as an outlook on future developments. Related work accessible via the DOIs [1] 10.11604/pamj.2018.30.240.14000, [2] 10.1002/jmri.26638, [3] 10.1002/nbm.4846, and at [4] who.int/publications/. . . /9789241504546.

Technologies for task-shifting: novel optical instrumentation and computer vision algorithms to bridge health access gaps in low resource settings

Soumyadipta Acharya*
*Johns Hopkins University, USA

Abstract At the Johns Hopkins Center for Bioengineering Innovation and Design, we focus on developing novel approaches towards transforming health care access in Low- and Middle-Income Countries, using computer vision, machine intelligence, and optical instrumentation. A common thread that ties together our work across multiple clinical and public health domains is the concept of 'task-shifting', wherein lesser skilled providers (such as semi-literate community health workers and midwives), are empowered to provide diagnostic or therapeutic services that are traditionally possible through providers with a much higher skill level (such as physicians). We will present several examples of technologies for task shifting, field validated in East Africa, and India, including: (1) VectorCam, a computer vision enabled virtual entomologist, for smart malaria vector surveillance. Vector surveillance is the process of trapping and identifying different disease carrying mosquito species in each locality, in order to drive species-specific malaria control interventions. Expertise and resource bottlenecks in surveillance can be overcome by an AI enabled handheld tool that could automate and task-shift the process of vector identification to Community Health Workers. We have co-created multiple designs of this imaging tool with health workers in Uganda, Ghana, and Zambia, and validated its performance on major mosquito species of interest in those geographies, with a $> 93\%$ overall accuracy. (2) HemoGlobe, a non-invasive (bloodless), cloud connected method for maternal anemia screening in rural communities. We have developed and clinically validated a novel non-invasive optical sensor and machine learning algorithm that can task shift anemia screening to community health workers to detect women with severe and moderate gestational anemia. In a study on 1583 women in rural India, we obtained a Pearson's CC of 0.81 in estimating true Hemoglobin levels. Other examples will also be presented.

Computational spectroscopy to enable compact medical imaging for low-resource settings

Audrey K Bowden[*], **Laura Waller**[†], **Joseph Malone**[*], **Neerja Aggarwal**[†], and **Eric N Markley** [†]
[*]Vanderbilt University, USA, [†]University of California - Berkeley, USA

Abstract Traditional spectrometers utilize an optical element such as a grating or prism to linearly disperse a broadband optical signal into its constituent wavelength components. The resulting one-to-one spectral-to-spatial (camera pixel) mapping simplifies measurement of the amplitudes of the underlying spectral components. A major limitation of these spectrometers, however, is the inherent tradeoff between bandwidth, dispersion angle (or system footprint) and spectral resolution, which leads to bulky designs with large footprints and expensive components. Computational spectroscopy breaks the inherent one-to-one spatial-to-spectral pixel mapping of traditional spectrometers by multiplexing spectral data over a given sensor region. Most computational spectrometers require components that are complex to design, fabricate or both. DiffuserSpec is a simple computational spectrometer that utilizes the inherent spectral dispersion of commercially available diffusers to generate speckle patterns that are unique to each wavelength. Using Scotch tape as a diffuser, we demonstrate narrowband and broadband spectral reconstructions with 2-nm spectral resolution over an 85-nm bandwidth in the near-infrared, limited only by the bandwidth of the calibration dataset. We also investigate the effect of spatial sub-sampling of the 2D speckle pattern. Finally, we show the ability to integrate DiffuserSpec into complex imaging systems like optical coherence tomography to enable new designs for medical imaging systems. Our ultimate goal is to develop a compact OCT system for low-resource settings.

Multiscale Integration of Widefield and High-Resolution Imaging Data for Cervical Precancer and Cancer Detection

David Brenes[*], **Jackson Coole**[*], **Yajur Maker**[*], **Jenny Carns**[*], **Richard Schwarz**[*], **Mila P. Salcedo**[†], **Karen C. Borba**[‡], **Julio C. Possati-Resende**[‡], **Kathleen M. Schmeler**[†], and **Rebecca Richards-Kortum**[*]
[*]Rice University, USA, [†]University of Texas MD Anderson Cancer Center, USA, [‡]Barretos Cancer Hospital, Brazil

Abstract In 2020, the World Health Assembly adopted a Global Strategy for cervical cancer elimination which includes the goal to provide treatment to 90% of women with precancer. To help enable this goal, we are developing a multimodal system to perform low-cost colposcopy and high-resolution microendoscopy (HRME) imaging at the point-of-care. The system will use deep learning models to interpret and integrate the multimodal, multiscale imaging data with the goal of diagnosing patients and guiding biopsy acquisition by a non-specialist. In this work we present, the results of integrating two modality-specific deep learning models for colposcopy and HRME images for cervical precancer detection. We collected imaging data from over 1600 patients in a population with a high prevalence of cervical precancer (Hunt et al. Int J Cancer 2021). Lesions observed in colposcopy images were annotated by two expert clinicians. These data were used to train a U-Net model for lesion segmentation. We also developed a multi-task (segmentation and classification) network for pathology prediction using HRME images (Brenes et al. Comp Med Imaging & Graphics 2022). The two architectures were integrated with both mid- and late-fusion strategies. Our preliminary results in lesion segmentation show a mean intersection over union (mIOU) above 0.40 for high-grade lesions when compared to the union of the clinician annotations. For reference, the mIOU between the two experts was 0.36 for high-grade lesions. With respect to our HRME pathology prediction model, we demonstrate a performance comparable to expert colposcopy impression with a sensitivity of 0.94 (p=0.3) and specificity of 0.54 (p=1.0). Preliminary prospective results show consistent performance with previous work (Coole et al. Biom Optics Express 2022).

