Imperial College London

Inverse Problems in the Age of AI

Pier Luigi Dragotti, Imperial College London

20 November 2024

Imperial College London

Motivation: Inverse Problem in Imaging

Inverse problems involve reconstructing unknown physical quantities from indirect measurements.

The growing complexity of modern imaging workflows calls for a more holistic approach to inverse problems where sensing, physics and computation are analized jointly

Key in inverse problem is the development of the interplay between physical and learned models

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

Model-Based Deep Learning

Plato: models, priors

Aristotle: data

Need to find the right balance between **data** and **prior** models to develop methods that

- reduce complexity,
- increase generalizability
- can handle lack of training data
- can handle complex settings

Imperial College London

Three Case Studies in Imaging Science

Image restoration problems: Invertible neural networks and diffusion models

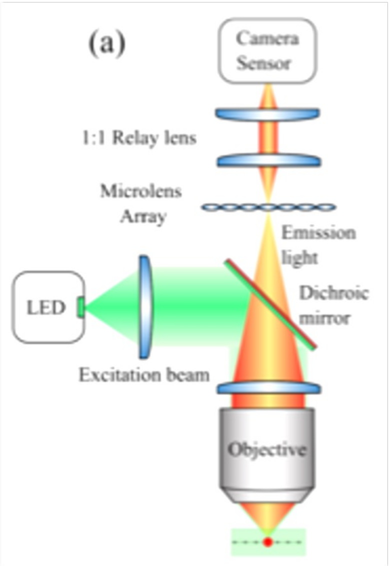

Light field microscopy for neuroscience

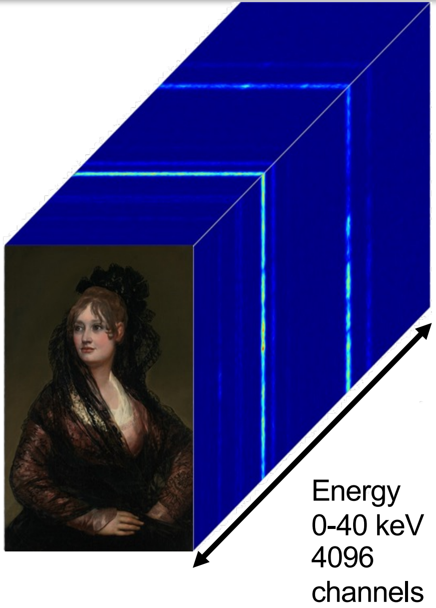

Technical study of Old Masters paintings

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

Imperial College London

Plug and Play

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

  consistency term   prior

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

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

- Solve by alternating between x and v

  - Consistency step: $\hat{x} = \min_{x} ||H(x) - y||^2 + \beta||x - v||^2$

  - 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

INN and Diffusion Models for Inverse Problems

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

  consistency term prior

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

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

- Solve by alternating between x and v

  - Consistency step: $\hat{x} = \min_{x} \|H(x) - y\|^2 + \beta\|x - v\|^2$

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

Use INN to impose consistency

Use Diffusion Models to impose the prior

Invertible Neural Networks

Invertible Neural Networks are bijective function approximators with a forward mapping

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

and inverse mapping

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

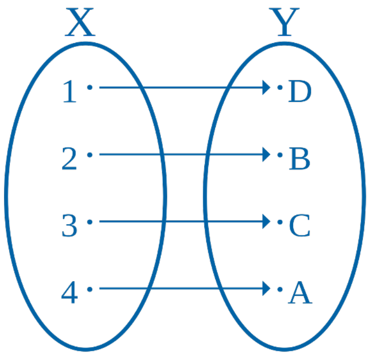

A bijective function (or invertible function)

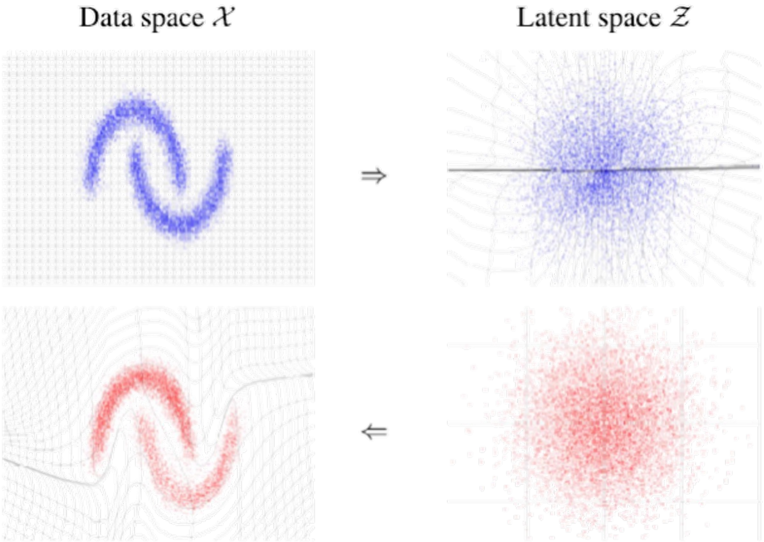

Invertible Neural Networks

How to Achieve Invertibility?

