Imperial College London

Model-Based Deep Learning for Inverse Problems

Pier Luigi Dragotti, Imperial College London

20 November 2023

Three Case Studies in Imaging Science

Imperial College London

Image restoration problems

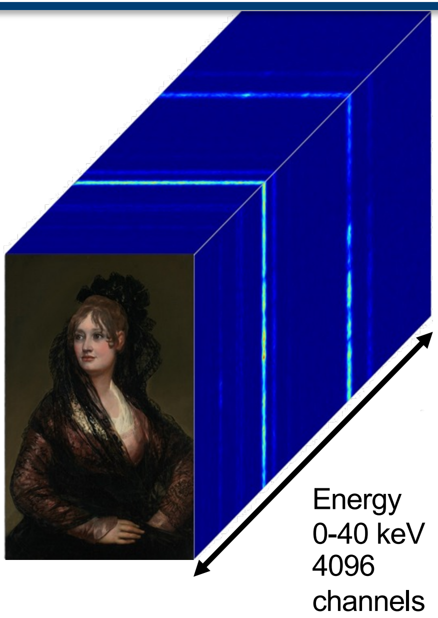

Technical study of Old Masters paintings

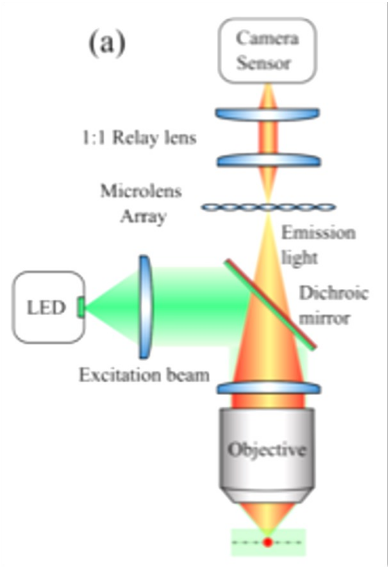

Light field microscopy for neuroscience

Imperial College London

Motivation: Computational Imaging

The complexity of modern imaging workflows calls for a rethink of imaging as an integrated sensing and inference model.

Seeing imaging as a whole is the domain of Computational Imaging

Key in computational imaging is the development of the interplay between physical and learned models

- Model-based approaches more interpretable and predictable, can reduce complexity
- Data-driven approaches can handle more complex settings

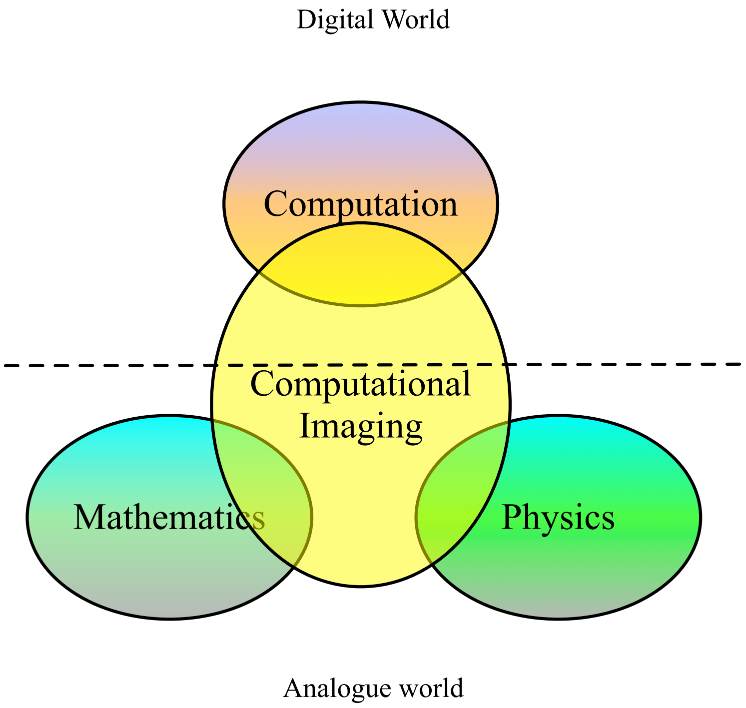

Imperial College London

Model-based Deep Learning

Plato: models, priors

Aristotle: data

Need to find the right
balance between
data and prior models

Model-based Deep Learning

- In inverse problems one looks for the right trade-off between a fidelity term and a prior
- $\hat{x} = \min_x ||H(x) - y||^2 + \lambda\rho(x)$

  fidelity term prior
- Models/physics can help with H and sometimes with $\rho(x)$
- Two key approaches to embed systematically priors and models into deep neural network architectures:
  - **Plug-and-play approach** → use neural networks as regularizers
  - **Deep Unfolding** → embed models and priors in the network architecture

Plug-and-play

- $\hat{x} = \min_{x} \|Hx - y\|^2 + \rho(x)$

  likelihood prior

- $\hat{x} = \min_{x,v} \|Hx - y\|^2 + \rho(v)$ s.t $x = v$

- Turn the constraint into a penalty: $\hat{x} = \min_{x,v} \|Hx - y\|^2 + \rho(v) + \beta\|x - v\|^2$

- Solve by alternating between x and v

  - Least-square: $\hat{x} = \min_{x} \|Hx - y\|^2 + \beta\|x - v\|^2$

  - A denoiser: $\hat{v} = \min_{v} \rho(v) + \beta\|x - v\|^2$

Use Deep Learning for denoising

- Venkatakrisnhan et al. Plug-and-play priors for model-based reconstruction, GlobalSip 2013
- Kamilov et al, Plug-and-Play Methods for Integrating Physical and Learned Models in Computational Imaging IEEE Signal Processing Magazine, 2023

Imperial College London

Wavelets and Invertible Neural Networks

- Wavelets provide sparse representations of piecewise smooth signals.
- This is why they have been successfully used in many imaging applications

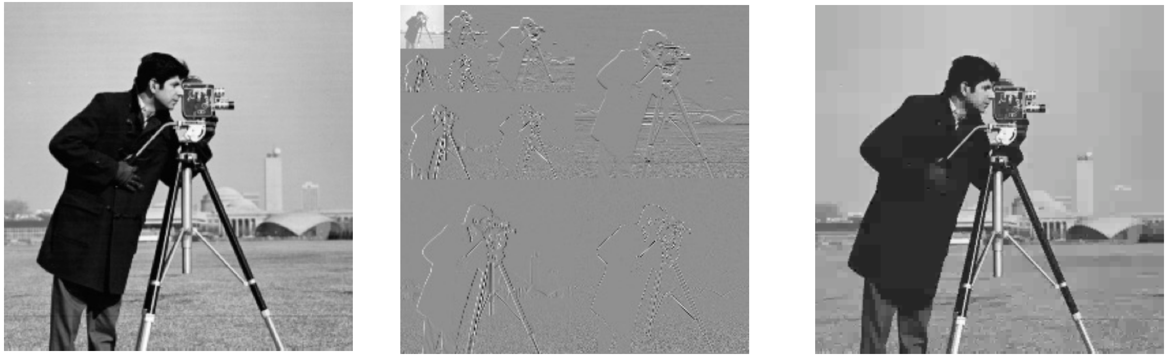

Figure: Cameraman is reconstructed using only 8% of the wavelet coefficients

Imperial College London

Implementation of the 2-D Wavelet Transform

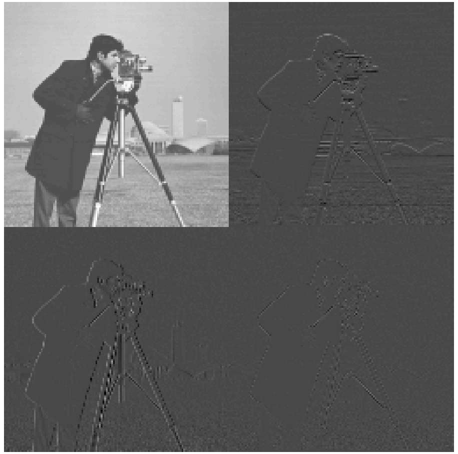

Implementation of the 2-D Wavelet Transform

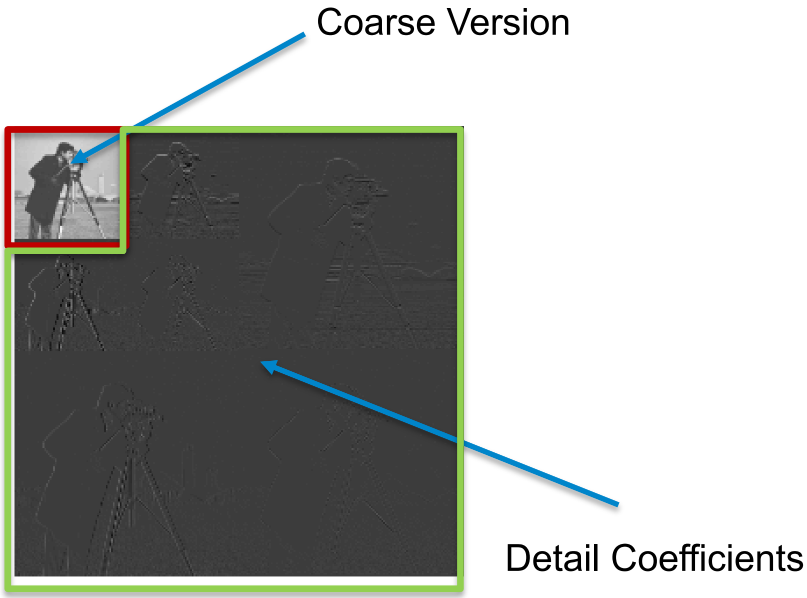

Wavelet-based Denoising

- Principles of wavelet denoising:

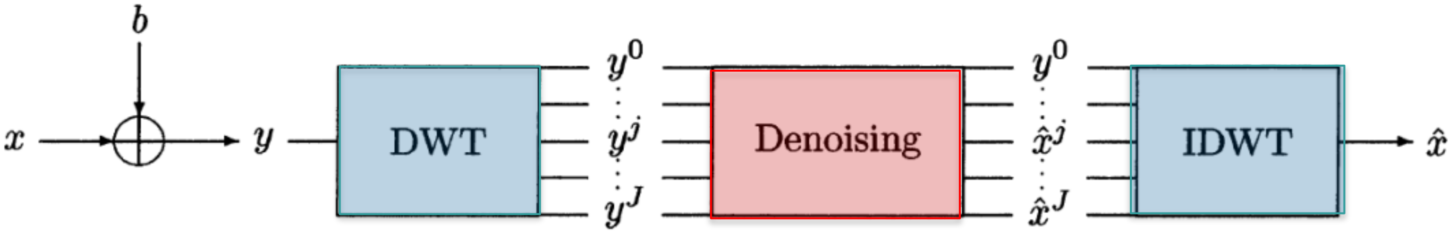

Wavelet transform

- Multi-resolution analysis
- Perfect reconstruction
- Noise is uniformly spread through the coefficients
- Image information is concentrated on small number of large coefficients

Denoising

- Element-wise thresholding, e.g. soft-thresholding

Wavelet-based Denoising

Imperial College London

1-D Example

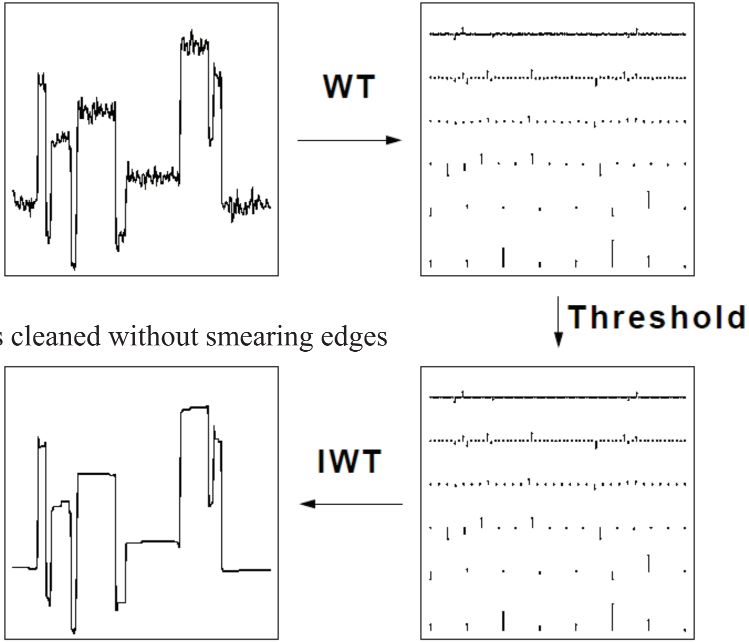

Smooth parts cleaned without smearing edges

What are Invertible Neural Networks?

