Workshop

Mathematical Models for Plug-and-play Image Restoration

December, 7-8th, 2022

Detailed Schedule

Wednesday December 7th

- 8:45am-9am: Welcome of participants
- 9am-9:15am : Workshop opening

Morning (Chairs : Arthur Leclaire & Andrés Almansa)

- 9:15am-10am : **Samuel Hurault** (Institut de Mathématiques de Bordeaux)

  Tutorial session : *Introduction to Plug-and-Play methods and to their convergence analysis*

  Plug-and-Play (PnP) methods constitute a class of iterative algorithms for imaging problems where regularization is performed by an off-the-shelf denoiser. Specifically, given an image dataset, optimizing a function (e.g. a neural network) to remove Gaussian noise is equivalent to approximating the gradient or the proximal operator of the log prior of the training dataset. Therefore, any off-the-shelf denoiser can be used as an implicit prior and inserted into an optimization scheme to restore images. After introducing the PnP and Regularization by Denoising (RED) frameworks, we will explore the different convergence analyses that have been proposed in the literature. From monotone operators theory to convex and non-convex analysis, we will propose various tools and deep denoisers that allow theoretical PnP convergence guarantees and state-of-the-art IR performance.

- 10am-10:45am : **Regev Cohen** (Verily Israel, remotely)

  The Fixed Points of Plug-and-Play Methods

  The success of the frameworks Plug-and-Play Prior (PnP) and Regularization by Denoising (RED) has led to the rapid development of numerous inverse methods that exploit denoisers as priors. Yet, the convergence of RED and PnP, as well as the nature of their solutions, are not completely understood. In the first part of the talk we aim to shed some light on the solution of RED and PnP. To that end, we present RED-PRO: an inverse problem formulation which constrains the solutions to the fixed-point set of a denoiser. Surprisingly, RED-PRO is a convex optimization problem which can be solved via a simple iterative method. Through convergence analysis we establish the connections of RED-PRO to both PnP and RED, showing their solutions are directly related to the fixed-points of the denoiser. In the second part, we introduce potential-driven denoisers designed as gradients of neural networks. This new type of denoisers can be interpreted as maximum a posteriori or minimum mean square error esitmators, and can be used within any PnP and RED algorithms to ensure fast convergence.

- 10:45am - 11:15am : Break
- 11:15am - 11:40am : **Sebastian Neumayer** (EPFL Lausanne)

  Convex Regularizers based on Shallow Neural Networks

  In this talk, we will revisit the state-of-the-art in learned convex regularisation. As comparison, we propose a regulariser based on a one hidden layer neural network with (almost) free-form activation functions. For training this network, we rely on some connection to gradient based denoisers. Our numerical experiments indicate that this shallow architecture already achieves the best performance, which is very different from the nonconvex case. Moreover, even when learning both the filters and the activation functions, we recover wavelet-like filters and thresholding-like activation functions. These observations raise the question if the fundamental limit is already reached in the convex setting.

- 11:40am - 12:05pm : **Rita Fermanian** (Inria Rennes)

  PnP-ReG: Learned Regularizing Gradient for Plug-and-Play Gradient Descent

  The Plug-and-Play (PnP) framework makes it possible to integrate advanced image denoising priors into optimization algorithms, to efficiently solve a variety of image restoration tasks generally formulated as Maximum A Posteriori (MAP) estimation problems. The Plug-and-Play alternating direction method of multipliers (ADMM) and the Regularization by Denoising (RED) algorithms are two examples of such methods that made a breakthrough in image restoration. However, while the former method only applies to proximal algorithms, it has recently been shown that there exists no regularization that explains the RED algorithm when the denoisers lack Jacobian symmetry, which happens to be the case of most practical denoisers. We show that it is possible to train a network directly modeling the gradient of a MAP regularizer while jointly training the corresponding MAP denoiser. We use this network in gradient-based optimization methods and obtain better results comparing to other generic Plug-and-Play approaches. We also show that the regularizer can be used as a pre-trained network for unrolled gradient descent. Lastly, we show that the resulting denoiser allows for a better convergence of the Plug-and-Play ADMM.