Invertible via lifting scheme like architectures

- Signal splitting
- Alternate prediction and update

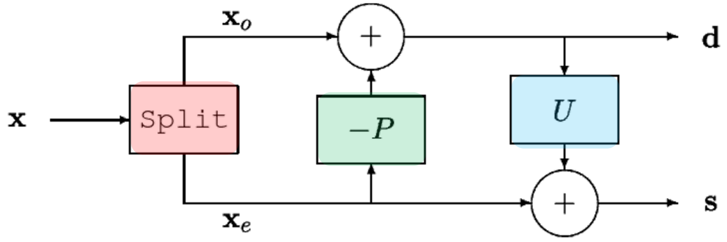

Split → $$\begin{cases} d = x_o - P(x_e) \\ s = x_e + U(d) \end{cases}$$

Forward pass

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

Backward pass

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

INN and Diffusion Models for Inverse Problems

Invertible Neural Networks are ideal architectures to address inverse problems

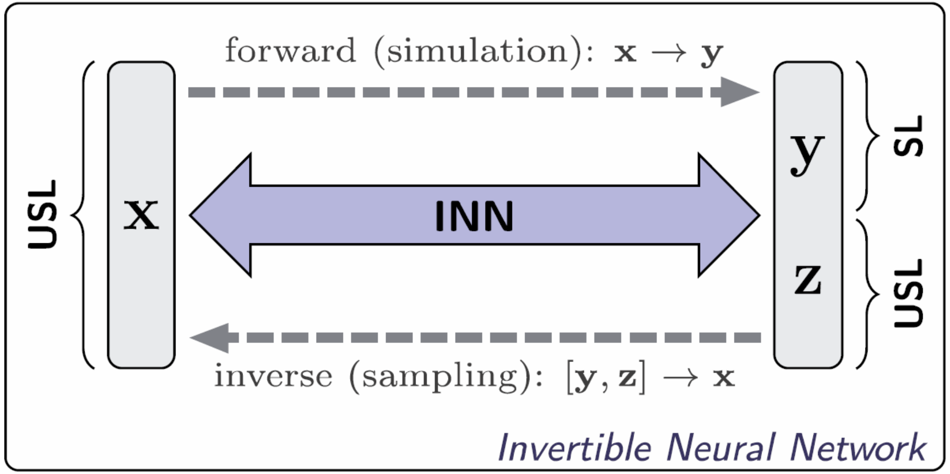

Figure from: Ardizzone, Lynton, Jakob Kruse, Sebastian Wirkert, Daniel Rahner, Eric W. Pellegrini, Ralf S. Klessen, Lena Maier-Hein, Carsten Rother, and Ullrich Köthe. "Analyzing inverse problems with invertible neural networks." in Proc. of *ICLR*, 2019.

Imperial College London

INN and Diffusion Models for Inverse Problems

Diffusion Models are good for “unconditional” generation of new samples (e.g., Denoising Probabilistic Diffusion Models)

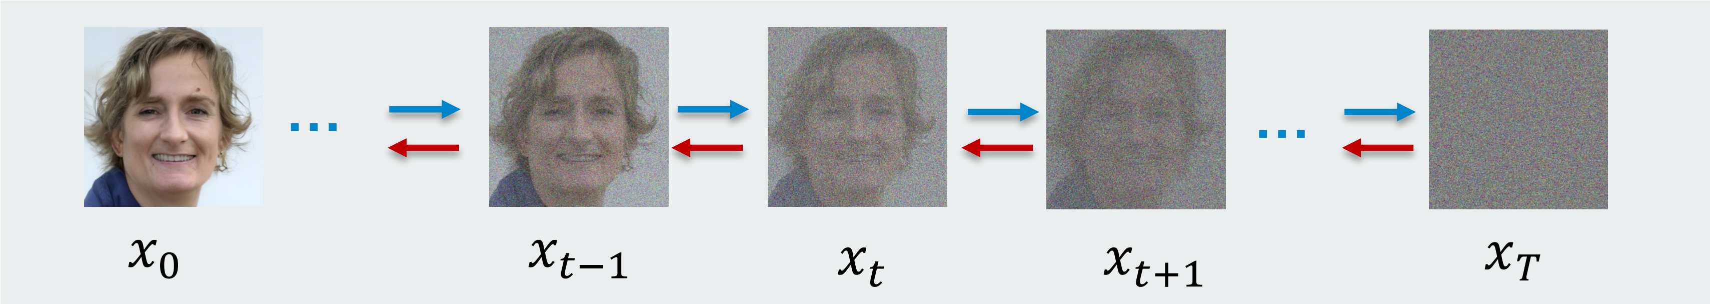

Motivation: Can we use a pretrained “unconditional” diffusion model for inverse problems?

$$\mathbf{x}_{t-1} = \frac{1}{\sqrt{\alpha_t}}\left(\mathbf{x}_t - \frac{1-\alpha_t}{\sqrt{1-\bar{\alpha}_t}}\boldsymbol{\epsilon}_\theta(\mathbf{x}_t, t)\right) + \sigma_t \mathbf{z}. \qquad \mathbf{x}_{0,t} = \frac{1}{\sqrt{\bar{\alpha}_t}}(\mathbf{x}_t - \sqrt{1-\bar{\alpha}_t}\boldsymbol{\epsilon}_\theta(\mathbf{x}_t, t))$$