Imperial College London

- Bijective (invertible) function approximators that have a forward mapping

$$F_\theta : \mathbb{R}^d \rightarrow \mathbb{R}^l$$
$$x \mapsto z$$

- and inverse mapping

$$F_\theta^{-1} : \mathbb{R}^l \rightarrow \mathbb{R}^d$$
$$z \mapsto x$$

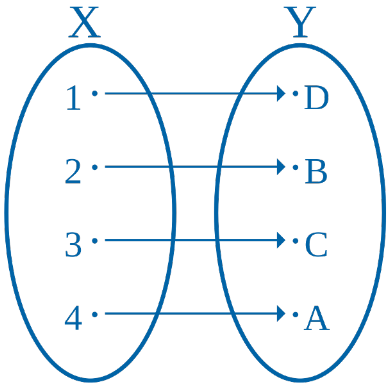

A bijective function (or invertible function)

Imperial College London

What are Invertible Neural Networks?

- INNs are bijective function approximators

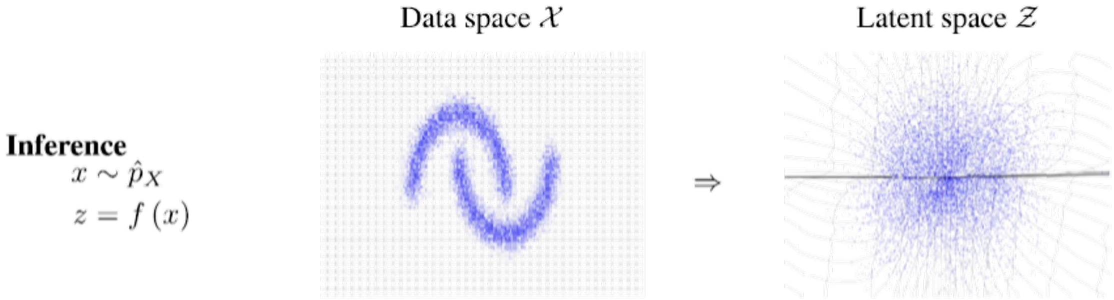

Imperial College London

How to Achieve Invertibility?

- Invertible via lifting scheme like architectures
  - Signal splitting
  - Alternate prediction and update

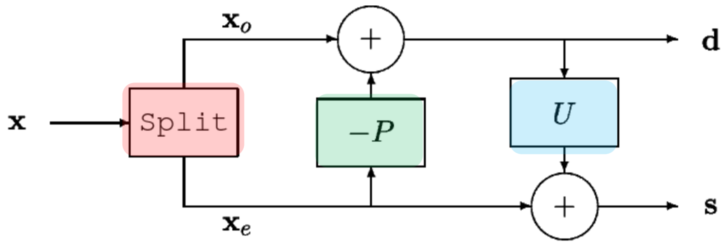

Split → $\begin{cases} d = x_o - P(x_e) \\ s = x_e + U(d) \end{cases}$

Forward pass

$\begin{cases} x_o = d + P(x_e) \\ x_e = s - U(d) \end{cases}$ → Merge

Backward pass

Imperial College London

Wavelets and INN

- The wavelet transform can be implemented using the lifting scheme

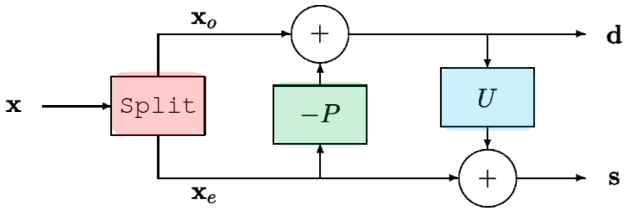

Split → $\begin{cases} d = x_o - P(x_e) \\ s = x_e + U(d) \end{cases}$

Forward pass

$\begin{cases} x_o = d + P(x_e) \\ x_e = s - U(d) \end{cases}$ → Merge

Backward pass

- The predictor (P) predicts the odd samples using the even, the update (U) uses the prediction error to smooth the even samples
- Predictor/update are fixed
- The scheme is perfectly invertible

I. Daubechies and W. Sweldens, “Factoring Wavelet Transforms into lifting Steps” 1997

Wavelets and INN

- Can we learn a wavelet-like non-linear sparsifying transform?

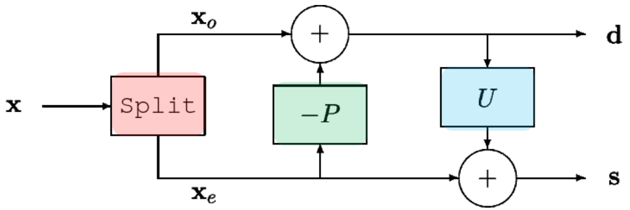

Split → $\begin{cases} d = x_o - P(x_e) \\ s = x_e + U(d) \end{cases}$ $\begin{cases} x_o = d + P(x_e) \\ x_e = s - U(d) \end{cases}$ → Merge

Forward pass Backward pass

- Approach:
  - convert the P/U operators into two deep networks and learn them
  - Use denoising as the bottleneck to impose sparsity

Imperial College London

Wavelets and INN

- Can we learn a wavelet-like non-linear sparsifying transform?

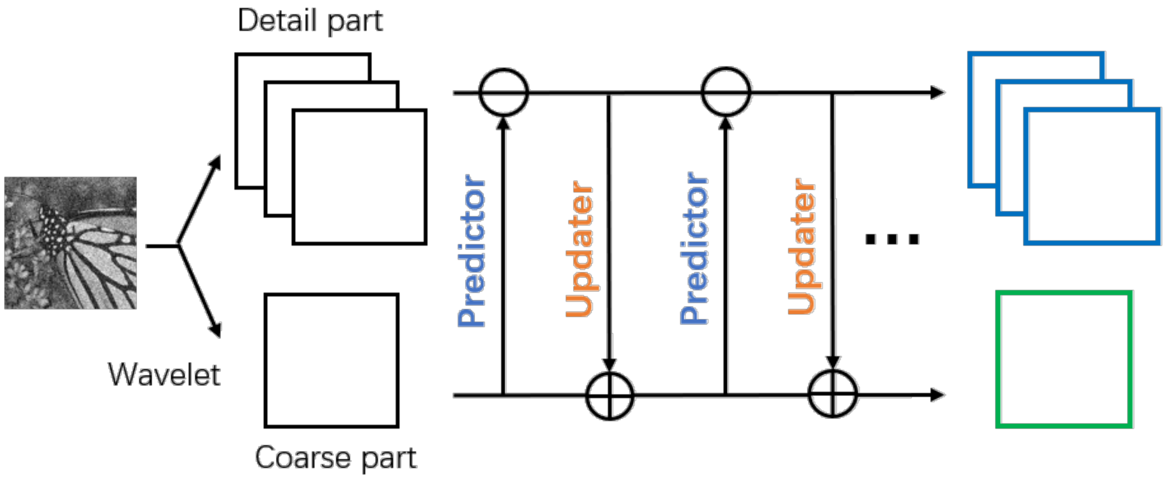

- Approach:
  - convert the P/U operators into two deep networks and learn them
  - Use denoising as the bottleneck to impose sparsity

Imperial College London

Wavelets and INN

- To make sure P acts as a sparsifying predictor:
  - Train the network with noisy/noiseless image pairs
  - Add a denoising network on the details

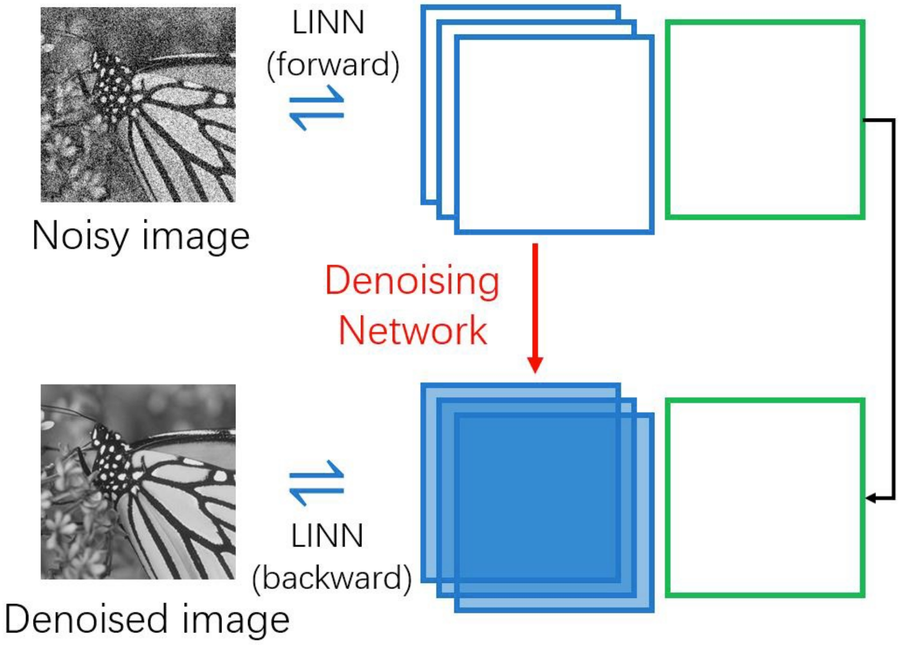

Denoising - Overall Method

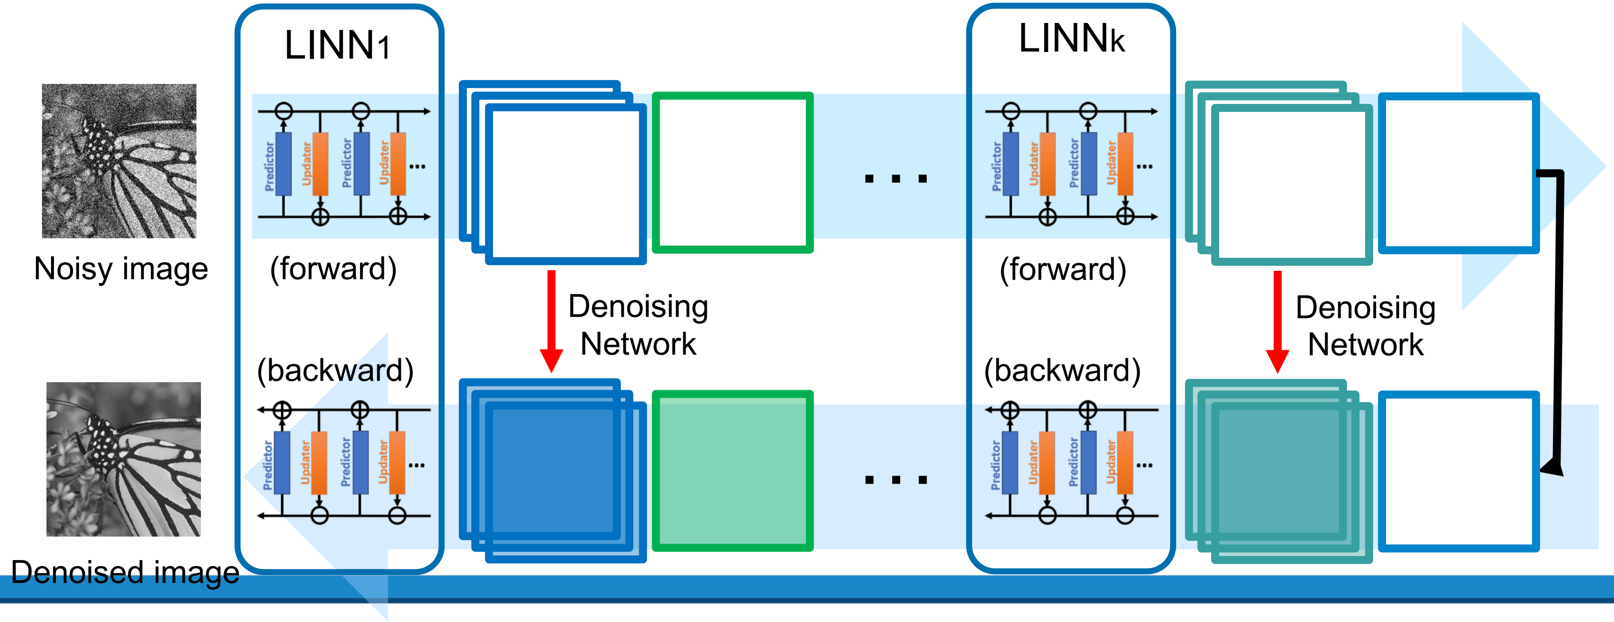

Imperial College London

Results

Denoising:

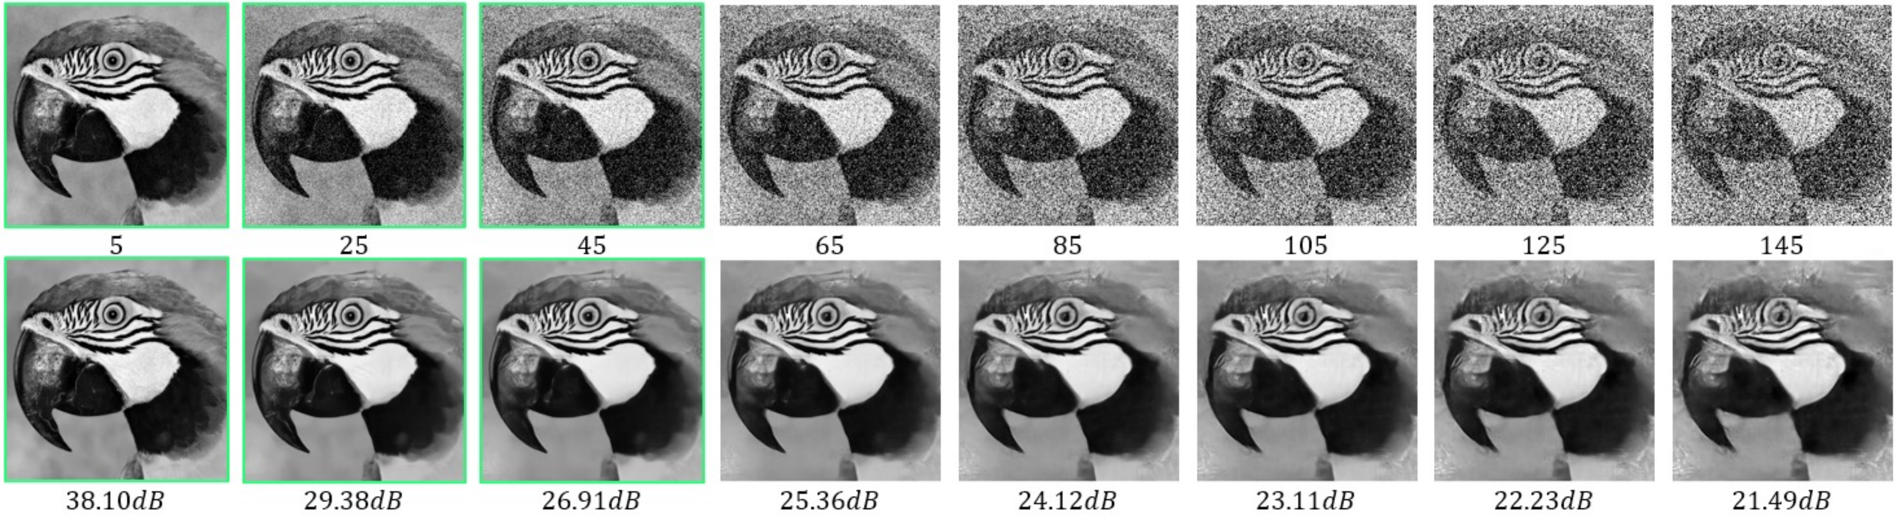

J. Huang and P.L. Dragotti, “WINNet: Wavelet-inspired Invertible Network for Image Denoising”, IEEE Trans. on Image Proc., 2022

Image Deblurring

Imperial College London

Results

Deconvolution:

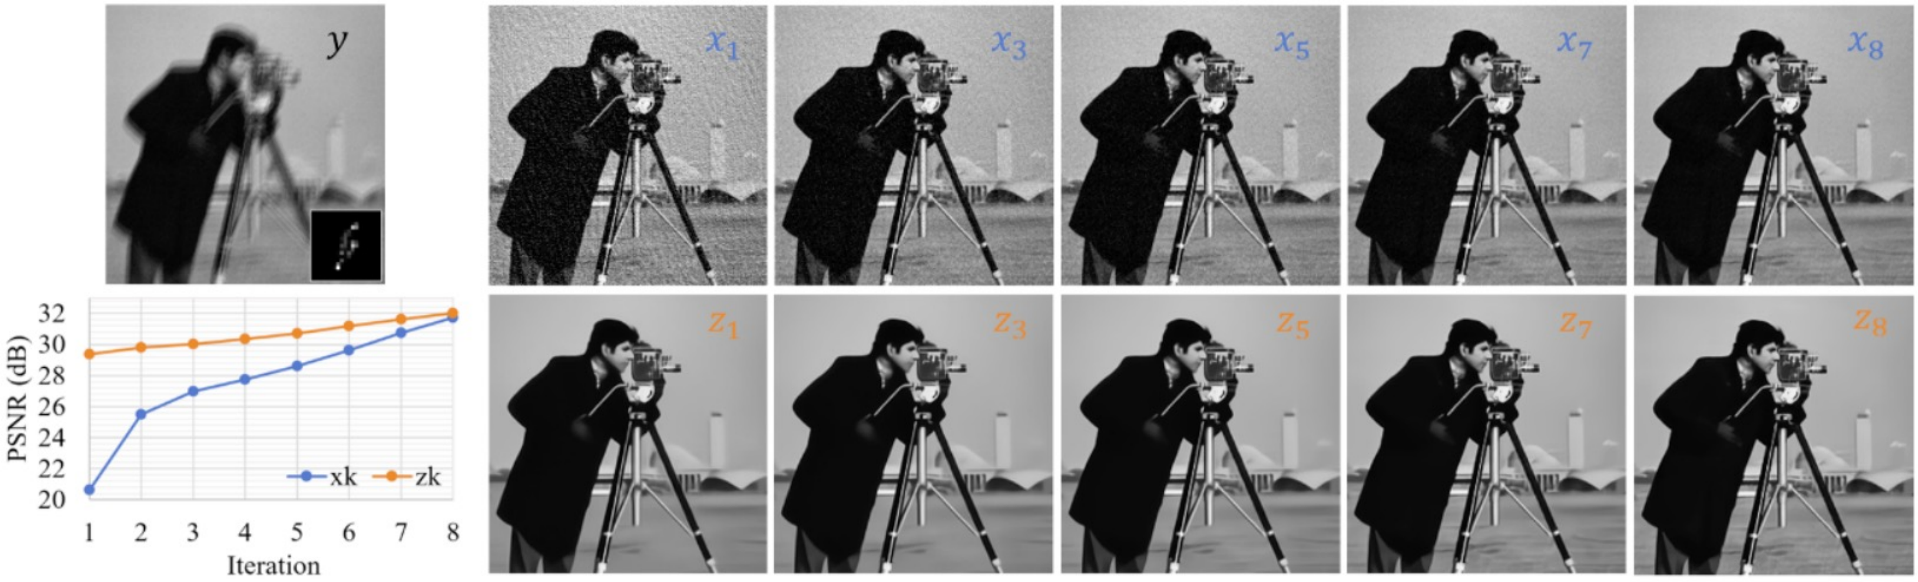

J. Huang and P.L. Dragotti, "WINNet: Wavelet-inspired Invertible Network for Image Denoising", IEEE Trans. on Image Proc., 2022

Imperial College London

INN + Diffusion Models for Inverse Problems

- $\hat{x} = \min_{x} ||H(x) - y||^2 + \lambda\rho(x)$

  consistency term prior

- Impose consistency using the forward part of the INN
- Impose the prior using diffusion models
- Iterate

INN + Diffusion Models for Inverse Problems

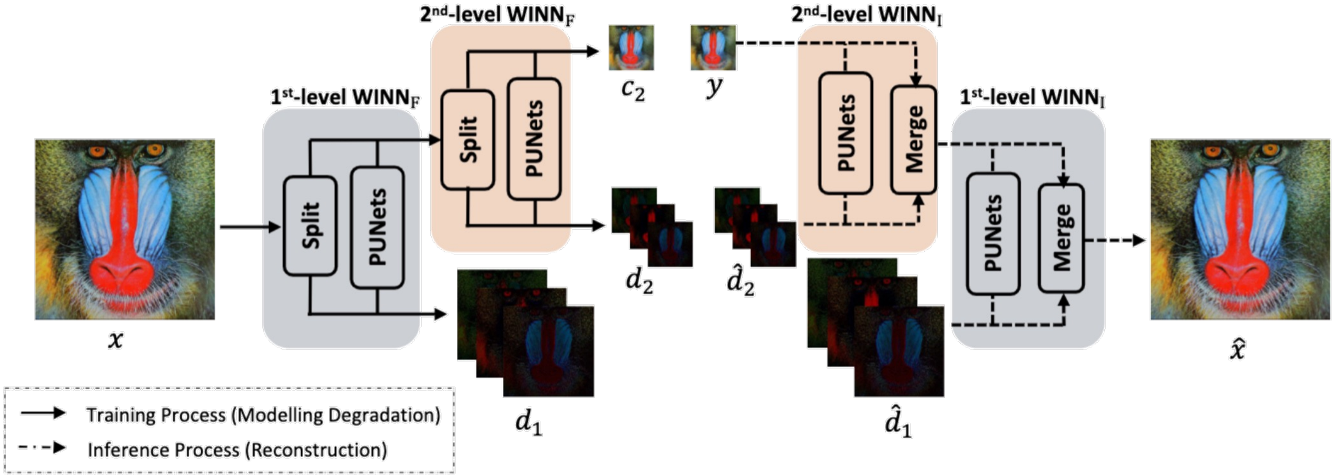

Imperial College London

Ground Truth

Degraded

Reconstructed

Ground Truth

Degraded

Reconstructed

Imperial College London

First Set of Conclusions

- Invertible Neural Networks are an interesting new concept
- Designing INN using wavelets/lifting leads to more interpretable and simpler architectures
- Good generalization ability
- Potential for further developments by combining INNs with diffusion models