High Dynamic Range Tomography via Modulo Radon Transform

Ayush Bhandari[*]
[*]Imperial College London, UK

Abstract Recently, experiments have been reported where researchers were able to perform high dynamic range (HDR) tomography in a heuristic fashion, by fusing multiple tomographic projections. This approach to HDR tomography has been inspired by HDR photography and inherits the same disadvantages. Taking a computational imaging approach to the HDR tomography problem, we here suggest a new model based on the modulo Radon transform (MRT), which we rigorously introduce and analyze. By harnessing a joint design between hardware and algorithms, we present a single-shot HDR tomography approach, which to our knowledge, is the only approach that is backed by mathematical guarantees. On the hardware front, instead of recording the Radon transform projections that may potentially saturate, we propose to measure modulo values of the same. This ensures that the HDR measurements are folded into a lower dynamic range. On the algorithmic front, our recovery algorithms reconstruct the HDR images from folded measurements. Beyond mathematical aspects such as injectivity and inversion of the MRT for different scenarios including band-limited and approximately compactly supported images, we also provide a first proof-of-concept demonstration. To do so, we implement MRT by experimentally folding tomographic measurements available as an open source dataset using our custom designed modulo hardware. Our reconstruction clearly shows the advantages of our approach for experimental data. In this way, our MRT based solution paves a path for HDR acquisition in a number of related imaging problems. Related work accessible via the DOIs: 10.48550/arXiv.2105.04194, and 10.23919/Eusipco47968.2020.9287586.

Single-Modality Supervised Joint PET-MR Reconstruction

Guillaume G. Corda-D'Incan[*], **Julia A. Schnabel**[†‡], and **Andrew Reader**[*]
[*]King's College London, UK, [†]Technical University of Munich, Germany, [‡]Helmholtz Center Munich, Germany

Abstract We present a new approach for deep learned joint PET-MR reconstruction. The proposed network is inspired by conventional synergistic methods using a joint regulariser. The maximum a posteriori expectation-maximization algorithm for PET and the Landweber algorithm for MR are unrolled and interconnected through a deep learned joint regularization step. The parameters of the joint U-Net regulariser and the respective regularization strengths are learned and shared across all the iterations. Along with introducing this framework, we propose an investigation of the impact of the loss function selection on network performance. We explored how the network performs when trained with a single or a multi-modality loss. Finally, we explored under which settings a joint reconstruction was beneficial for MR reconstruction by using various undersampling factors. The results obtained on 2D simulated data show that the joint networks outperform conventional non-deep learned reconstruction methods, both synergistic and independent. For PET, the network trained with only a PET loss achieves a better global reconstruction accuracy than the version trained with a weighted sum of PET and MR loss terms. More importantly, the former further improves the reconstruction of PET-specific features where deep learned MR-guided methods show their limit. Therefore, using a single-modality loss to supervise the training while still reconstructing the two modalities in parallel leads to better reconstructions and improved modality-unique structure recovery. For MR, while the same effect is observed, the gain from joint reconstruction only occurs in the presence of highly undersampled data. Single modality loss joint reconstruction results are also demonstrated on 3D clinical PET-MR datasets. Related work accessible via the DOIs: 10.1109/NSS/MIC42677.2020.9508086, and 10.1109/NSS/MIC44867.2021.9875927.

DriPP: Driven Point Processes to Model Stimuli Induced Patterns in M/EEG Signals

Thomas Moreau*
*INRIA, France

Abstract Temporal neuroimaging resort mostly on non-invasive electrophysiology signals from electroencephalography (EEG) and magnetoencephalography (MEG). The quantitative analysis of these signals boils down to the identification of temporal patterns such as evoked responses, transient bursts of neural oscillations but also blinks or heartbeats for data cleaning. Several works have shown that these patterns can be extracted efficiently in an unsupervised way, e.g. using Convolutional Dictionary Learning, even in low data regimes. This leads to an event-based description of the data. Given these events, a natural question is to estimate how their occurrences are modulated by certain cognitive tasks and experimental manipulations. To address it, we propose a point process approach. While point processes have been used in neuroscience in the past, in particular for single cell recordings (spike trains), techniques such as Convolutional Dictionary Learning make them amenable to human studies based on EEG/MEG signals. We develop a novel parametric statistical point process model – called driven temporal point processes (DriPP) – where the intensity function of the point process model is linked to a set of point processes corresponding to stimulation events. We derive a fast and principled expectation-maximization (EM) algorithm to estimate the parameters of this model. Simulations reveal that model parameters can be identified from long enough signals. Results on standard MEG datasets demonstrate that our methodology reveals event-related neural responses – both evoked and induced – and isolates non-task-specific temporal patterns. The use of a parametric kernels with few parameters for our model makes the method sound even in the low-resource setting in the MEG experiment, with only a few dozens of stimuli events. These results were presented in ICLR22; openreview.net/forum?id=d_2lcDh0Y9c.