J. Ho, J. Ajay and P. Abbeel. "Denoising diffusion probabilistic models." in Proceedings of *(NeurIPS)* 2020.

Imperial College London

INN and Diffusion Models for Inverse Problems

- Diffusion Models are good for “unconditional” generation of new samples (e.g., Denoising Probabilistic Diffusion Models)

- From x_T to x_0 :

- From $x_{0,T}$ to $x_{0,1}$:

J. Ho, J. Ajay and P. Abbeel. "Denoising diffusion probabilistic models." in Proceedings of *(NeurIPS)* 2020.

INN and Diffusion Models for Inverse Problems

- Given a training set $\{x_i, y_i\}$ which contains N high-quality images and their low-quality counterparts, we learn the forward part of the INN using the following loss:

$$L(\Theta) = \frac{1}{N}\sum_{i=1}^{N} \left\|\mathbf{c}^i - \mathbf{y}^i\right\|_2^2,$$

- Consequently, d models the lost details that need to be recovered with the diffusion model

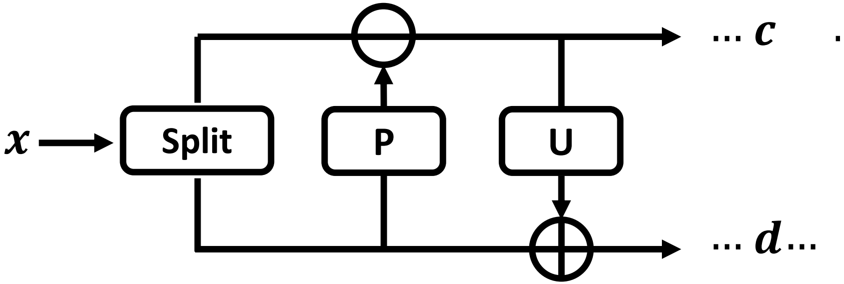

INN and Diffusion Models for Inverse Problems

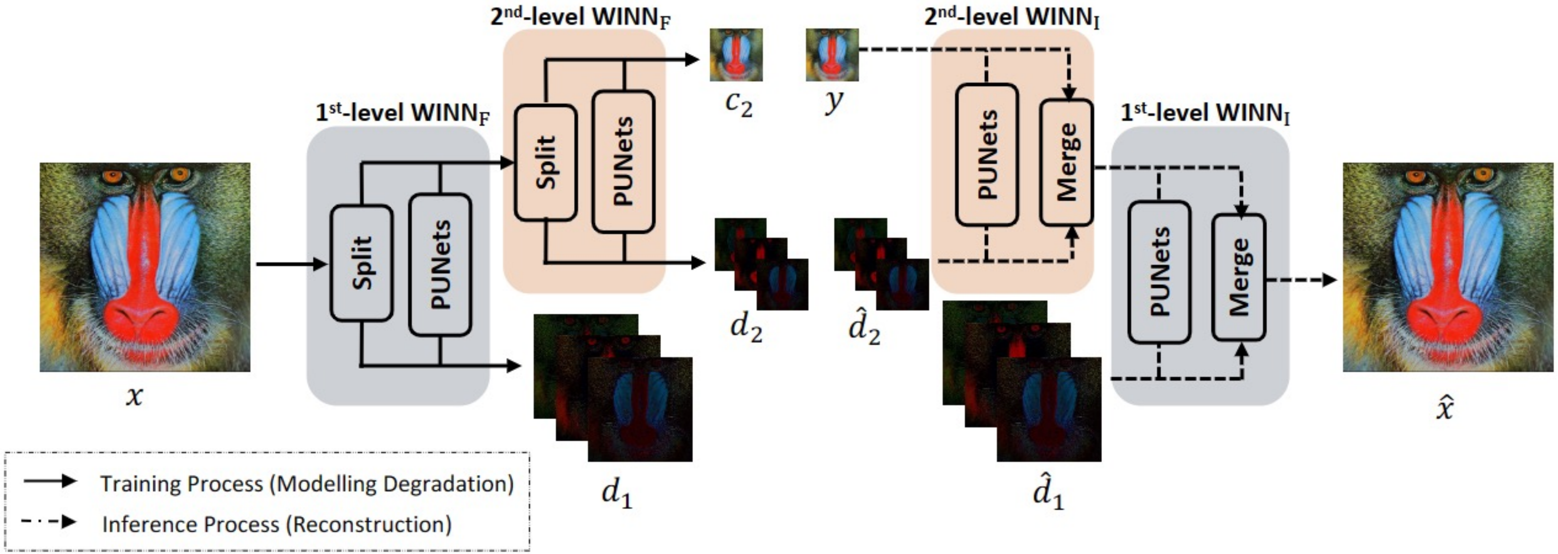

INN and Diffusion Models for Inverse Problems

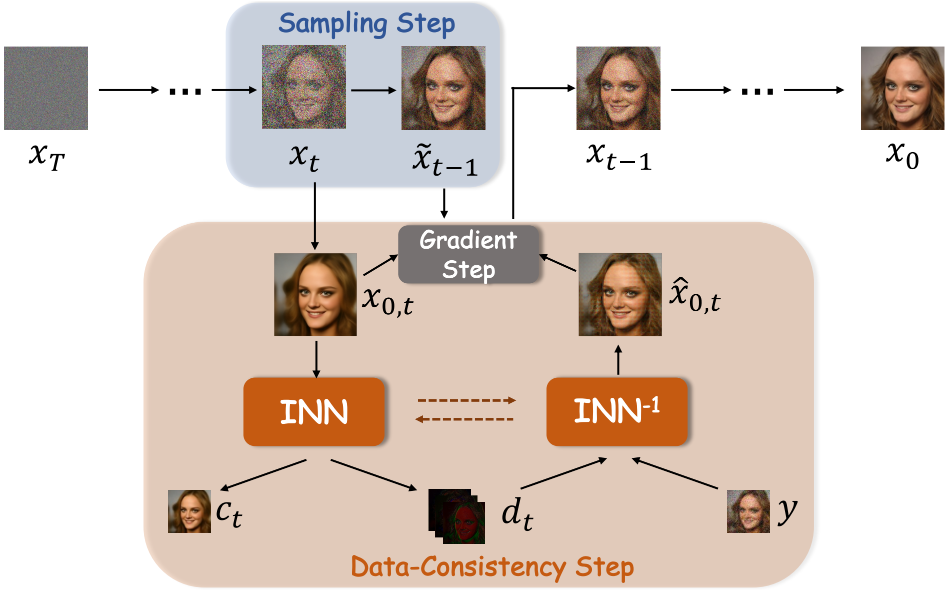

Algorithm 1 INDigo

$\mathbf{x}_T \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$