Sparsity and Deep Unfolding Strategy

Imperial College London

Explicit embedding of priors and constraints in deep networks

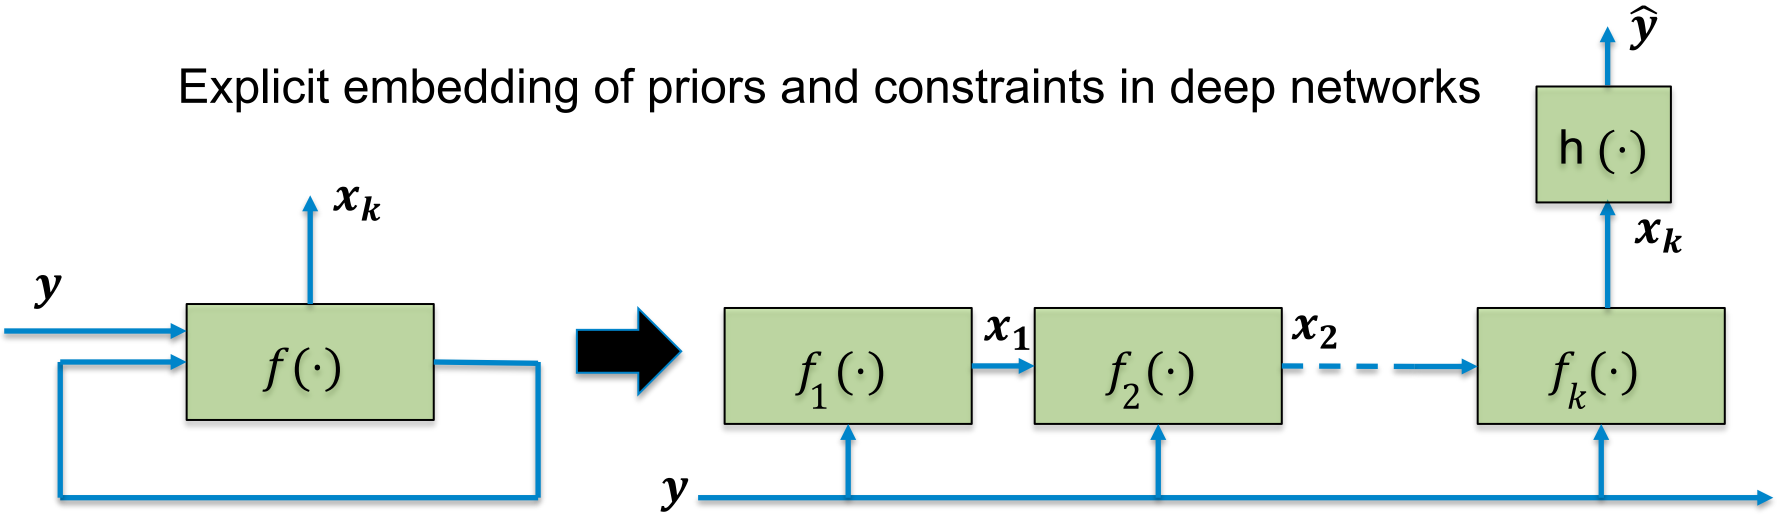

Iterative algorithm with $\boldsymbol{x}$ as input and $\boldsymbol{I}$ as output

Unfolded version of the iterative algorithm with learnable parameters

Need to re-synthesize the input, if self-supervised

Imperial College London

Sparsity as the model for deep unfolding

- The dictionary is usually learned

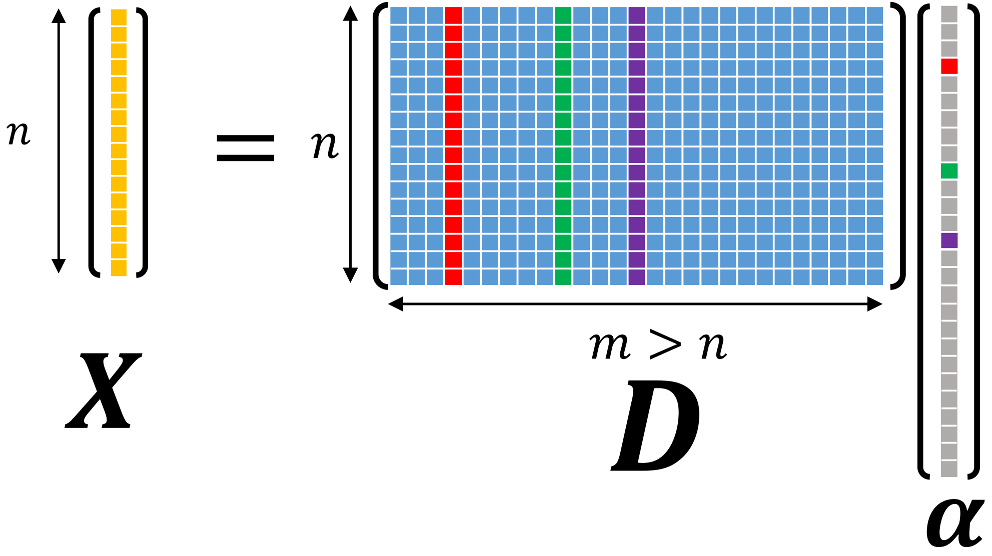

Deep Unfolding Strategy

- The sparse vector α can be found using ISTA: $\alpha_k = S_{\lambda_k}(\alpha_{k-1} + D_x^T(x - D_x\alpha_{k-1}))$

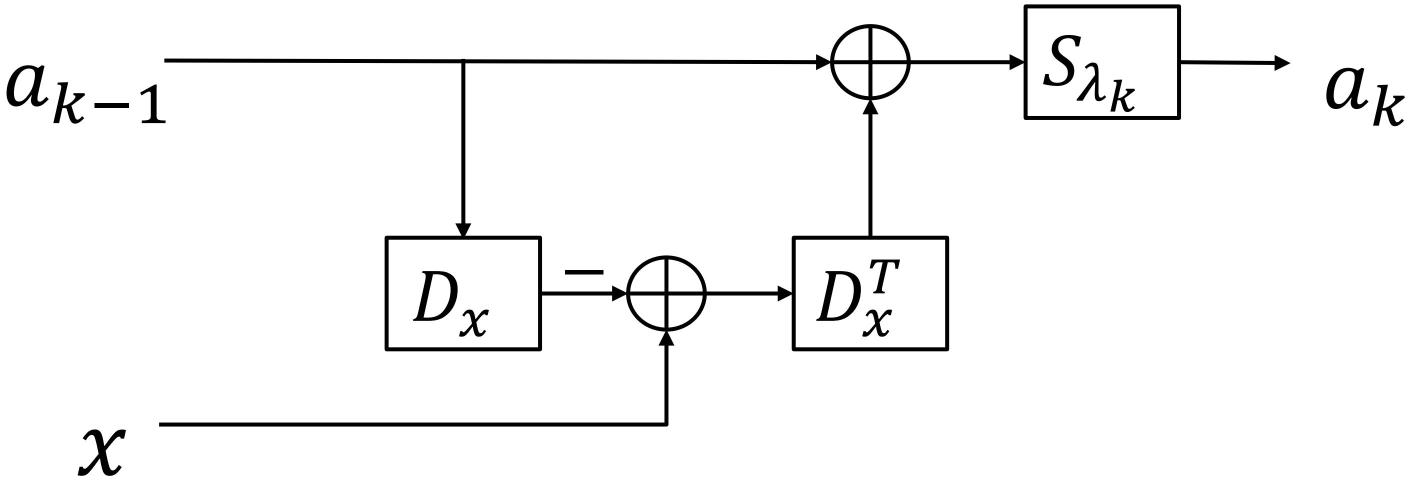

Imperial College London

ISTA network

- Solving by ISTA algorithm through unfolding:

$$\boldsymbol{a}_k = S_{\lambda_k}(\boldsymbol{a}_{k-1} + \boldsymbol{D}_x^T(\boldsymbol{x} - \boldsymbol{D}_x\boldsymbol{a}_{k-1}))$$

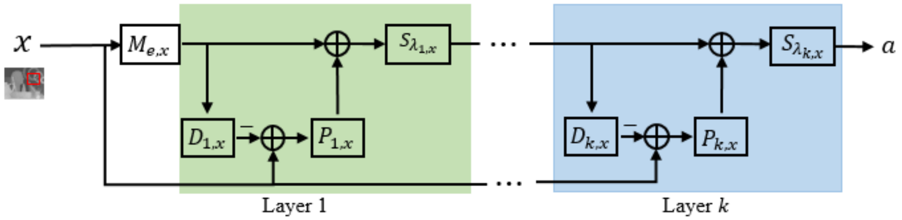

- Gregor Karol and LeCunYann, “Learning fast approximations of sparse coding ”, Proceedings of the 27th International Conference on International Conference on Machine Learning, 2010
- Y. Eldar et al, “Algorithm Unrolling: Interpretable, Efficient Deep Learning for Signal and Image Processing”, IEEE Signal Processing Magazine, 2021

Imperial College London

Art-Investigation

- Goal: we want to separate the two x-ray images
- Approach:
  - Use the visible RGB image as side information (x-ray visible similar to RGB image)
  - Exclusion loss: the “contours” of the two x-ray images should be as different as possible

Visible

X-ray

Francisco de Goya, Dona Isabel de Porcel (NG1473), before 1805. Oil on canvas, Images © The National Gallery

Imperial College London

X-ray Separation – Proposed Sparsity Model

$$\boldsymbol{x}_1 = \sum_{k=1}^{K} \boldsymbol{\Xi}_k * \boldsymbol{z}_{1,k}, \quad \boldsymbol{x}_2 = \sum_{k=1}^{K} \boldsymbol{\Xi}_k * \boldsymbol{z}_{2,k},$$

$$\boldsymbol{r}_{1,s} = \sum_{k=1}^{K} \boldsymbol{\Omega}_{k,s} * \boldsymbol{z}_{1,k}, \quad \boldsymbol{x} = \sum_{k=1}^{K} \boldsymbol{\Xi}_k * (\boldsymbol{z}_{1,k} + \boldsymbol{z}_{2,k}),$$

- The visible image and the two separated X-ray images have a sparse representation in proper dictionaries
- RGB image and visible X-ray share the same sparse representation
- The two X-rays x_1, x_2 share the same dictionary
- The measured X-ray is $x = x_1 + x_2$

Visible

X-ray

Francisco de Goya, Dona Isabel de Porcel (NG1473), before 1805. Oil on canvas, Images © The National Gallery

Imperial College London

X-ray Separation – Exclusion Loss

- Given the reconstructed X-ray images x_1, x_2, we expect that their edges are as different as possible we therefore add an “exclusion term” in the optimization

$$
\begin{aligned}
\min_{\boldsymbol{y}_1,\boldsymbol{y}_2,\boldsymbol{z}_{1,k},\boldsymbol{z}_{2,k}} \quad & \|\boldsymbol{x} - \boldsymbol{\Psi} * \boldsymbol{y}_1 - \boldsymbol{\Psi} * \boldsymbol{y}_2\|_F^2 \\
& + \tau_1 \|\boldsymbol{y}_1 - \sum_{k=1}^{K} \boldsymbol{\Theta}_k * \boldsymbol{z}_{1,k}\|_F^2 \\
& + \tau_2 \|\boldsymbol{y}_2 - \sum_{k=1}^{K} \boldsymbol{\Theta}_k * \boldsymbol{z}_{2,k}\|_F^2 \\
& + \gamma \sum_{s=1}^{3} \|\boldsymbol{r}_{1,s} - \boldsymbol{\Phi}_s * \boldsymbol{y}_1\|_F^2 \\
& + \lambda_1 \sum_{k=1}^{K} \|\boldsymbol{z}_{1,k}\|_1 + \lambda_2 \sum_{k=1}^{K} \|\boldsymbol{z}_{2,k}\|_1 \\
& + \sum_{i=1}^{I} \mu_i \|(\boldsymbol{W}_i * \boldsymbol{y}_1) \odot (\boldsymbol{W}_i * \boldsymbol{y}_2)\|_1,
\end{aligned}
$$


Visible


X-ray

Imperial College London

One Layer of the Network

- The sparsity model and the exclusion constraint leads to an iterative optimization method which leads to a network through unfolding