Interferometric Lensless Imaging - Rank-one Projections of Image Frequencies with Speckle Illuminations

Olivier Leblanc*, **Matthias Hofer**†, **Siddharth Sivankutty**‡, **Hervé Rigneault**†, and **Laurent Jacques***
*Université catholique de Louvain, Belgium, †Institut Fresnel, France, ‡Université de Lille, France

Abstract Lensless endoscopy with multicore fibers (MCF) is a technology that enables fluorescent imaging of biological samples at the cellular scale. In this work, we show that under a common far-field approximation, the corresponding speckle imaging process is tantamount to collecting multiple symmetric rank-one projections (SROP) of an Hermitian interferometric matrix—a matrix encoding a subsampling of the Fourier transform of the sample image. Specifically, each SROP of this matrix is achieved with the complex vector shaping the incident wavefront using a spatial light modulator (SLM), and, for specific MCF core arrangements, the interferometric matrix collects as many image frequencies as the square of the core number. We improve the stability of these SROP measurements (a process known as debiasing) by dropping out the mean of the observed sample, hence avoiding doubling the number of SROP measurements as proposed in the literature. Configuring the SLM randomly allows us to characterize the sample complexity of the system through a two-component sensing perspective. In particular, by inspecting the separate dimensional conditions ensuring the specific restricted isometry properties of the two composing sensing models in the observation of sparse images, we show that a basis pursuit denoising (BPDN) program associated with an L1-fidelity cost provably provides a stable and robust image estimate. In parallel, numerical and experimental reconstruction results demonstrate the effectiveness of this imaging procedure for a TV-penalized formulation of the inverse problem. In experimental conditions, we develop a second SROP model fed by a calibration phase to predict any illuminating speckle in the sample plane from the complex amplitudes programming the SLM. This model avoids the need to know the cores positions and makes the imaging system robust to crosstalk, modefield deffects, and near-field conditions.

Session 8

Medical Image Reconstruction and Analysis

Organizers: Kerstin Hammernik & **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.*

Medical Image Reconstruction and Analysis - with a little help from AI

Kerstin Hammernik*
*Technical University of Munich, Germany

Abstract The aim of this talk is to guide the audience through the medical imaging pipeline - from acquisition and reconstruction to analysis - with a little help from AI. We start with formulating the inverse problem for medical image reconstruction, with the goal of reconstructing high-quality images from, i.e., low-dose Computed Tomography (CT) images, or undersampled k-space data as in Magnetic Resonance Imaging (MRI). In the next step, we show how traditional reconstruction approaches are linked to deep neural networks for solving inverse problems in imaging. Here, we showcase examples from MRI reconstruction and focus on the clinical application of the aforementioned approaches. Finally, we highlight the current challenges of deep neural networks for inverse problems and give an overview of current literature in linking both medical image reconstruction with further downstream tasks. Related work accessible via the DOIs: 10.48550/arXiv.2203.12215; 10.3389/fcvm.2022.826283; 10.1016/B978-0-12-816176-0.00007-7 and at elsevier.com/books/magnetic-resonance-image-reconstruction. Lecture notes on Inverse Problems in Medical Imaging by C. Graf and K. Hammernik accessible at youtube.com/channel/UCrE30RPowTwb6Lys_Iz34zg.

Data-driven model corrections and learned iterative reconstruction

Andreas Hauptmann*
*University of Oulu, Finland

Abstract Iterative model-based reconstruction approaches for high-dimensional problems in inverse problems with non-trivial forward operators can be highly time consuming. Thus, it is desirable to employ model reduction techniques to speed-up reconstructions in variational approaches as well as to enable training of learned model-based iterative reconstruction techniques. Nevertheless, reduced or approximate models can lead to a degradation of reconstruction quality and need to be accounted for. In this talk we discuss the possibility of learning data-driven model corrections for inverse problems. We make a distinction between implicit and explicit corrections. In the former, correction and regularization are trained simultaneously with the data-driven iterative updating network. This way a substantial speed-up, compared to classical variational approaches, can be achieved with improved reconstruction quality. Nevertheless, such implicit corrections offer limited insights into how approximate models are corrected for and hence we additionally discuss the latter case of learning and explicit correction that can be subsequently used in a variational framework. We will discuss the conceptual difficulty of learning such an explicit forward model correction and present conditions under which solutions to the variational problem with a learned correction converge to solutions obtained with the accurate model. Finally, we will leverage the obtained insights on learned model corrections and present recent work on a model corrected learned-primal dual scheme with applications to in-vivo 3D limited-view photoacoustic tomography. On Learned Operator Correction in Inverse Problems see epubs.siam.org/doi/10.1137/20M1338460.

Freeze it: Estimating and compensating motion in MRI

Thomas Küstner*
*University Hospital Tübingen, Germany

Abstract The challenge of motion in magnetic resonance (MR) imaging will be presented, starting with the various sources of motion in MR and how they impact the acquisition and reconstruction processes. The mechanisms behind motion artifacts will be described, recognizing motion artifacts as superposition of aliased images in different motion states. Various types of motion models and physiological motion challenges will be considered, along with the typical artifacts encountered. General prospective and retrospective approaches to deal with motion are depicted. The fundamentals of estimating and correcting motion in MR will be discussed, with a strong focus on retrospective motion correction strategies, including deep learning-based solutions. The generalized forward model for an MR acquisition considering motion will be introduced. This model will be used to: (1) characterize the profile of the aliasing introduced due to motion and understand how the relationship between the geometry of the motion and the geometry of the sampling trajectory determines these artifacts; (2) provide a reconstruction model to correct for generalized motion occurring during the MR acquisition. The talk will further discuss practical and general approaches to estimate and correct for motion, including techniques like triggering and gating, motion binning, image registration, k-space motion correction, deep learning-based image registration and motion corrected reconstructions. Additionally, hands-on code and in-vivo datasets are provided (github.com/lab-midas/ismrm-moco-workshop), with examples focusing on retrospective motion correction covering three aspects: motion artifact appearance (evaluate their behavior for various model geometries and trajectories), motion estimation via image registration (conventional and deep learning) and motion corrected reconstructions (study the properties of a motion model in the presence of noise, undersampling or model errors).