Afternoon (Chair : Samuel Hurault)

- 1:35pm - 2:20pm : **Johannes Hertrich** (TU Berlin)

  Convolutional Proximal Neural Networks and Plug-and-Play Algorithms

  We introduce Proximal neural networks (PNNs) which are by construction averaged operators. The weight matrices in PNNs are located on the Stiefel manifold. For the training of dense PNNs and convolutional PNNs with full filter length, we propose a stochastic gradient descent on (a submanifold of) the Stiefel manifold. In the case of filters with limited length, we minimize a functional that peanalizes the distance to the Stiefel manifold and perform an additional projection step at the end of the training. Afterwards, we investigate how scaled cPNNs with a prescribed Lipschitz constant can be used for denoising images, where the achieved quality depends on the Lipschitz constant.

  We present two kinds of applications. First, we apply denoisers based on convolutional PNNs within a Plug-and-Play (PnP) framework and provide convergence results for the corresponding PnP forward-backward splitting algorithm based on an oracle construction. Second, we use the averagedness property of PNNs in order to construct a new architecture of normalizing flows based on invertible residual flows. Finally, we demonstrate the performance of PNNs in both applications by numerical examples.

- 2:20pm - 3:05pm : **Matthieu Terris** (Heriot-Watt University)

  Plug-and-play algorithms through the lens of monotone operators

  The lack of convergence guarantees for PnP algorithms can be problematic in some imaging applications. In this talk, we show that monotone operator theory provides a convenient theoretical framework for studying PnP algorithms. In particular, it allows to derive theoretical Lipschitz constraints on the denoiser to ensure convergence and characterisation of the limit point. These constraints can be relaxed and incorporated as a differentiable Jacobian (ℓ^2) norm regularisation term in the training loss function. This strategy can be applied to any feedforward architecture and is effective to enforce convergence despite being approximate. We show an application of the proposed method to (radio) astronomical imaging, where the accuracy of the reconstruction is key.

- 3:05pm - 3:30pm : **Jonathan Chirinos Rodriguez** (University of Genova)

  Learning Firmly Nonexpansive operators

  In the past years, Plug and Play (PnP) algorithms have received a huge attention in the Inverse Problems community because of their strong computational performance. In a few words, the main objective of a PnP method is, given an algorithm involving a proximal operator, to replace it by a (typically parametrized) function that acts as a denoiser. Such function may be either handcrafted or learned. In both situations, it has been observed that the problem of "learning a proximal operator of a convex function" can be generalized to and rephrased as "learning a firmly nonexpansive operator". Starting from this point, and despite some pioneering works (for instance, Ryu et al., Plug-and-Play Methods Provably Converge with Properly Trained Denoisers, 2019), very few theoretical studies have been developed: either the nonexpansivity condition is difficult to impose or no convergence result is provided.

  In this project, we model the problem from a statistical learning point of view. Specifically, we aim at finding the minimizer of the empirical error in the class of firmly nonexpansive operators from a given set of noisy/ground truth terms. Since such class is infinite-dimensional, we propose instead a constrained finite dimensional problem by considering classic numerical analysis techniques: simplicial partitions. Such solution will be ensured to be firmly nonexpansive. Moreover, we are able to provide a convergence result of our constructed operator to the solution of the empirical error. Finally, in the experimental results, we consider a primal-dual algorithm where the proximal map is replaced by the learned functional, and compare our method to classical denoising methods.

- 3.30pm - 4pm : Break

- 4pm - 4:45pm : **Kai Zhang** (ETH Zurich)

  Deep Plug-and-Play and Deep Unfolding Methods for Image Restoration

  Model-based methods and learning-based methods have been the two dominant strategies for solving various image restoration problems in low-level vision. In recent years, deep plug-and-play methods and deep unfolding methods have shown great promise by leveraging both learning-based methods and model-based methods. In this talk, I will present my work on deep plug-and-play and deep unfolding image restoration. I will also present my recent work on blind image super-resolution and denoising for real applications.

- 4:45pm - 5:10pm : **Alban Gossard** (Institut de Mathématiques de Toulouse)

  Training adaptive reconstruction networks for inverse problems and applications to blind inverse problems

  Neural networks are full of promises for the resolution of ill-posed inverse problems. In particular, physics informed approaches already progressively replace carefully hand-crafted reconstruction algorithms, for their superior quality. The aim of this presentation is threefold. First we show a limitation of these networks: when trained on a given forward operator, they do not generalize well to a different one. Second, we show that training the network with a family of forward operators solves the adaptivity problem without compromising the reconstruction quality significantly. Finally, we illustrate that this training procedure allows solving different blind inverse problems. This approach is compared with Plug & Play priors and some numerical illustrations are given.