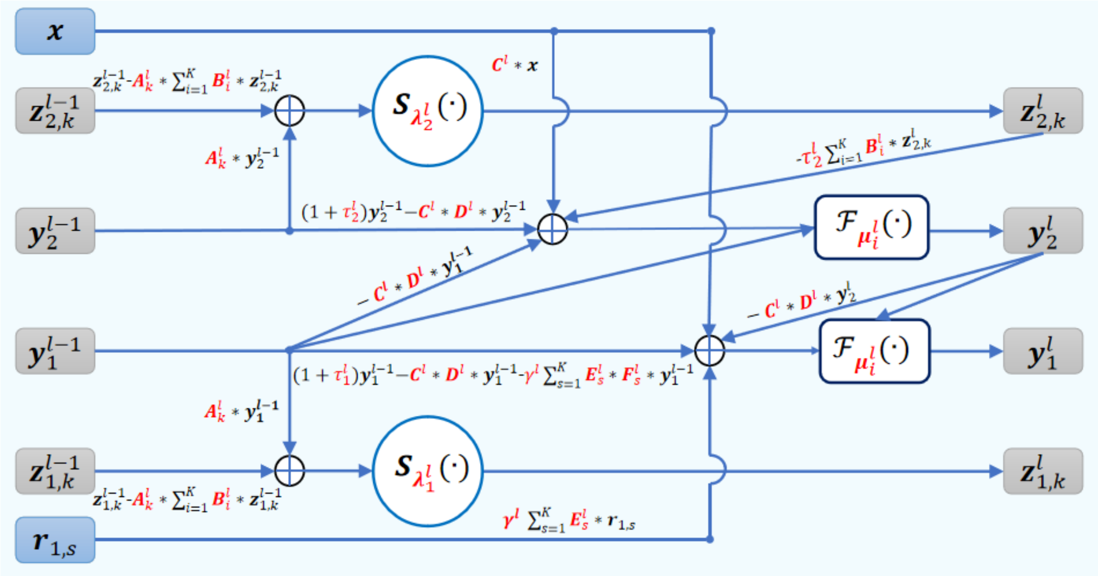

Separation Results

W. Pu et al “Mixed x-ray image separation for artworks with concealed designs”, IEEE Trans. on Image Processing, 2022

Imperial College London

MA-XRF Datacube and Spectrum

- Macro X-ray provides volumetric data and the locations of the pulses in the energy direction are related to the chemical elements present in the painting.
- This potentially allows us to create maps that show the distribution of different chemical elements

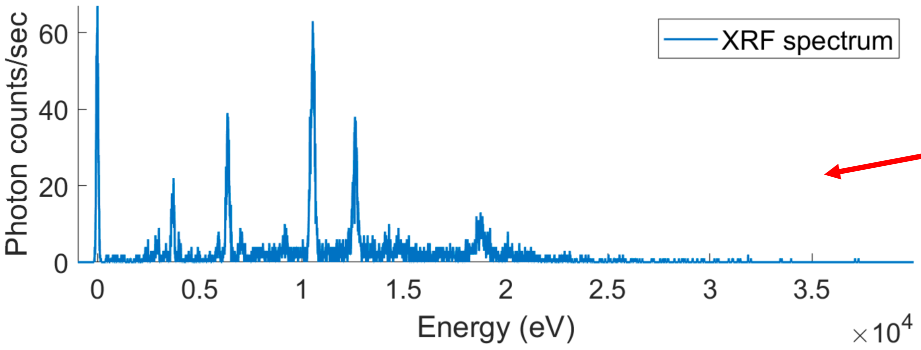

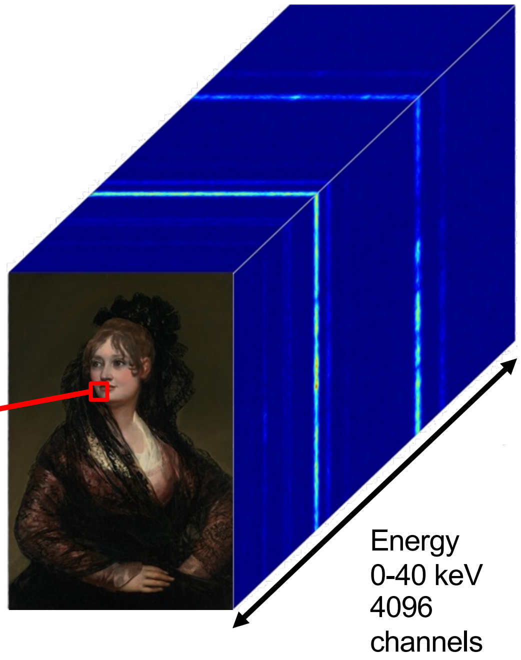

Imperial College London

Extraction of Elemental Maps

Our XRF
Deconvolution
Algorithm

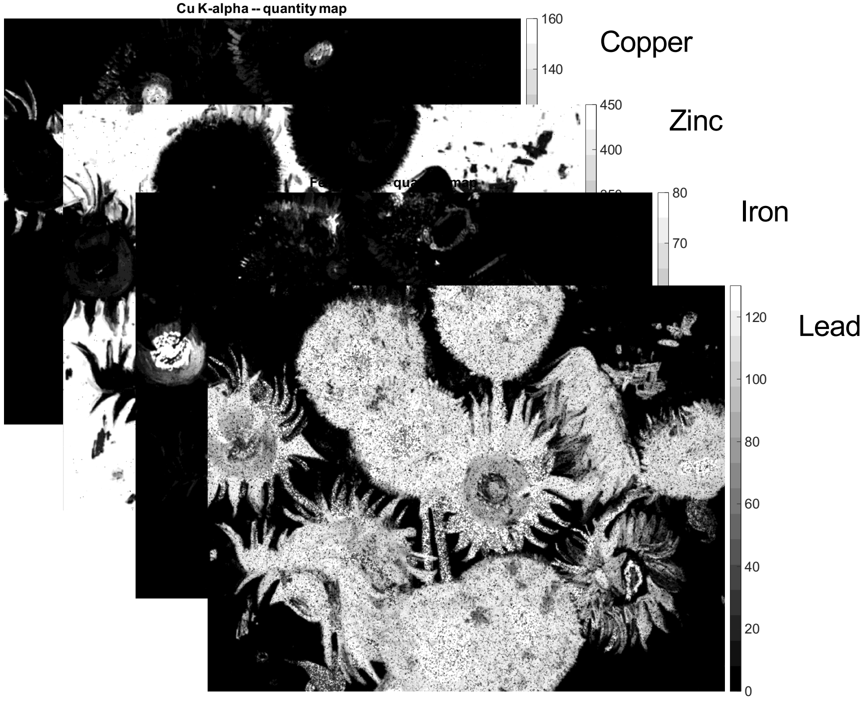

Imperial College London

Results

Leonardo da Vinci's "The Virgin of the Rocks"

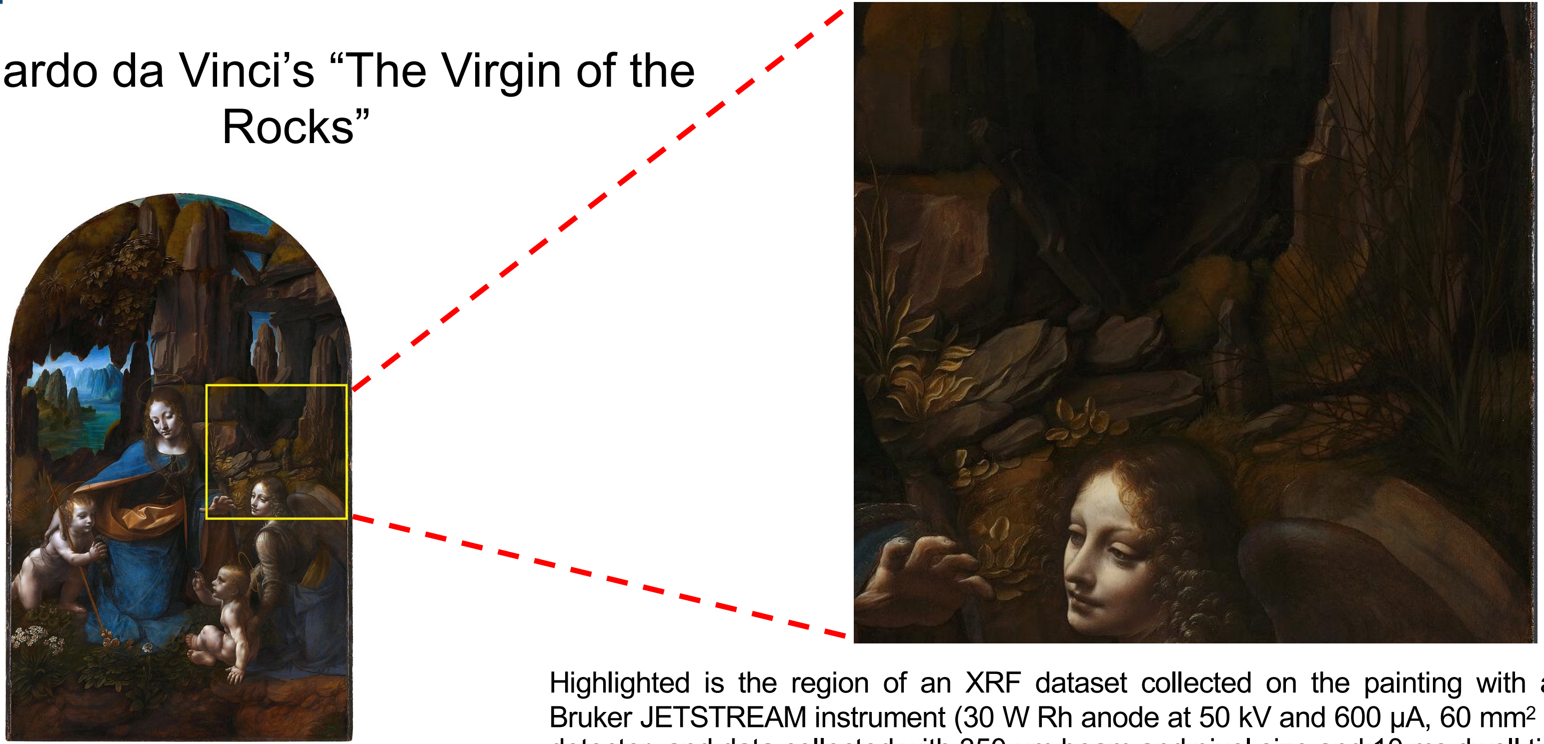

Highlighted is the region of an XRF dataset collected on the painting with an M6 Bruker JETSTREAM instrument (30 W Rh anode at 50 kV and 600 µA, 60 mm^2 Si drift detector, and data collected with 350 µm beam and pixel size and 10 ms dwell time).

Leonardo da Vinci, "The Virgin of the Rocks (NG1093)," about 1491/2-9 and 1506-8, oil on poplar, 189.5 x 120 cm, The National Gallery, London.