for $t = T, \ldots, 1$ **do**
  $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ if $t > 1$, else $\mathbf{z} = \mathbf{0}$
  $\mathbf{x}_{0,t} = \frac{1}{\sqrt{\bar{\alpha}_t}}(\mathbf{x}_t - \sqrt{1 - \bar{\alpha}_t}\boldsymbol{\epsilon}_\theta(\mathbf{x}_t, t))$
  $\tilde{\mathbf{x}}_{t-1} = \frac{\sqrt{\alpha_t}(1-\bar{\alpha}_{t-1})}{1-\bar{\alpha}_t}\mathbf{x}_t + \frac{\sqrt{\bar{\alpha}_{t-1}}\beta_t}{1-\bar{\alpha}_t}\mathbf{x}_{0,t} + \sigma_t \mathbf{z}$
  $\mathbf{c}_t, \mathbf{d}_t = f_\phi(\mathbf{x}_{0,t})$
  $\hat{\mathbf{x}}_{0,t} = f_\phi^{-1}(\mathbf{y}, \mathbf{d}_t)$
  $\mathbf{x}_{t-1} = \tilde{\mathbf{x}}_{t-1} - \zeta \nabla_{\mathbf{x}_t} \|\hat{\mathbf{x}}_{0,t} - \mathbf{x}_{0,t}\|_2^2$
end for
return $\mathbf{x}_0$

INN and Diffusion Models for Inverse Problems

Ground Truth

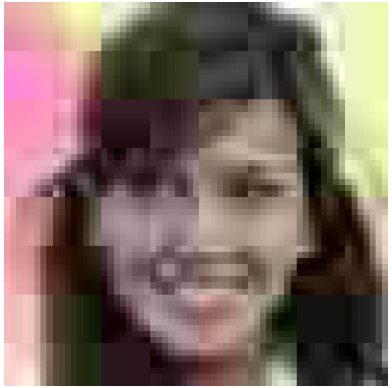

Degraded

Reconstructed

- This approach is simple, flexible and effective
  - No-need to know the degradation process
  - The degradation process can be highly non-linear
  - No need to retrain the diffusion model for every new degradation (just need to train the INN)