- 5:10pm - 5:35pm : **Valentin Debarnot** (Basel University)

  Deep-Blur: Blind Identification and Deblurring with Convolutional Neural Networks

  We propose a neural network architecture and a training procedure to estimate blurring operators and deblur images from a single degraded image. Numerical experiments reveal that the proposed method recovers the blur parameters robustly even for large noise levels. This estimate can then be used as an input of an unrolled neural network to deblur the image. It is trained with a specific sampling procedure adapted to a family of parameterized operators. It considerably improves alternative blind deblurring softwares in the examples treated.

- 5:35pm - 6pm : **Charles Laroche** (GoPro & MAP5, Université Paris Cité)

  Deep Model-Based Super-Resolution with Non-uniform Blur

  We propose a state-of-the-art method for super-resolution with non-uniform blur. Single-image super-resolution methods seek to restore a high-resolution image from blurred, subsampled, and noisy measurements. Despite their impressive performance, existing techniques usually assume a uniform blur kernel. Hence, these techniques do not generalize well to the more general case of non-uniform blur. Instead, in this paper, we address the more realistic and computationally challenging case of spatially-varying blur. To this end, we first propose a fast deep plug-and-play algorithm, based on linearized ADMM splitting techniques, which can solve the super-resolution problem with spatially-varying blur. Second, we unfold our iterative algorithm into a single network and train it end-to-end. In this way, we overcome the intricacy of manually tuning the parameters involved in the optimization scheme. Our algorithm presents remarkable performance and generalizes well after a single training to a large family of spatially-varying blur kernels, noise levels and scale factors.

Thursday December 8

Morning (Chair : Andrés Almansa)

- 9am - 9:45am : **Julie Delon** (MAP5, Université Paris Cité)

  PnP sampling for inverse problems in imaging

  In a Bayesian framework, image models are used as priors or regularisers and combined to explicit likelihood functions to define posterior distributions. These posterior distributions can be used to derive Maximum A Posteriori (MAP) estimators, leading to optimization problems that are generally well studied and understood. Sampling schemes can also be used to explore more finely these posterior distributions, derive other estimators, quantify uncertainties or perform other advanced inferences. In a manner akin to Plug & Play (PnP) methods in optimization, these sampling schemes can be combined with denoising neural networks approximating the gradient of a log-prior on images. In this talk, after reminders on Bayesian sampling for inverse problems in imaging, we will focus on these PnP sampling schemes, which raise important questions concerning the correct definition of the underlying Bayesian models or the computed estimators, as well as their regularity properties, necessary to ensure the stability of the numerical schemes.

- 9:45am - 10:30am : **Marcelo Pereyra** (Institute for Mathematical Sciences & Heriot-Watt University)

  Bayesian inference with learnt generative image priors encoded by neural networks

  This talk presents a mathematical and computational methodology for performing Bayesian inference in problems where prior knowledge is available in the form of a training dataset or set of training examples. This prior information is encoded into the model by using a generative model, represented by deep neural network, which is combined with an explicit likelihood function by using Bayes' theorem to derive the posterior distribution for the quantities of interest given the available data. Bayesian computation is then performed by using appropriate Markov chain Monte Carlo stochastic algorithms. We study the proposed models and computation algorithms theoretically and empirically, and illustrate their performance on a range of imaging inverse problems involving point estimation, uncertainty quantification, hypothesis testing, and model misspecification diagnosis. Our experiments show that this is a highly promising line of research that needs to be explored carefully and with rigorous statistical tests, as some models appear excellent form a generative viewpoint but behave poorly as priors because they fail to properly capture the marginal distribution of the unknown image of interest. Joint work with Matthew Holden and Kostas Zygalakis.

- 10:30am - 11am : Break

- 11am - 11:25pm : **Jean Prost** (Institut de Mathématiques de Bordeaux)

  Diverse super-resolution with pretrained hierarchical Variational Autoencoders

  Image generative models provide a strong prior about natural images that can be exploited to solve ill-posed inverse problems. In this work we propose a method to perform super-resolution by reusing a pretrained hierarchical variational autoencoders. Specifically, we train a probabilistic encoder on low-resolution images, and we combine this new low-resolution encoder with VD-VAE high-resolution decoder to sample high-resolution images conditionned on a low-resolution input. We find that the latent hierarchical representation learned by VD-VAE naturally separates high-frequency from low-frequency information, allowing us to efficiently train the low-resolution encoder, by restraining it to the subset of the latent variables encoding the low-frequency information. Our method provides a model of the posterior distribution of high resolution images, allowing us to sample diverse plausible and realistic solutions. We demonstrate the ability of our method to generate diverse solutions to the super-resolution problem on face super-resolution with upsampling factors x4, x8, and x16.