Imperial College London

Zinc (Zn) distribution maps

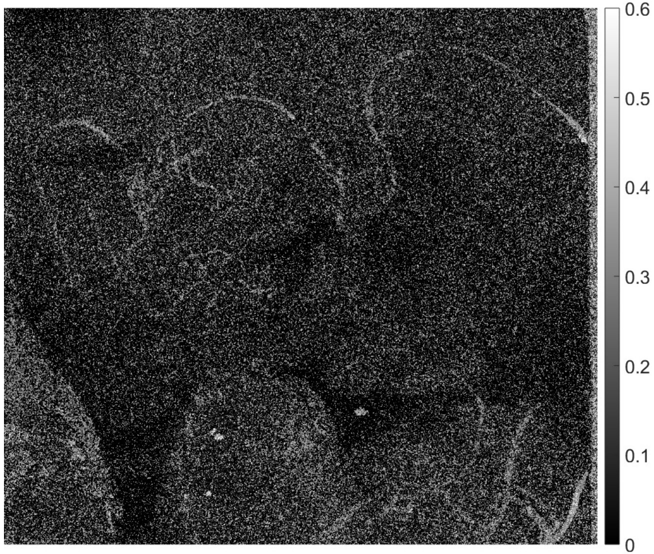

Zn confidence map

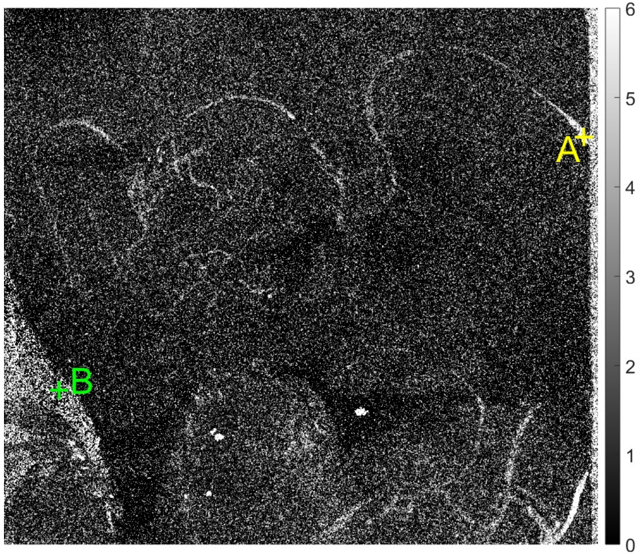

Zn quantity map

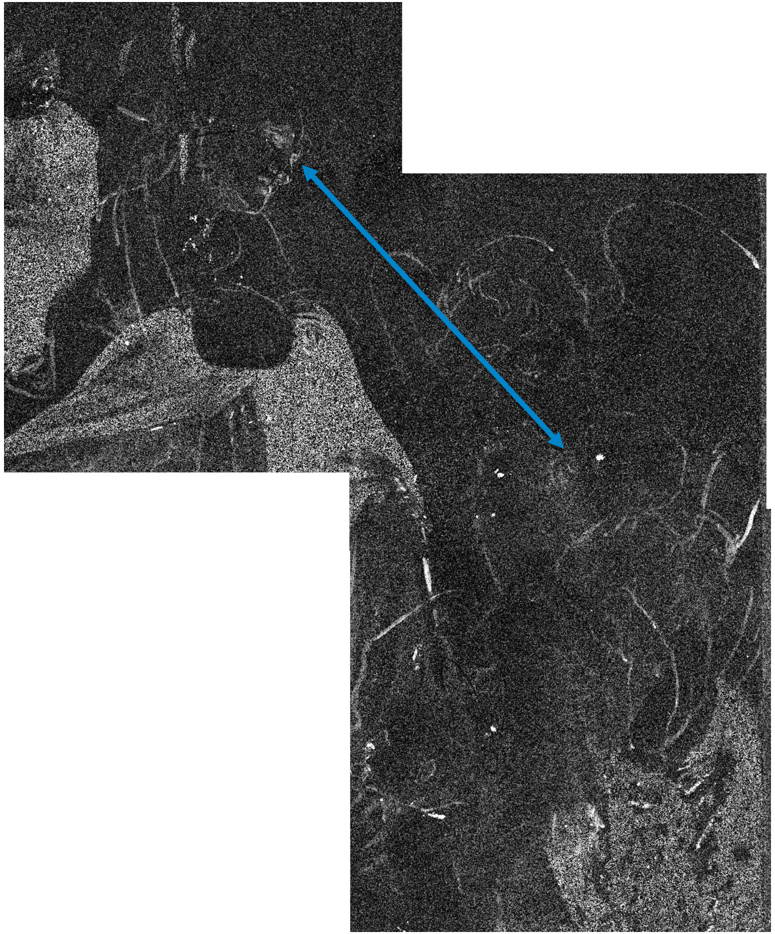

S. Yan, J.-J. Huang, N. Daly, C. Higgitt, and P. L. Dragotti, “When de Prony Met Leonardo: An Automatic Algorithm for Chemical Element Extraction in Macro X-ray Fluorescence Data”, IEEE Transactions on Computational Imaging, vol.7, 2021.

Imperial College London

Two-Photon Microscopy for Neuroscience

- Goal of Neuroscience: to study how information is processed in the brain
- Neurons communicate through pulses called Action Potentials (AP)
- Need to measure in-vivo the activity of large populations of neurons at cellular level resolution
- Two-photon microscopy combined with right indicators is the most promising technology to achieve that

Imperial College London

Two-Photon Microscopy

- Fluorescent sensors within tissues
- Highly localized laser excites fluorescence from sensors
- Photons emitted from tissue are collected
- Focal spot sequentially scanned across samples to form image
- Two-photon microscopes in raster scan modality can go deep in the tissue but are slow

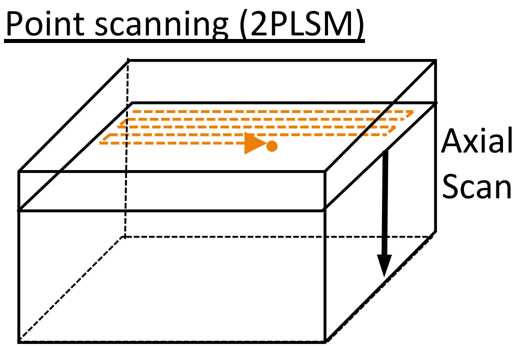

Imperial College London

Two-Photon Microscopy

- In order to speed up acquisition one can change the illumination strategy
- This mitigates the issue but does not fix it
- Issue with scattering

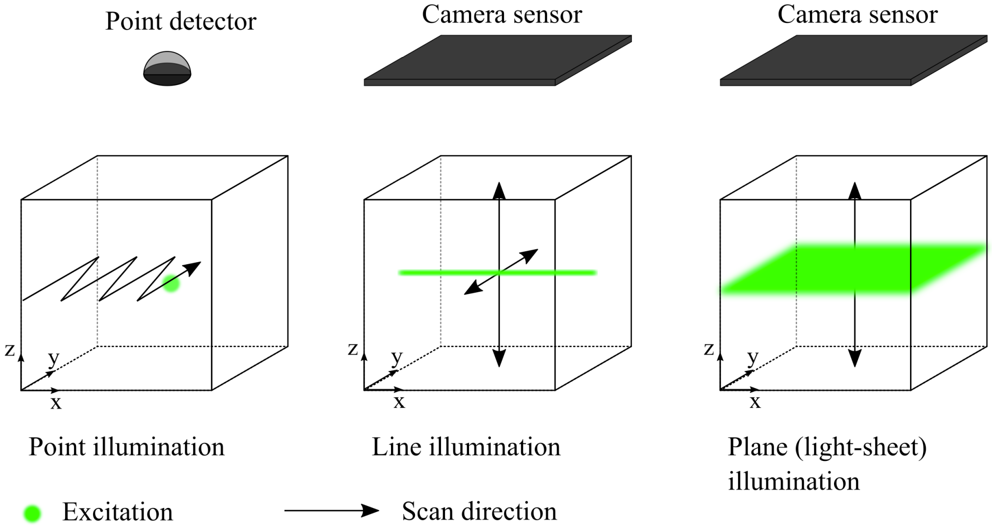

Imperial College London

Light-field Microscopy

Light-Field Microscopy (LFM) is a high-speed imaging technique that uses a simple modification of a standard microscope to capture a 3D image of an entire volume in a single camera snapshot