Image quality transfer and democratization of MRI

Daniel Alexander*
*University College London, UK

Abstract My talk will focus on Image Quality Transfer [1,2,3] and its role in international efforts to democratize MRI expertise and capability, such as CAMERA network [4]. The technique aims to 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. We have recently be adapting the idea towards the problem of estimating from an MRI acquired using a low-field system in clinics in lower-middle-income countries (LMICs), the image we would have obtained by flying the patient to e.g. London to image them on a state of the art high field system. I will talk through the history of development 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 (see marketing.webassets.siemens-healthineers.com/...ISMRM_2022, and cds.ismrm.org/.../2021-02-25-Africa, for background information). Related work accessible via the DOIs: [1] 10.1016/j.neuroimage.2017.02.089, [2] 10.1016/j.neuroimage.2020.117366, [3] 10.48550/arXiv.1909.06763, and [4] 10.1101/2022.05.02.22274588,

Uncertainty Estimation in Medical Image Analysis and Reconstruction

Ender Konukoglu[*], **Kerem Tezcan**[*], and **Katarina Tothova**[*]
[*]ETH Zurich, Switzerland

Abstract Estimating uncertainties associated with algorithmic predictions is an important and challenging task in medical image computing. The importance mainly stems from the fact that inaccurate, yet confident predictions can misguide users and lead to mistakes. The challenge arises from the complexity of the state-of-the-art algorithms and ambiguity of the data on which inference is usually based. In this work, we briefly describe the problem of uncertainty estimation using neural networks and highlight the challenges. We then describe two sampling-based approaches: (i) model-based and (ii) post-hoc sampling with an approximate posterior. In both of approaches we use a generic Bayesian formulation of the problem. We particularly focus on the posterior $p(y|x, w, D, M)$ (with y, x, w, D and M being the output, input, parameters, training set and model, respectively) that captures the ambiguity with respect to data on which inference is based. For the model-based approach, we focus on reconstructing MRI from undersampled acquisitions. We define the posterior through integrating a neural network-based probabilistic prior on the expected output y and a forward model that describes how the input x can be generated from an output y. This approach allows posterior sampling in a latent space that is of lower dimension than the imaging data, facilitating efficient sampling. For the post-hoc sampling approach, we focus on predicting shape from images. Here, a prior over the outputs can be modeled as before, but a forward model that describes how the input can be generated from the output is not trivially modeled. This is a case that is more common among applications of neural networks in medical imaging and beyond. We define a likelihood model through approximately inverting a trained neural network and use the resulting posterior for sampling. Associated publications are accessible via the DOIs: 10.1109/TMI.2022.3150853, and 10.1007/978-3-031-16749-2_2.

The role of data and models for deep-learning based image reconstruction

Reinhard Heckel[*], **Mohammad Zalbagi Darestani**[†], **Tobit Klug**[*], **Jiayu Liu**[*], and **Akshay S. Chaudhari**[‡]
[*]Technical University of Munich, Germany, [†]Rice University, USA, [‡]Stanford University, USA

Abstract Deep-learning methods give state-of-the-art performance for a variety of imaging tasks, including accelerated magnetic resonance imaging. In this work, we discuss whether improved models and algorithms or training data are the most promising way forward. First, we ask whether increasing the model size and the training data improves performance in a similar fashion as it has in domains such as language modeling. We find that scaling beyond relatively few examples yields only marginal performance gains. Second, we discuss the robustness of deep learning based image reconstruction methods. Perhaps surprisingly, we find no evidence for neural networks being any less robust than classical reconstruction methods (such as l1 minimization). However, we find that both classical and deep learning based approaches perform significantly worse under distribution shifts, i.e., when trained (or tuned) and tested on slightly different data. Finally, we show that the out-of-distribution performance can be improved through more diverse training data, or through an algorithmic intervention called test-time-training. Associated papers available on ArXiv (2209.13435, 2204.07204, 2102.06103).