- 11:25am - 11:50pm : **Charlesquin Kemajou** (Heriot-Watt University)

  Maximum marginal likelihood estimation of regularisation parameters in Plug-and-Play Bayesian estimation

  This talk presents an empirical Bayesian extension of the Plug & Play (PnP) Bayesian inference method of Laumont et al., 2022. A main novelty is that the amount of regularisation enforced by the prior, determined by the noise level parameter of the PnP denoiser, is estimated directly from the observed data by maximum marginal likelihood estimation (MMLE). The MMLE problem is computationally intractable. In a manner akin to Vidal et al., 2019, we address this difficulty by incorporating the PnP unadjusted Langevin algorithm of Laumont et al., 2022 within a stochastic approximation proximal gradient scheme that simultaneously calibrates the regularisation parameter by MMLE and generates samples asymptotically distributed according to the empirical Bayes (pseudo-)posterior distribution of interest. In addition, the method can also estimate other unknown model parameters by MMLE, such as the observation noise level or parameters of the forward operator. The proposed method is demonstrated with a range of non-blind and semi-blind image deconvolution problems.

Afternoon (Chair : Arthur Leclaire)

- 2pm - 2:45pm : **Ulugbek Kamilov** (Washington University in St Louis)

  Plug-and-Play Models for Large-Scale Computational Imaging

  Computational imaging is a rapidly growing area that seeks to enhance the capabilities of imaging instruments by viewing imaging as an inverse problem. Plug-and-Play Priors (PnP) is one of the most popular frameworks for solving computational imaging problems through integration of physical and learned models. PnP leverages high-fidelity physical sensor models and powerful machine learning methods to provide state-of-the-art imaging algorithms. PnP models alternate between minimizing a data-fidelity term to promote data consistency and imposing a learned image prior in the form of an "artifact reducing" deep neural network. This talk presents a principled discussion of PnP under inexact physical and learned models. Inexact models arise naturally in computational imaging when using approximate physical models for efficiency or when test images are from a different distribution than images used for training the image prior. We present several successful applications of our theoretical and algorithmic insights in bio-microscopy, computerized tomography, and magnetic resonance imaging.

- 2:45pm - 3:30pm : **Zahra Kadkhodaie** (New York University)

  Stochastic Solutions for Linear Inverse Problems using the Prior Implicit in a Denoiser

  Deep neural networks have provided state-of-the-art solutions for problems such as image denoising, which implicitly rely on a prior probability model of natural images. Two recent lines of work–Denoising Score Matching and Plug-and-Play–propose methodologies for drawing samples from this implicit prior and using it to solve inverse problems, respectively. Here, we develop a parsimonious and robust generalization of these ideas. We rely on a classic statistical result that shows the least-squares solution for removing additive Gaussian noise can be written directly in terms of the gradient of the log of the noisy signal density. We use this to derive a stochastic coarse-to-fine gradient ascent procedure for drawing high-probability samples from the implicit prior embedded within a CNN trained to perform blind denoising. A generalization of this algorithm to constrained sampling provides a method for using the implicit prior to solve any deterministic linear inverse problem, with no additional training, thus extending the power of supervised learning for denoising to a much broader set of problems. The algorithm relies on minimal assumptions and exhibits robust convergence over a wide range of parameter choices. To demonstrate the generality of our method, we use it to obtain state-of-the-art levels of unsupervised performance for deblurring, super-resolution, and compressive sensing.

- 3:30pm - 4pm : Break

- 4pm - 4:25pm : **Christophe Kervazo** (Télécom Paris)

  Unrolling PALM for sparse semi-blind source separation

  Unrolling PALM for sparse semi-blind source separation Sparse Blind Source Separation (BSS) has become a well established tool for a wide range of applications – for instance, in astrophysics and remote sensing. Classical sparse BSS methods, such as the Proximal Alternating Linearized Minimization (PALM) algorithm, nevertheless often suffer from a difficult hyperparameter choice, which undermines their results. To bypass this pitfall, we propose in this work to build on the thriving field of algorithm unfolding/unrolling. Unrolling PALM enables to leverage the data-driven knowledge stemming from realistic simulations or ground-truth data by learning both PALM hyperparameters and variables. In contrast to most existing unrolled algorithms, which assume a fixed known dictionary during the training and testing phases, this article further emphasizes on the ability to deal with variable mixing matrices (a.k.a. dictionaries). The proposed Learned PALM (LPALM) algorithm thus enables to perform semi-blind source separation, which is key to increase the generalization of the learnt model in real-world applications. We illustrate the relevance of LPALM in astrophysical multispectral imaging: the algorithm not only needs up to 10^4-10^5 times fewer iterations than PALM, but also improves the separation quality, while avoiding the cumbersome hyperparameter and initialization choice of PALM. We further show that LPALM outperforms other unrolled source separation methods in the semi-blind setting.