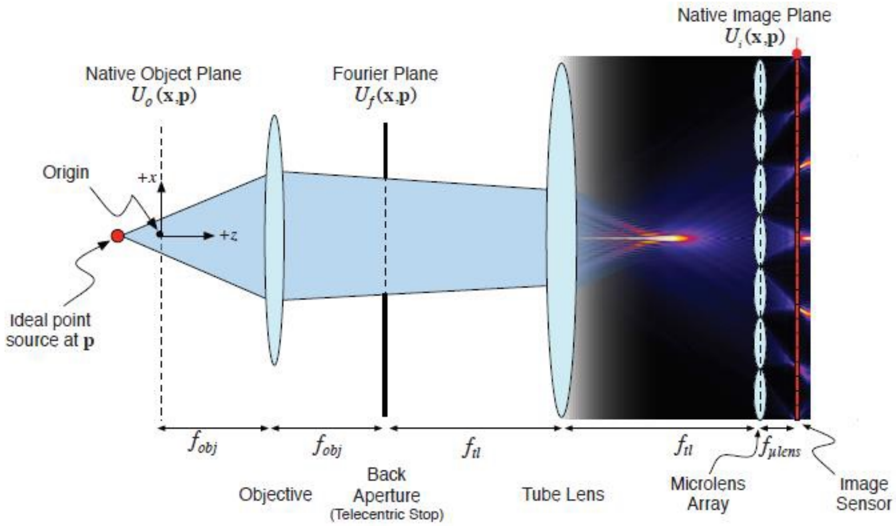

Imperial College London

Light-field Microscopy and EPI

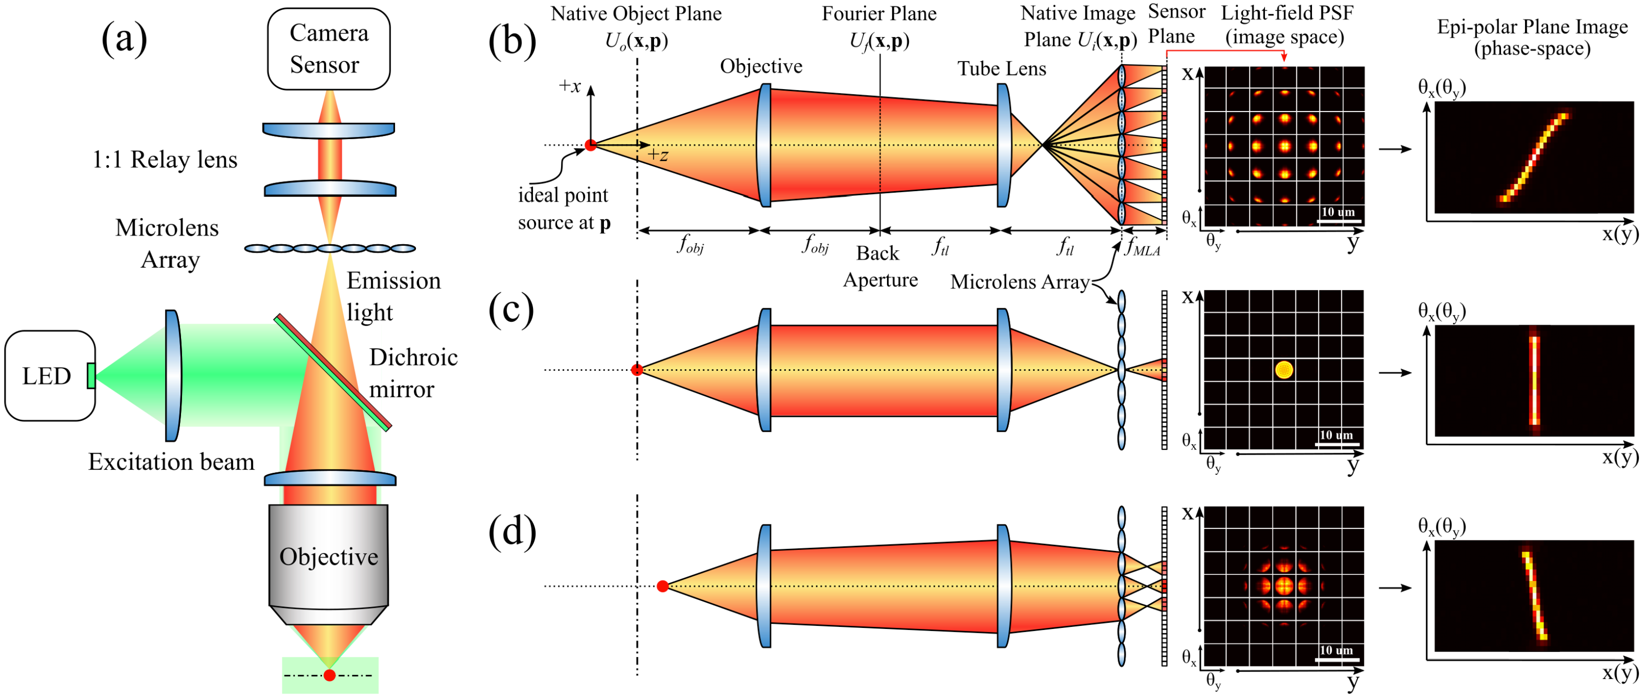

Imperial College London

Light-field Microscopy and Illumination Strategies

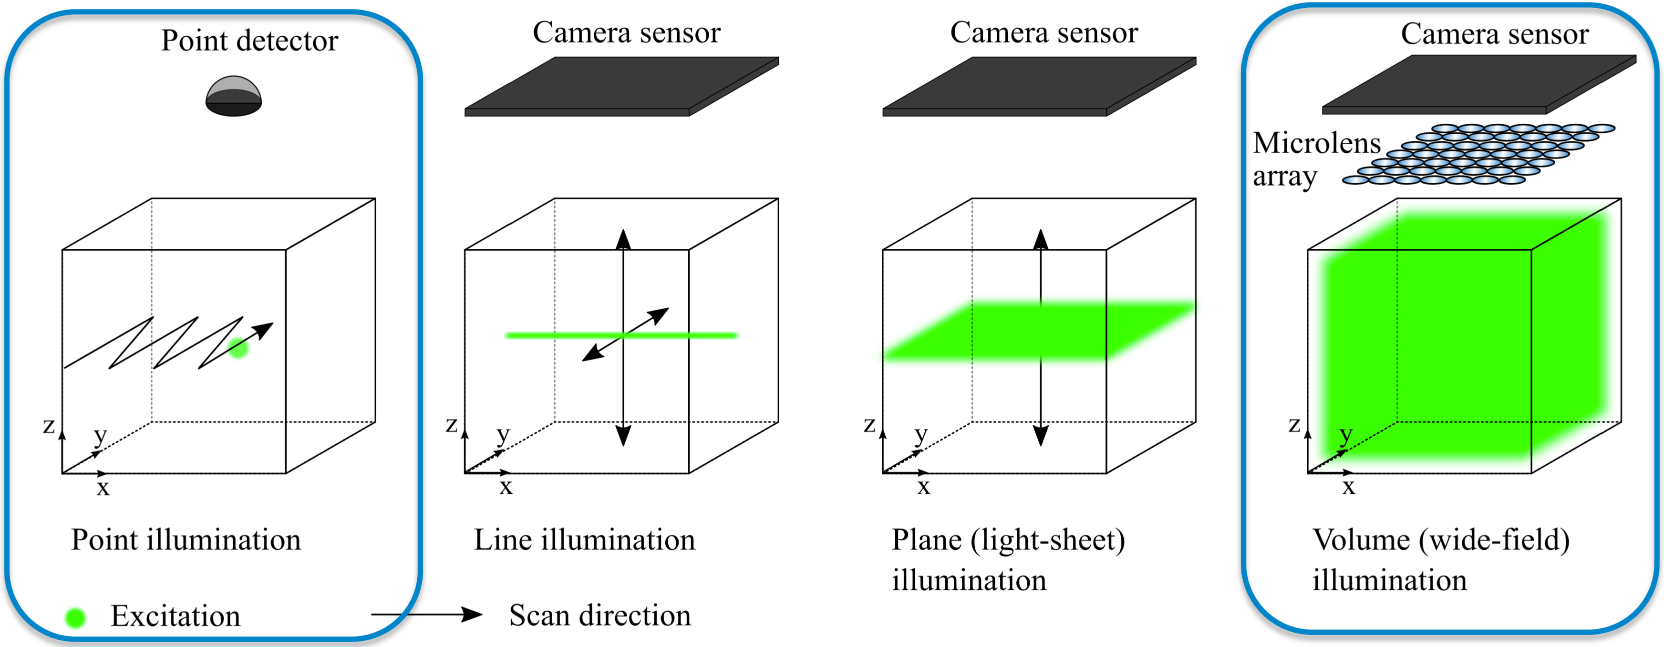

Key insight: use the 2P microscope for high-resolution structural information and the LFM for monitoring the activity of neurons.

Imperial College London

Light-field Microscopy

Challenge: given a sequence of lightfields (2-D signals), need to reconstruct a sequence of volumes (3-D+t)

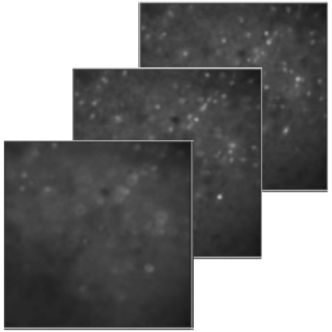

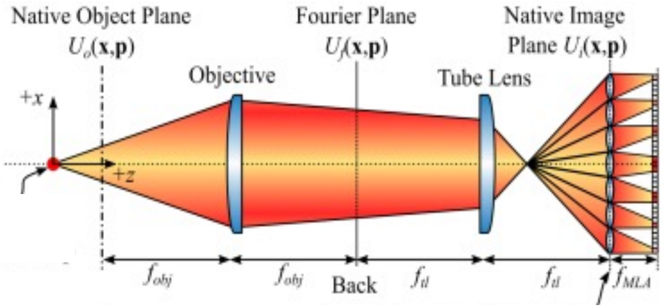

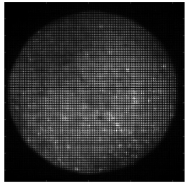

2D Measured LF image

3D Input

Computational Algorithm

Imperial College London

Volume reconstruction from LF Data

- **Challenges**
  - Scattering induces blur, making inversion more challenging
  - Lack of ground-truth data for learning

- **Opportunities**
  - Forward model structured and linear
  - Data is sparse (neurons fire rarely and are localized in space)
  - Occlusion can be ignored

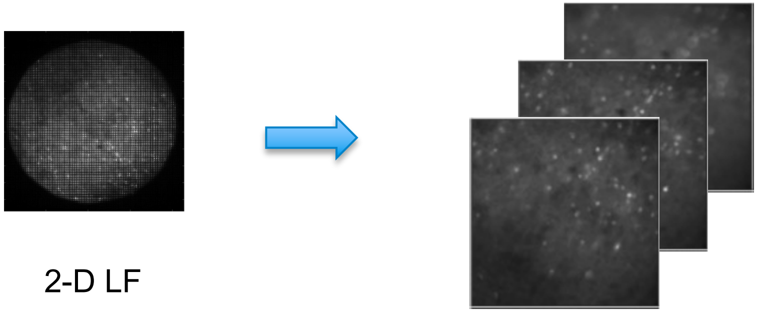

2-D LF

Volume

Forward Model

- Forward model is linear which means $\boldsymbol{y} = \boldsymbol{H}\boldsymbol{x}$
  - $\boldsymbol{H}$ is estimated using wave-optics
  - For each depth, $\boldsymbol{H}$ is block-circulant (periodically shift invariant) and can be modelled with a filter-bank
  - The entire forward model can be modelled using a linear convolutional network with known parameters (given by the wave-optics model)

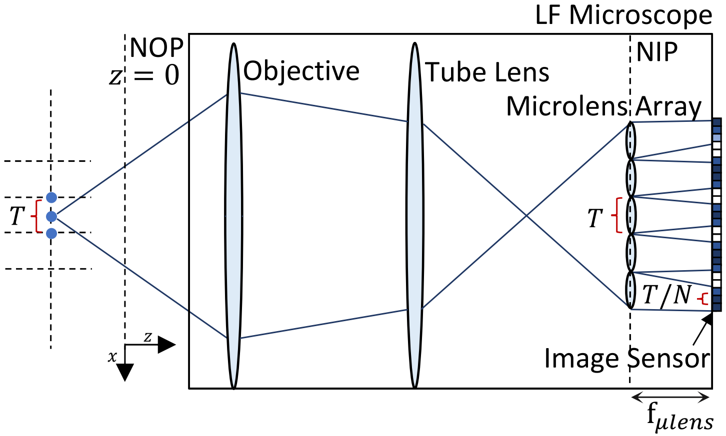

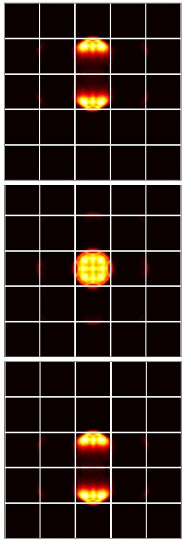

Imperial College London

Neural network for volume reconstruction

- Data is sparse (neurons fire rarely and are localized in space)
- Solve $\min_{x}(\|y - Hx\|^2 + \|x\|_1)$ s.t $x \geq 0$
- This leads to the following iteration:

$$x_{k+1} = ReLU(x_k - H^T H x_k + H^T y + \lambda)$$

- Approach: Convert the iteration in a deep neural network using the unfolding technique

Imperial College London

Neural network for volume reconstruction

- Convert the iteration in a deep neural network using the unfolding technique

$$x^{k+1} = ReLU(x^k - H^T H x^k + H^T y + \lambda)$$

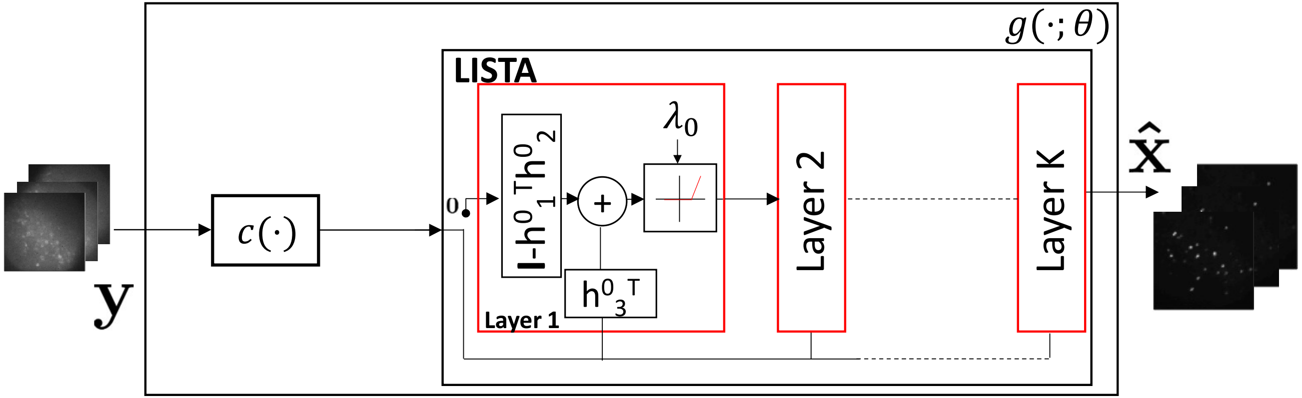

Training of the neural network

- Training, in this context, is difficult due to lack of ground-truth data
- Our approach: semi supervised learning
  - Small ground truth dataset
  - Adversarial network for adversarial loss
  - Light-field loss based on re-synthesizing light-field from reconstructed volume