Hybrid learning of Non-Cartesian k-space trajectory and MR image reconstruction networks

Chaithya Giliyar Radhakrishna[*] and **Philippe Ciuciu**[*†]
[*]CEA - Saclay, France, [*]INRIA Parietal, France

Abstract Compressed sensing in MRI involves the optimization of (1) the sampling pattern in k-space under MR hardware constraints and (2) image reconstruction from the undersampled k-space data. To tackle the first sub-problem, the Spreading Projection Algorithm for Rapid K-space sampLING (SPARKLING) algorithm was introduced which optimizes non-Cartesian sampling patterns compatible with MR hardware constraints on maximum gradient amplitude and slew rate. However, one major drawback of SPARKLING was the need for a target sampling density as an input. Through retrospective studies, we show how a data-driven and reconstruction algorithm-aware method performs the best in learning a good density. Recently, deep learning methods have allowed the community to address both the above sub-problems simultaneously, especially in the non-Cartesian acquisition setting. As a next step, we try to prevent the 2 step process of learning the density and then generating SPARKLING trajectories by developing a generic model to directly learn hardware-compliant k-space trajectories. We perform ablation studies using parameter-free reconstructions like the density compensated (DCp) adjoint operator of the nonuniform fast Fourier transform (NUFFT) to ensure that the learned k-space trajectories actually sample the center of k-space densely. Inspired by SPARKLING, we optimize these trajectories by embedding a projected gradient descent algorithm over the MR hardware constraints. Later, we introduce a novel hybrid learning (HL) approach that operates across multiple resolutions to jointly optimize the deep learning architecture for image reconstruction and the k-space trajectory. This HL method presents an improved image reconstruction quality at 20-fold acceleration factor on the fastMRI dataset with SSIM scores of nearly 0.92-0.95 in our retrospective studies as compared to corresponding Cartesian reference. Further, we observe a 3-4dB gain in PSNR as compared to earlier state-of-the-art methods.

Joint Cryo-ET Alignment and Reconstruction with Neural Deformation Fields

Valentin Debarnot[*], **Sidharth Gupta**[†], **Konik Kothari**[†], and **Ivan Dokmanić**[*]
[*]University of Basel, Switzerland, [†]University of Illinois at Urbana-Champaign, USA

Abstract We propose a framework to jointly determine the deformation parameters and reconstruct the unknown volume in electron cryotomography (CryoET). CryoET aims to reconstruct three-dimensional biological samples from two-dimensional projections. A major challenge is that we can only acquire projections for a limited range of tilts, and that each projection undergoes an unknown deformation during acquisition. Not accounting for these deformations results in poor reconstruction. The existing CryoET software packages attempt to align the projections, often in a workflow which uses manual feedback. Our proposed method sidesteps this inconvenience by automatically computing a set of undeformed projections while simultaneously reconstructing the unknown volume. We achieve this by learning a continuous representation of the undeformed measurements and deformation parameters. We show that our approach enables the recovery of high-frequency details that are destroyed without accounting for deformations. More detail can be found in the paper "Joint Cryo-ET Alignment and Reconstruction with Neural Deformation Fields" under review. Preprint available on ArXiv (2211.14534).

A framework for self-supervised MR image reconstruction using sub-sampling via Noisier2Noise

Charles Millard*, and **Mark Chiew***
*University of Oxford, UK

Abstract In recent years, there has been attention on leveraging the statistical modeling capabilities of neural networks for reconstructing sub-sampled MRI data. Most proposed methods assume the existence of a representative fully-sampled dataset and use fully-supervised training. However, for many applications, fully sampled training data is not available, and may be highly impractical to acquire. The development and understanding of self-supervised methods, which use only sub-sampled data for training, are therefore highly desirable. In this work we extend the Noisier2Noise framework, which was originally constructed for self-supervised denoising tasks, to variable density sub-sampled MRI data. We demonstrate that the framework can be used to analytically explain the performance of Self-Supervised Learning via Data Undersampling (SSDU), a recently proposed method that performs well in practice but until now lacked theoretical justification. In particular, we prove that a network trained with SSDU and an L2 loss correctly estimates k-space in expectation in regions not sampled in the input. Further, we show that improved reconstruction quality is possible when the distribution of the mask used to further sub-sample k-space is matched to the distribution of the original mask. We trained a Variational Network (VarNet) on the fastMRI brain dataset at acceleration factors 4 and 8 using Noisier2Noise, the original and modified versions of SSDU, and, as a benchmark, fully-supervised training. Our proposed modified SSDU had substantially improved reconstruction quality over the competing self-supervised methods, offering a test set mean-squared-error within 1% of fully supervised training for both acceleration factors. We also found that modified SSDU was more robust to parameter tuning of the second mask than competing methods. Our PyTorch implementation is available at github.com/charlesmillard/Noisier2Noise_for_recon and a full paper preprint is on ArXiv (2205.10278).

iFind, you find ... On fetal ultrasound compounding

Julia A. Schnabel*†
*Technical University of Munich, Germany, †Helmholtz Center Munich, Germany

Abstract Fetal imaging is conventionally carried out using ultrasound as the modality of choice, due to its non-ionising nature, real-time acquisition, portability, low cost and wide availability. However, it also requires significant operator skills and can be of variable image quality, making accurate manual measurements that are required for fetal biometrics often challenging. Deep learning has proven to be a game changer in fetal ultrasound, as it can directly operate on the incoming ultrasound video stream in near-real time, allowing for online semantic detection, labeling, measurements, and ultimately, clinical reporting. However, most fetal ultrasound is 2D, hampering volumetric measurements. In contrast, 3D ultrasound has poorer image resolution and is therefore of limited diagnostic value in the clinic. In general, both 2D and 3D US suffer from a limited field of view (FoV) at late stages of gestation. As part of the interdisciplinary ifindproject.com on intelligent fetal imaging and diagnosis, we have focused on advancing the quality and diagnostic value of 3D ultrasound imaging using deep learning frameworks. As part of the standard screening scans at 20 weeks, we provide a motion-reconstructed fetal head compounding method using free-hand 3D US, yielding unprecedented image quality. To address the problem of limited FoV at late gestation, we use multiple US probes to allow, for the first time, 3D whole placenta US imaging and volumetric measurements at late gestation. Related work accessible via the DOIs: 10.1007/978-3-030-32248-9_43, and 10.1016/j.media.2022.102639.