- 4:25pm - 4h50pm : **Maud Biquard** (ISAE-SUPAERO & CNES)

  Deep regularization for inverse problems in satellite imaging

  The most efficient image restoration methods currently available learn to invert a given degradation model in a supervised manner, with neural networks. Due to the large variety of on-board sensors and the drift of their characteristics over time that modify the degradation, this method seems hard to apply in the field of satellite imagery: it would require a model trained for each satellite considered and updated at regular time intervals. Instead, recent methods for image restoration propose an approach independent of the degradation model, by looking for the solution of the inverse problem in the latent space of a generative neural network. Generating all possible images would require a particularly powerful generative model, given the diversity and complexity of the satellite images. An alternative solution is to search for the restored image in the image space rather than in the latent space of the generative model. We compare this method, which we call analysis approach, to the synthesis method, which is the classical approach. Then, we propose to replace the variational autoencoder by an autoencoder that is more accurate in its reconstruction, but not necessarily generative, and to regularize the inverse problem with this network. We test our approach on reference datasets and satellite images.

- 4:50pm - 5:15pm : **Vasiliki Stergiopoulou** (Université Côte d'Azur)

  PnP COL0RME : Fluorescence image deconvolution microscopy using a Plug-and-Play Denoiser

  The spatial resolution of images of living samples obtained by fluorescence microscopes is physically limited due to the diffraction of visible light, making it difficult to study entities of size less than the diffraction barrier that is around 200 nm in the lateral plane. To overcome this limitation, super-resolution fluorescence microscopy techniques have been proposed in the literature. One of them, the method COvariance-based l0 super-Resolution Microscopy with intensity Estimation (COL0RME), forms a super-resolved image given a short temporal stack of frames obtained by common fluorescence microscopes. In COL0RME, we exploit the independence between distinct emitters as well as the sparse distribution of the fluorescent molecules via the use of an appropriate regularization term defined on the emitters' covariance matrix. Globally, the cost function to minimize is non-differential and non-convex. For its numerical minimization we use a proximal gradient-based iterative schema. With the interest of obtaining a reconstruction more adapted to sample geometries (e.g. filaments), in this work, we replace the proximity operator of the regularization term with an image denoiser (i.e. a pretrained network), following the Plug-and-Play (PnP) reconstruction framework. The image denoiser, differently from other PnP algorithms, is trained on the temporal autocovariance matrices of simulated sequences of fluctuating fluorescent molecules. The proposed method has been verified on simulated but also Ostreopsis fluorescence microscopy images.

- 5:15pm - 6pm : **Valentin De Bortoli** (CNRS, ENS Ulm)

  Tutorial session : *Denoising diffusion models for inverse problems*

  This tutorial is devoted to score-based generative modelling for inverse problems. These models to sample from given posterior distribution are adapted from methods used for generative modelling, i.e. the task of generating new samples from a data distribution. Score-based generative modelling is a recently developed approach to solve this problem and exhibits state-of-the-art performance on several image synthesis problems. These methods can be roughly described as follows. First, noise is incrementally added to the data to obtain an easy-to-sample distribution. Then, we learn the time-reversed denoising dynamics using a neural network. When initialized at the easy-to-sample distribution we obtain a generative model. These dynamics can be analyzed through the lens of stochastic analysis. In particular, it is useful to describe these processes as Stochastic Differential Equations (SDEs). The time-reversed SDE is a diffusion whose drift depends on the logarithmic gradients of the perturbed data distributions, i.e. the Stein scores. These scores are computed leveraging score-matching methods and in particular the Tweedie identity as well as neural network approximations. These generative models can be conditioned on observed data and give rise to efficient solvers for inverse problems. We will draw connections between these machine-learning models and the PnP methods introduced in image processing and present applications to some classical inverse problems in image processing.