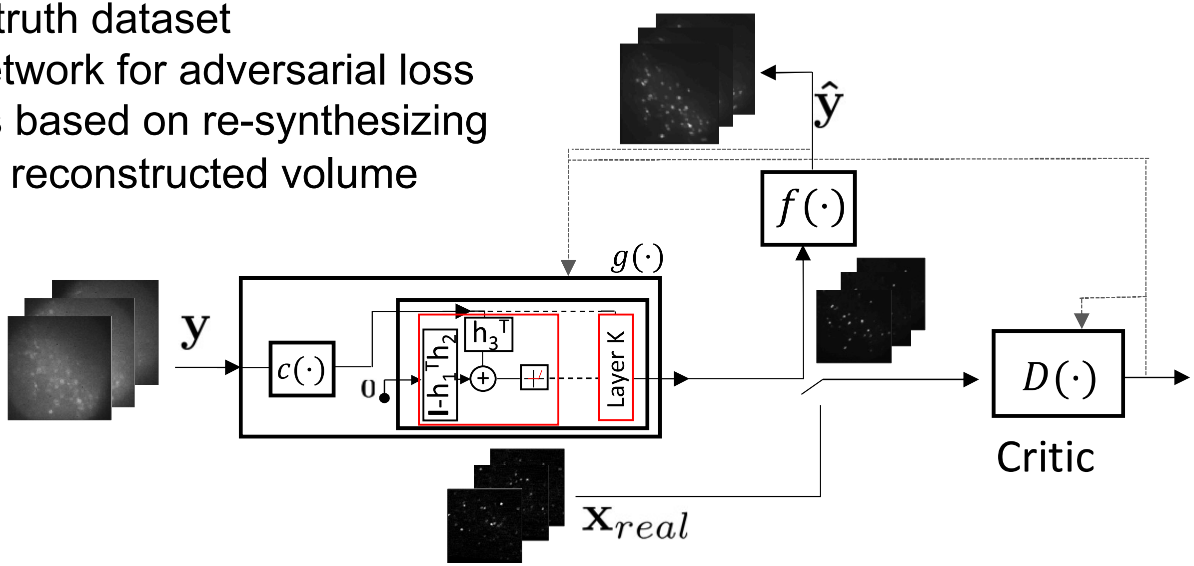

Training of the neural network

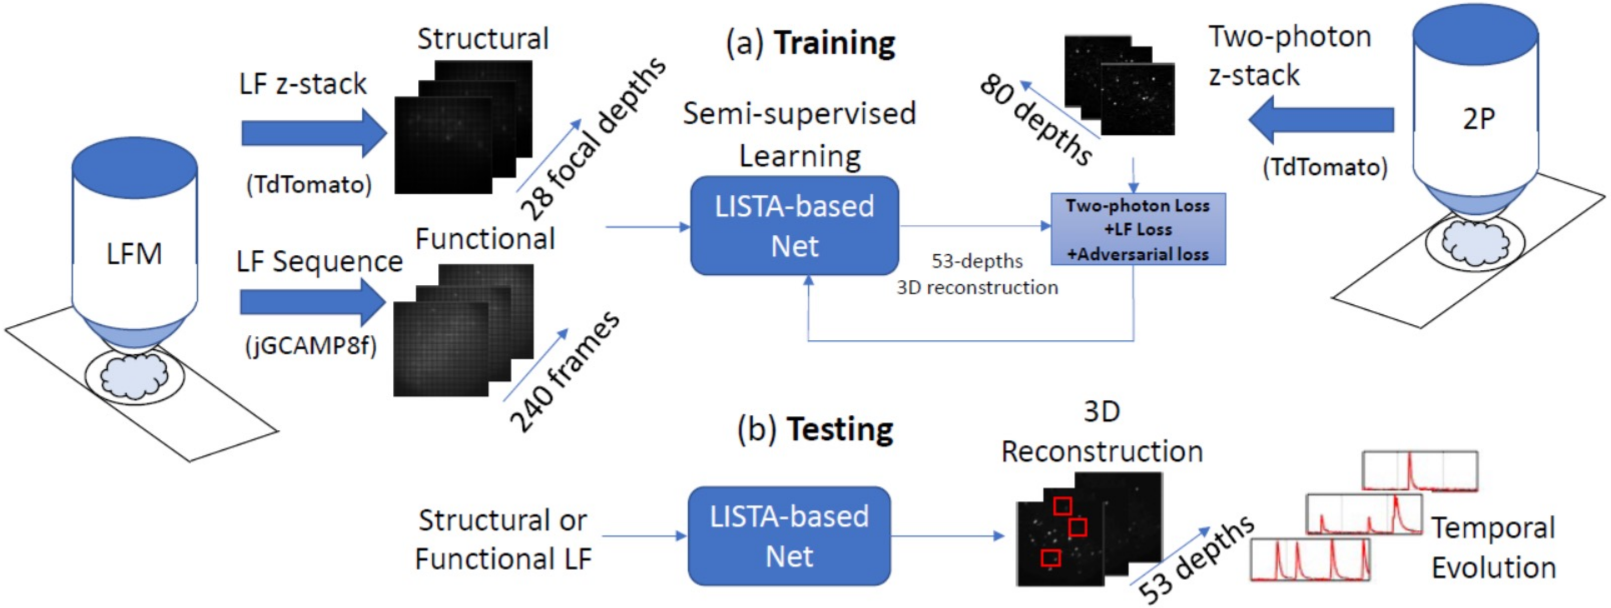

Imperial College London

Results – Structural Data

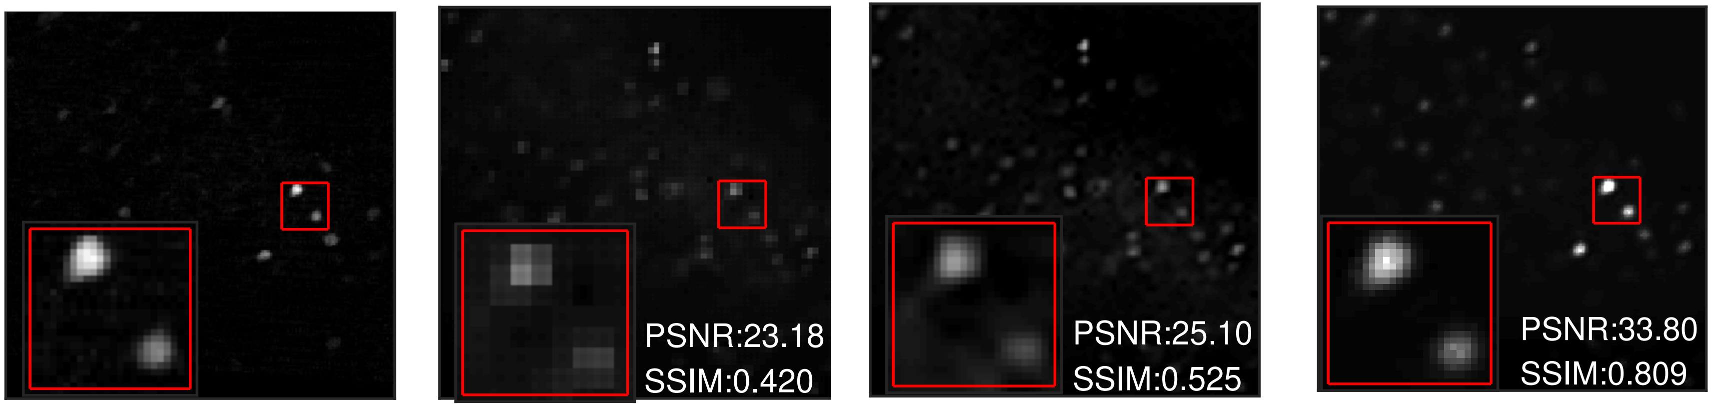

Ground-truth | ISRA | ADMM | New method (0.3s to reconstruct one volume)

H. Verinaz et al. "Physics-based Deep Learning for Imaging Neuronal Activity via Two-photon and Light-field Microscopy", IEEE Trans. on Computational Imaging, 2023.

Imperial College London

Results – Functional Data

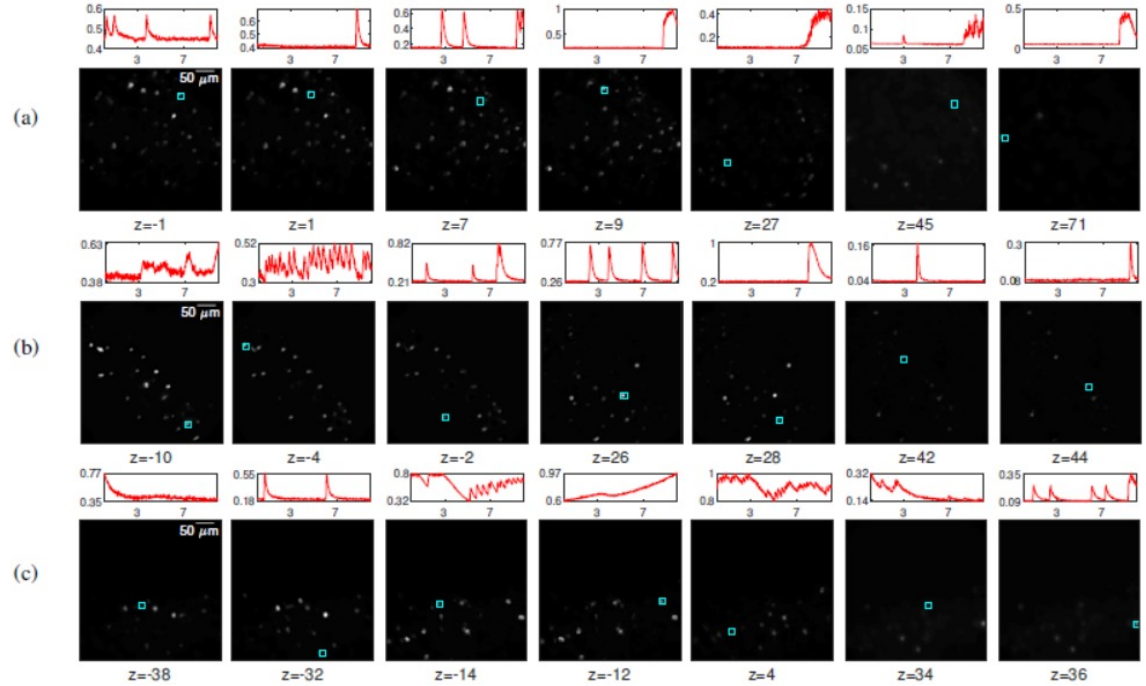

Three brain samples are shown in parts (a), (b), and (c)

Imperial College London

Conclusions

- Cross fertilization between model-based approaches and deep learning is fruitful
  - Models and priors can reduce complexity of a deep network and can lead to better results
  - Some computational approaches are transferable

- Computational Imaging:
  - is fun 🙂,
  - is inter-disciplinary,
  - is the right way to handle ‘big data’: joint sensing, representation, analysis and inference

Imperial College London

A special thank to:

Junjie Huang

Pingfan Song

Herman Verinaz

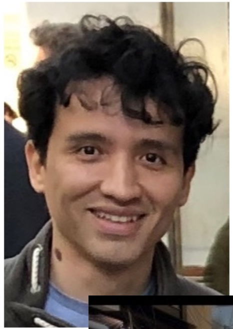

Peter Quicke

Carmel Howe

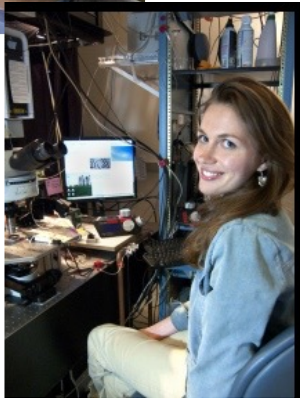

Amanda Foust

Consortium involving: UCL, ICL, Duke and National Gallery

Imperial College London

Thank you!

Imperial College London

Related Publications

- J. Huang and P.L. Dragotti, “LINN: Lifting Inspired Invertible Neural Network for Image Denoising”, in proc. of 29th European Signal Processing Conference, EUSIPCO 2021
- J. Huang and P.L. Dragotti, “WINNet: Wavelet-inspired Invertible Network for Image Denoising”, IEEE Transactions on Image Processing, 2022
- W. Pu, J. Huang et al., “Mixed X-Ray Image Separation for Artworks with Concealed Designs”, IEEE Transactions on Image Processing, 2022
- S. Yan, J.-J. Huang, N. Daly, C. Higgitt, and P. L. Dragotti, “When de Prony Met Leonardo: An Automatic Algorithm for Chemical Element Extraction in Macro X-ray Fluorescence Data”, IEEE Transactions on Computational Imaging, vol.7, 2021.
- H. Verinaz et al. "Physics-based Deep Learning for Imaging Neuronal Activity via Two-photon and Light-field Microscopy”, IEEE Trans. on Computational Imaging, 2023.
- D. You, A. Floros and P.L. Dragotti, “INDigo: An INN-Guided Probabilistic Diffusion Algorithm for Inverse Problems”, IEEE MMSP conference, September 2023, arxiv:2306.02949