Session 9

Potential Pitfalls of Deep Learning in Medical Image Reconstruction

Organizers: Efrat Shimron & Florian Knoll

Summary *This session focuses on scenarios in which deep learning algorithms developed for medical imaging might produce unreliable results, e.g. due to distribution shifts, bias, hallucinations, or other factors. The session is planned following the increasing interest in studying sensitivities and instabilities of such algorithms. The talks will discuss strategies for exposing algorithmic sensitivities and addressing them, preventing inverse crimes, and increasing algorithmic interpretability. The aim of the session is to raise awareness to the growing problem of unreliable AI performance in the context of medical imaging, suggest guidelines and solutions, and invoke community discussions.*

The Risk in Naive Training of Medical AI Algorithms: Pitfalls, Stability and Generalizability Issues

Efrat Shimron*
*University of California - Berkeley, USA

Abstract The development of deep learning (DL) algorithms for image reconstruction from limited measurements has shown highly promising results in a range of fields, including computational and medical imaging. However, their development depends on the availability of suitable data, which can be difficult to obtain. Unfortunately, the use of limited or unsuitable data for algorithm development and evaluation can lead to issues such as bias, instability, hallucinations, and limited generalizability. This talk will review such barriers in the development of DL algorithms, and will discuss techniques for identifying and mitigating them. The first half of the talk will be dedicated to the topic of implicit bias stemming from "off-label" data use, where data published for one task is used for a different one. We will demonstrate how such use could lead to overly optimistic results due to hidden data preprocessing pipelines that are commonly applied for open-access databases. We will also show that algorithms pre-trained on off-label data could exhibit significant performance gaps for real-world data. To raise awareness to the growing problem off-label data use, we will discuss the concept of "data crimes", which refers to publication of biased results. The second half of this talk will give an introduction to the next presentations in this session. It will highlight topics related to use of adversarial approaches for ensuring trust-worthy image reconstruction, investigation of generalization ability and robustness to distribution shifts, the risk of hallucinations in reconstructed images, and the importance of signal modeling for MRI artifact mitigation.

Trustworthy Image Reconstruction and beyond: Using Adversarial Approaches

Daniel Rueckert*, **Cheng Ouyang**†, **Kerstin Hammernik***, and **Chen Chen**†
*Technical University Munich, Germany, †Imperial College London, UK

Abstract We investigate the robustness of deep-learning-based MRI reconstruction and segmentation models when the model is tested on unseen image contrasts and pathologies. We then investigate how to best achieve the generalization of the networks. In the case of MRI reconstruction, we propose to generalize the network by training with large publicly-available natural image datasets with synthesized phase information which achieves high cross-domain reconstruction performance which is comparable with those from domain-specific training. In the case of MRI segmentation, we show that adversarial data augmentation which generates plausible and realistic signal corruptions outperforms conventional pixelwise adversarial data augmentation.

Generalizability (or not?) of patch-based image models

Jeffrey A Fessler[*], **Jason Hu**[*], and **Xiaojian Xu**[*]
[*]University of Michigan, USA

Abstract Under-determined inverse problems require some form of prior information. Traditional regularization methods like total variation (TV) do not fully exploit the information in databases of acquired images. Generative prior models have garnered great interest in computer vision and in the image reconstruction field. An expressive generative model of recent interest is score-based diffusion models (see, e.g., Ramzi et al., NeurIPS 2020, Jalel et al., NeurIPS 2021, Song et al, ICLR 2021 & 2022, Chung&Ye, MIA, 2022). Such models can provide uncertainty estimates. Existing approaches for applying score-based diffusion models to inverse problems have focused on priors for entire images. As one goes from 2D to 3D to dynamic 3D imaging, it seems increasingly difficult both computationally and statistically to learn reliable priors for entire images, possibly increasing the opportunity for hallucinations. Such priors also seem to be designed for a single image size and may not generalize readily to other sizes. That limitation could be particularly problematic in dynamic imaging. This work focuses on learning priors for image patches. Specifically we learn a patch-score model from a modest number of training images, because even a small number of images contain many patches. We then combine the patch scores in a way that is akin to how Markov Random Field priors are a product of clique terms that typically involve patches. This combination can accommodate essentially arbitrary image sizes. The patch size is a design variable that presumably controls trade-offs between computation time, expressiveness, and generalizability. The smallest 2-pixel patch size is akin to TV regularization, whereas the largest patch size would be the entire image. Thus the patch-score framework unifies these two extremes. We investigate the trade-offs between generalizability and image reconstruction quality for patch-based score models.

The importance of signal modeling for MRI artifact mitigation

Jonathan I. Tamir[*]
[*]The University of Texas at Austin, USA

Abstract While deep learning is a powerful tool for solving inverse problems, careful attention must be placed on the underlying assumptions during training and their applicability to inference. This is especially prudent for MRI, in which the raw data are a function of both signal and spatial encodings, which naturally vary across different scanning protocols and institutions. In this talk we highlight some of these considerations, and in particular focus on the effects of the temporal ordering as well as the signal decay present in acquired k-space. We highlight the consequences of improper signal modeling and demonstrate training and inference methods which correctly account for the variability in data acquisition. Finally, we propose methods for deep learning-based signal inversion that make weak assumptions about the signal model during training so as to reduce overfitting at inference. The framework can be used to handle various distribution shifts during inference, including due to anatomy, changes in sampling, and subject motion.