INN and Diffusion Models for Inverse Problems

Results for non-linear degradation models

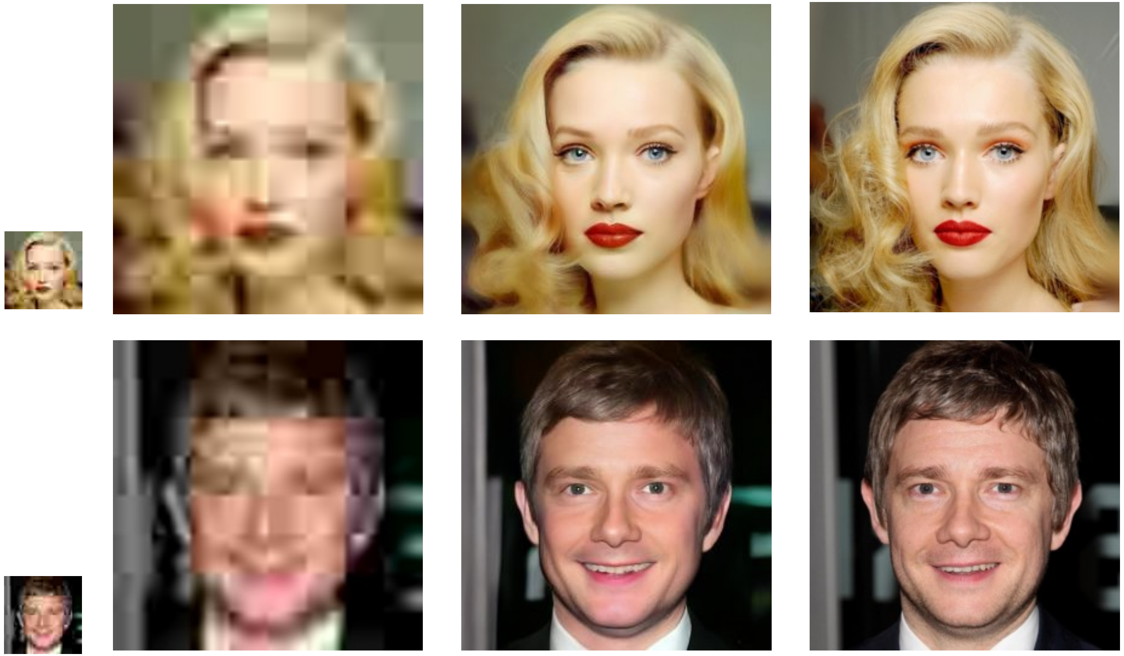

Input Bicubic Ours Ground Truth

Imperial College London

INN and Diffusion Models for Inverse Problems

Results on 4x super-resolution

Method	Noise σ	PSNR $\uparrow$	FID $\downarrow$	LPIPS $\downarrow$	NIQE $\downarrow$
ILVR	0	27.43	44.04	0.2123	5.4689
DDRM	0	28.08	65.80	0.1722	4.4694
DPS	0	26.67	32.44	0.1370	4.4890
Ours	0	**28.15**	**22.33**	**0.0889**	**4.1564**
ILVR	0.05	26.42	60.27	0.3045	4.6527
DDRM	0.05	27.06	45.90	0.2028	4.8238
DPS	0.05	25.92	31.71	0.1475	4.3743
Ours	0.05	**27.16**	**26.64**	**0.1215**	**4.1004**
ILVR	0.10	24.60	88.88	0.4833	4.4888
DDRM	0.10	26.16	45.49	0.2273	4.9644
DPS	0.10	24.73	31.66	0.1698	4.2388
Ours	0.10	**26.25**	**28.89**	**0.1399**	**3.9659**

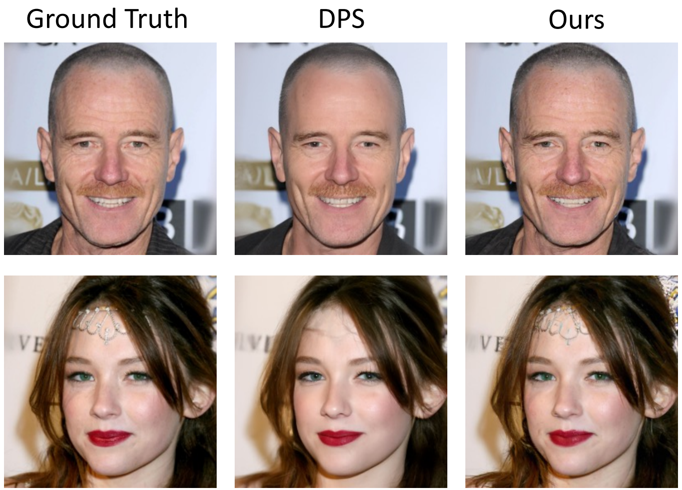

Results on blind unsupervised deconvolution

(a) Input (b) DR2 [1] (c) DifFace [2] (d) PGDiff [3] (e) StableSR [4] (f) **Ours**

D.You and P.L. Dragotti, “INDIGO+: A Unified INN-Guided Probabilistic Diffusion Algorithm for Blind and Non-Blind Image Restoration”, IEEE Journal of Selected Topics in Signal Processing, 2024

First Set of Conclusions

- Invertible Neural Networks are an interesting new concept
- Designing INN and combining them with diffusion models (plug-and-play) leads to more interpretable and simpler architectures
- Good performance and good generalization ability
- Potential for further developments