Generalized hardness of approximation and hallucinations – On barriers and paradoxes in image reconstruction

Anders Hansen*
*University of Cambridge, UK

Abstract AI techniques are transforming medical imaging with striking performance. However, these new methods are susceptible to AI generated hallucinations, the phenomenon where realistically looking Artifact are incorrectly added to the reconstructed image, causing serious concerns in the sciences. The basic question is therefore: can hallucinations be prevented? In this talk we will discuss how hallucinations can occur as a result of the reconstruction method being unstable, but also how perfectly stable methods can hallucinate. Moreover, there is a trade-off between performance and both instability and hallucination. Indeed, methods that perform too well on certain inputs will either become unstable or will hallucinate. Optimizing this trade-off turns out to be highly delicate. In particular, for inverse problems one can typically show that there exist optimal neural networks. However, the phenomenon of generalized hardness of approximation prevents in general the optimal neural network from being computed to any accuracy below a certain approximation threshold. Thus, paradoxically, there will exist deep learning methods that are provably optimal, but they can only be computed to a certain accuracy. This phenomenon is similar in spirit to hardness of approximation that transformed computer science over the last decades. Fortunately, it is possible to both show existence of optimal AI methods that can also be computed by algorithms, but only under specific conditions on the inverse problem. Related publications accessible via DOIs: 10.1073/pnas.2107151119, 10.1073/pnas.1907377117, 10.48550/arXiv.2110.15734, and 10.48550/arXiv.2209.06715.

Mitigating Data Paucity and Distributional Shifts for Accelerated MRI Alongside New Clinically-Relevant Evaluation Metrics

Akshay S. Chaudhari*, **Beliz Gunel***, **Batu Ozturkler***, **Phil Adamson***, **Zhongnan Fang**†, **Reinhard Heckel**‡, **Eric Gibbons**⋆, **Garry Gold***, **Brian Hargreaves***, **Shreyas Vasanawala***, **John Pauly***, and **Arjun Desai***
*Stanford University, USA, †LVIS Corporation, USA, ‡Technical University of Munich, Germany, ⋆Weber State University, USA

Abstract Datasets such as fastMRI and mridata.org that consist of fully-sampled k-space data for MRI acquisitions have enabled developing machine learning algorithms for high-quality accelerated MRI reconstruction. However, such datasets encompass only a small subset of all anatomies and sequences used in routine clinical practice. In this work, we demonstrate that even high-performing reconstruction algorithms are susceptible to distributional shifts when images at test-time correspond to different data distributions than those at train-time. We will describe a new semi-supervised learning framework that leverages MRI physics to tackle such distributional shifts in a data-efficient manner [1–3]. We further present the Stanford Knee MRI with Multi-Task Evaluation (SKM-TEA) dataset that consists of raw k-space data, along with image classification, detection, and segmentation labels for the reconstructed images. Such clinically-tailored evaluation metrics allow training reconstruction algorithms for true clinical utility rather than surrogate image quality metrics with weak correlation to radiologist-perceived image quality [4]. In scenarios where such clinically-relevant image labels are unavailable, we describe a new data-driven image quality metric that is trained in a self-supervised manner on MR images. This metric has significantly better correlation to radiologist-perceived image quality than commonly used metrics in computer vision [5]. Overall, our work presents a framework for accelerated MRI without requiring large fully-sampled datasets to produce algorithms that are have superior robustness to test-time distributional shifts and correspondence to downstream clinically-relevant metrics. Related work accessible on ArXiv ([1] 2102.06103, [2] 2110.00075, [3] 2111.02549, [4] 2203.06823), and at [5] openreview.net/orum?id=dgMvTzf6M_3.

Physics-Driven Data Priors for Robust Self-Supervised Accelerated MRI Reconstruction

Arjun Desai*, **Beliz Gunel***, **Batu Ozturkler***, **Harris Beg**†, **Deepro Pasha**†, **Brian Hargreaves***, **Shreyas Vasanawala***, **John Pauly***, **Christopher Re***, and **Akshay S. Chaudhari***
*Stanford University, USA, †California Institute of Technology, USA

Abstract Deep learning has enabled improved image quality and faster inference times for accelerated MRI reconstruction. However, these methods are predominantly supervised, making them intractable to train in cases where fully-sampled data is nonexistent. While few self-supervised methods have alleviated the need for fully-sampled scans, these methods are sensitive to distribution shifts resulting from routine acquisition-related perturbations (e.g. noise, motion) or changes in scan protocol and hardware. In this work, we propose a self-supervised method, termed VORTEX-SS, for accelerated MRI reconstruction that learns physics-aware data priors to increase robustness to clinically relevant distribution drifts. Specifically, VORTEX-SS leverages domain knowledge of the forward MRI data acquisition model with physics-driven, acquisition-based data augmentations for consistency training. Our method (1) does not require any fully-sampled training data, (2) is compatible with any network architecture, and (3) enables composing a heterogeneous set of image- and physics-based augmentations. Among in-distribution scans, VORTEX-SS achieved comparable performance to supervised and self-supervised methods. However, VORTEX-SS also substantially improved performance among scans with varying extents of perturbation-related distributions shifts, such as changes in SNR and roto-translational patient motion Artifacts. Additionally, VORTEX-SS outperformed all baselines when generalizing to distribution shifts not seen during training, including changes in anatomy, field strength, acceleration factors, and acquisition trajectories. In conclusion, we demonstrate our self-supervised method can increase robustness to diverse, unseen test-time distribution shifts without requiring fully-sampled training data. Preprints available on ArXiv (2111.02549, 2110.00075). Code accessible at github.com/ad12/meddlr.

MRI Reconstruction via Data-Driven Markov Chains with Joint Uncertainty Estimation

Guanxiong Luo*, **Martin Heide***, **Moritz Blumenthal***, and **Martin Uecker**†
*University Medical Center Göttingen, Germany, †Graz University of Technology, Austria

Abstract The application of generative models in MRI reconstruction is shifting researchers' attention from the unrolled reconstruction networks to the probablistically tractable iterations and permits an unsupervised fashion for medical image reconstruction. In our preprint, we formulate the image reconstruction problem from the perspective of Bayesian inference, which enables efficient sampling from the learned posterior probability distributions. Different from conventional deep learning-based MRI reconstruction techniques, samples are drawn from the posterior distribution given the measured k-space using the Markov chain Monte Carlo (MCMC) method. In addition to the maximum a posteriori (MAP) estimate for the image, which can be obtained with conventional methods, the minimum mean square error (MMSE) estimate and uncertainty maps can also be computed. The data-driven Markov chains are the reverse of a diffusion process that transfers the distribution of data to a known distribution. The reconstruction is the reverse process, conditional on the k-space measurement. The generative model learned from a given image database is independent of the forward operator that is used to model the k-space measurement. Then flexibility is guaranteed, which means the method can be applied to k-space acquired with different sampling schemes or receive coils using the same pre-trained models. Preprint available on ArXiv (2202.01479).

Partial privacy loss in machine learning: a statistical signal processing perspective

Tamara T. Mueller[*], **Stefan Kolek**[†], **Daniel Rueckert**[*], and **Georgios Kaissis**[*]
[*]Technical University Munich, Germany, [†]Ludwig Maximilian University of Munich, Germany

Abstract Machine learning has advanced medical imaging in numerous domains and yielded high benefits for diagnostics and therapy response prediction. However, when working with medical data, patient privacy needs to be considered. Differential privacy (DP) is the gold-standard to ensure private data protection. As a quantification tool, DP is concerned with the privacy loss of an individual. However, this metric is rather coarse, as different elements of the input dataset may contain attributes which are unequal in terms of "how private" they are. Also, as privacy loss is typically a worst-case metric over the entire dataset, it fails to capture the susceptibility of each individual to attacks trying to undermine their privacy. In this work, we introduce a novel metric, the partial privacy loss (PPL), quantifying the contribution of each input element to an individual's privacy loss under the Rényi DP interpretation. For images, this allows us to report a pixel-wise PPL, indicating which pixels contribute most to the overall privacy loss of an individual. By leveraging the connection between divergence measures and information geometry, we show that PPL is equivalent to the Fisher information from the model weights to the input pixels. Our work is similar to previous interpretations (Jacobian sensitivity and Fisher information loss) but is better aligned with the DP definition and much cheaper computationally, allowing its utilization even for deep networks. Experimentally, we demonstrate our metric for tabular and medical imaging data and show that PPL corresponds well to human-interpretable image features. In conclusion, we initiate an investigation into the connection between DP, statistical signal processing and information geometry. PPL can identify individuals at high risk of reconstruction attacks. In the future, we will investigate the connections of our work to algorithmic fairness and interpretability techniques. Preprint accessible on ArXiv (2211.10173).

Insights into the reliability of deep learning reconstructions with research challenges

Florian Knoll[*], **Matthew Muckley**[†], **Yvonne Lui**[‡], and **Daniel Sodickson**[‡]
[*]University of Erlangen-Nuremberg, Germany, [†]META AI, USA, [‡]New York University School of Medicine, USA

Abstract In 2016 deep learning techniques have been first introduced to solve the inverse problem of MR image reconstruction from accelerated acquisitions (Hammernik ISMRM 2016, Wang IEEE ISBI 2016, Hammernik MRM 2018). Since then, the field has grown substantially, and a wide range of methods have been developed and applied to a wide range of imaging applications. While they generally outperform non-data-driven methods by substantial margins, this comes at the cost of questions about generalizability and robustness (Antun PNAS 2020). In this work, we analyze the performance of deep learning image reconstruction with a series of experiments of an unrolled-iteration-based approach (Radmanesh Radiology AI 2022) as well as with an analysis of the results of the two fastMRI research challenges (Knoll MRM 2020, Muckley IEEE TMI 2021). Our results show that deep learning methods generally show the behavior of over-regularized compressed sensing reconstruction, because the influence of the prior on the reconstruction result becomes too strong. Artifacts produced by data-driven methods vary widely depending on the acceleration level and whether the test distribution matches the training distribution. For low accelerations within the training distribution, the dominating issue is the loss of fine image features, a behavior that is generally in line with results from compressed sensing. For high accelerations within-distribution, data-driven models exhibit a pseudo-normalization effect that arises from projecting to the model's feature-space representation of the training data. For high accelerations out-of-distribution, the models produce extreme artifacts that render the images completely undiagnostic.