Sparsity and Deep Unfolding Strategy

Explicit embedding of priors and constraints in deep networks

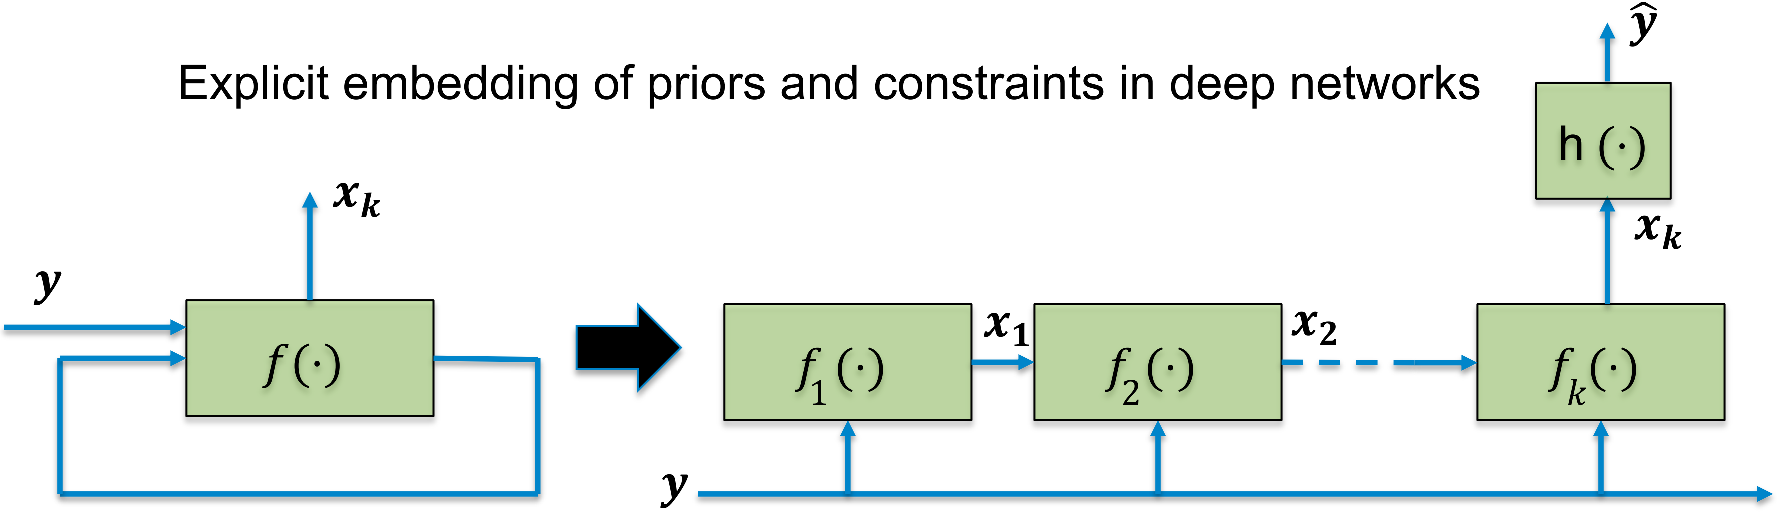

Iterative algorithm with $\boldsymbol{y}$ as input and $\boldsymbol{x}$ as output

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

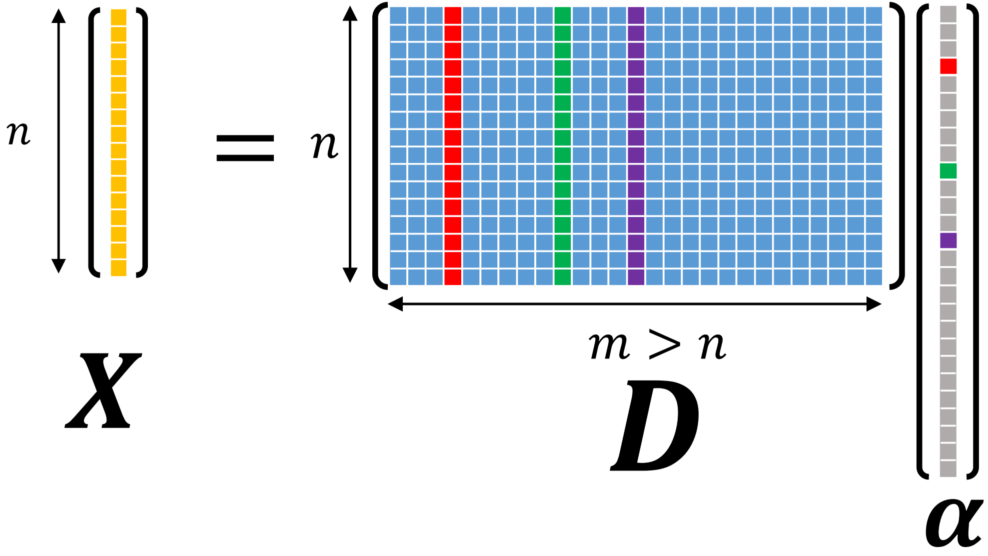

Imperial College London

Deep Unfolding Strategy

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

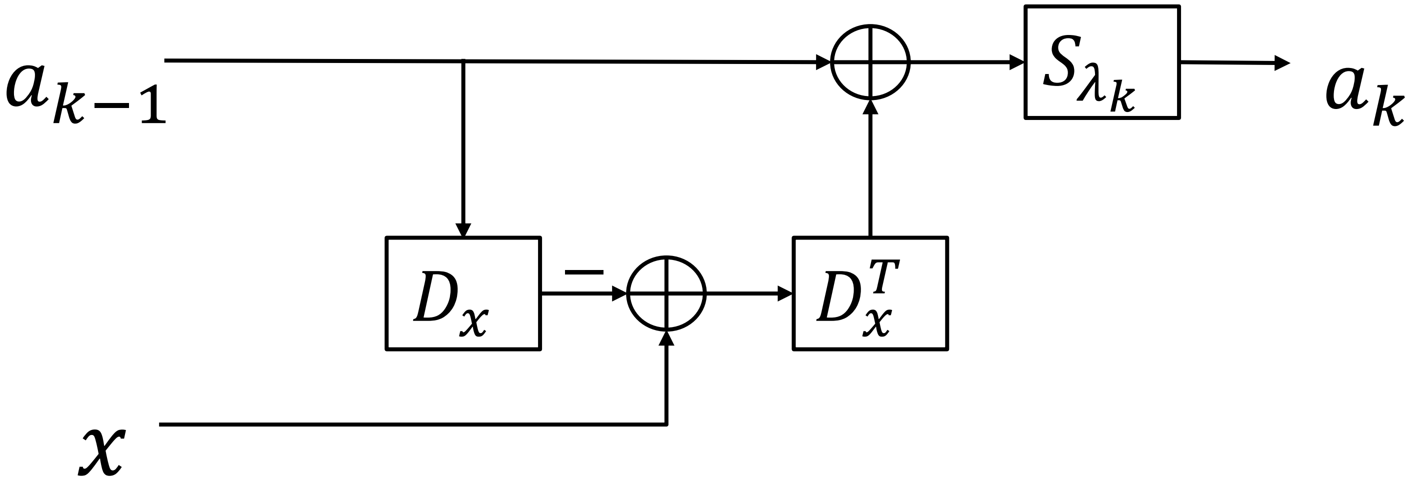

ISTA network

- Solving by ISTA algorithm through unfolding:

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

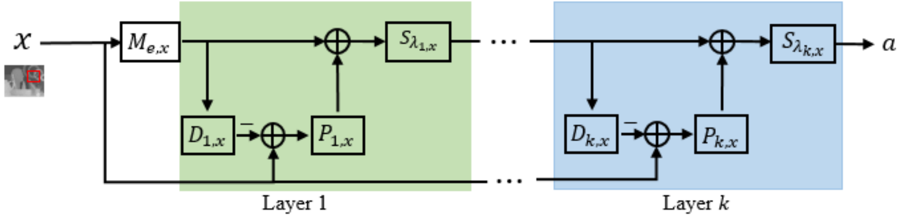

- Gregor Karol and LeCunYann, “Learning fast approximations of sparse coding ”, Proceedings of the 27th International Conference on International Conference on Machine Learning, 2010
- Y. Eldar et al, “Algorithm Unrolling: Interpretable, Efficient Deep Learning for Signal and Image Processing”, IEEE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

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,

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

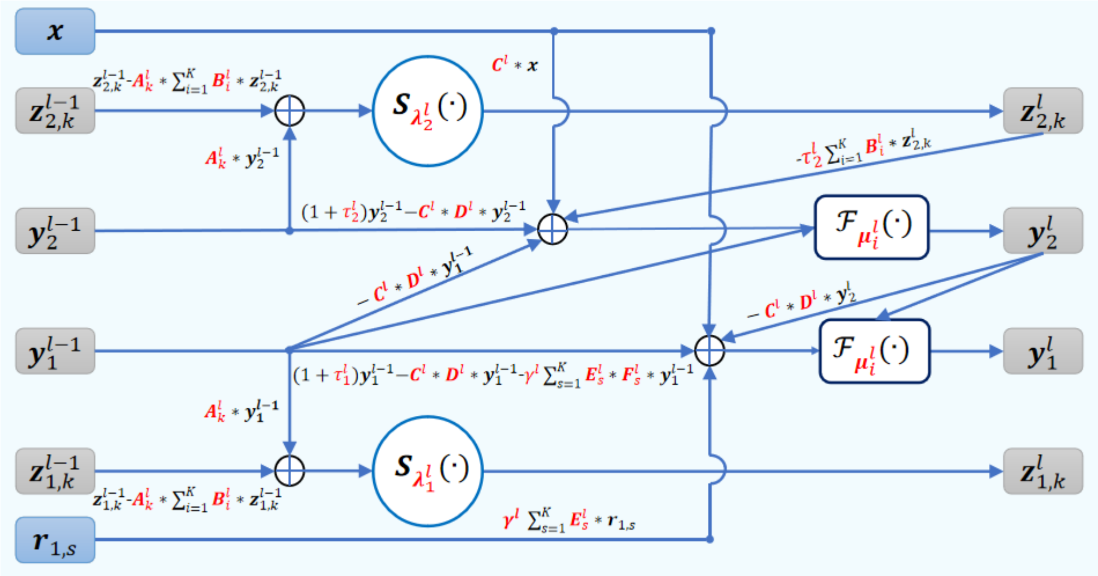

Separation Results

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

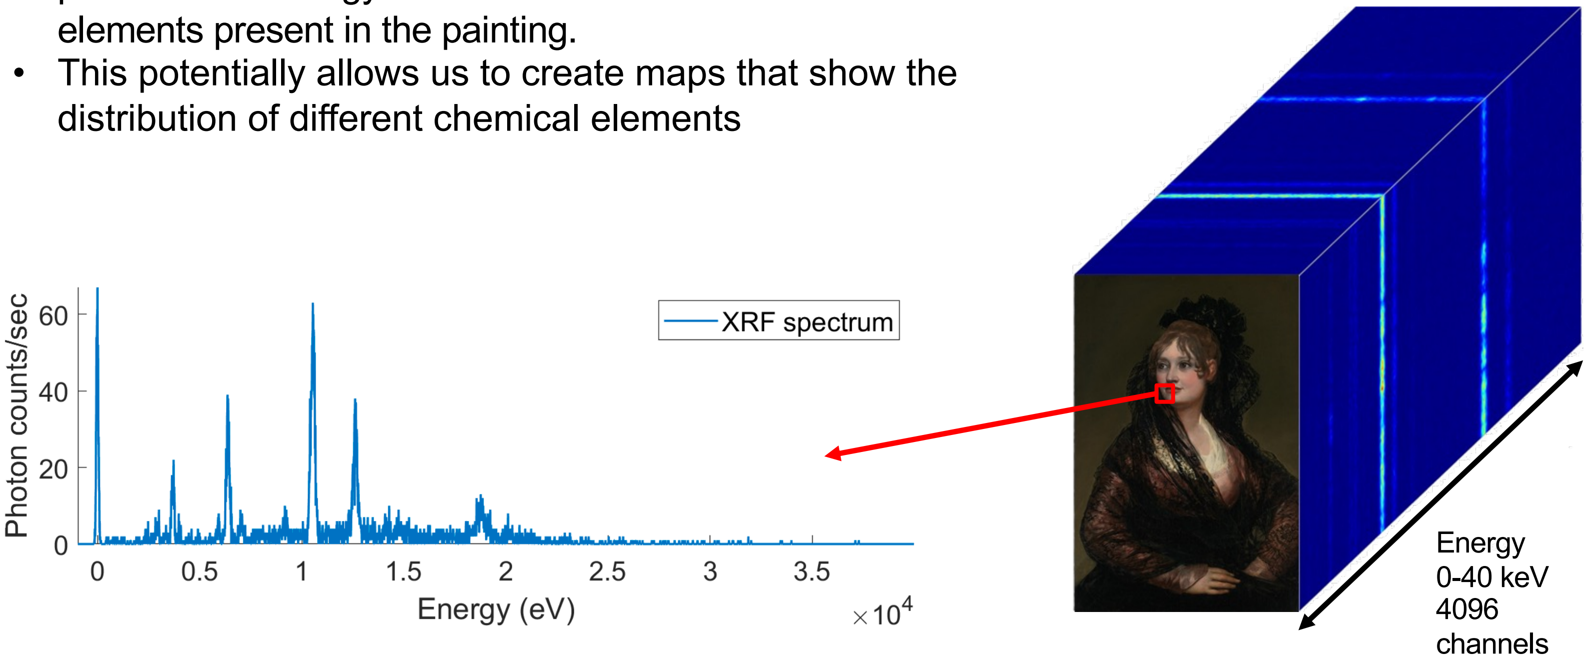

Imperial College London

Extraction of Elemental Maps

Our XRF Deconvolution Algorithm

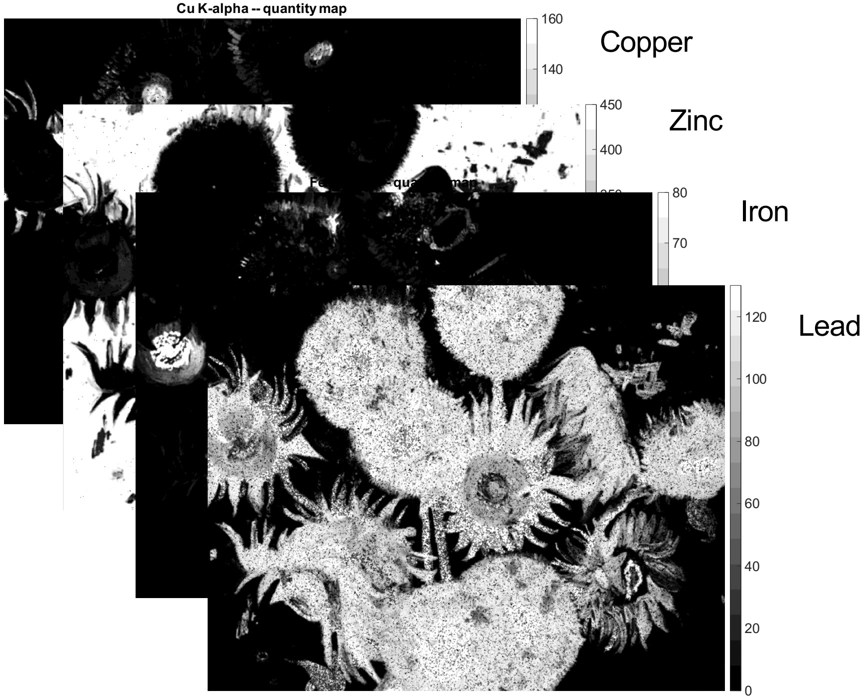

Results

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

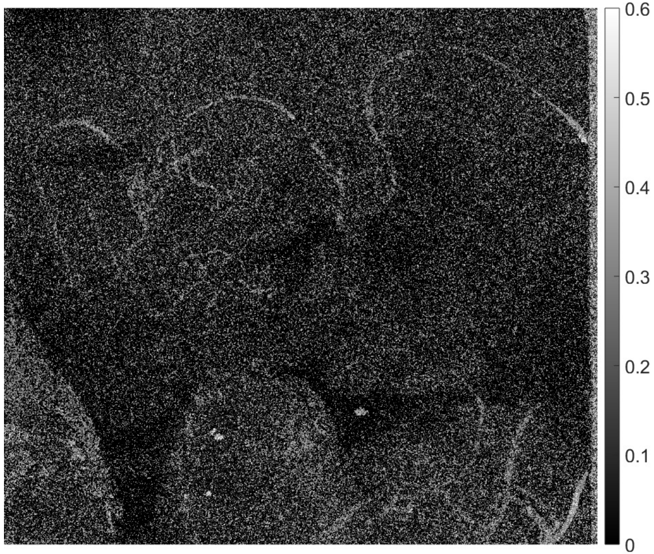

Zn confidence map

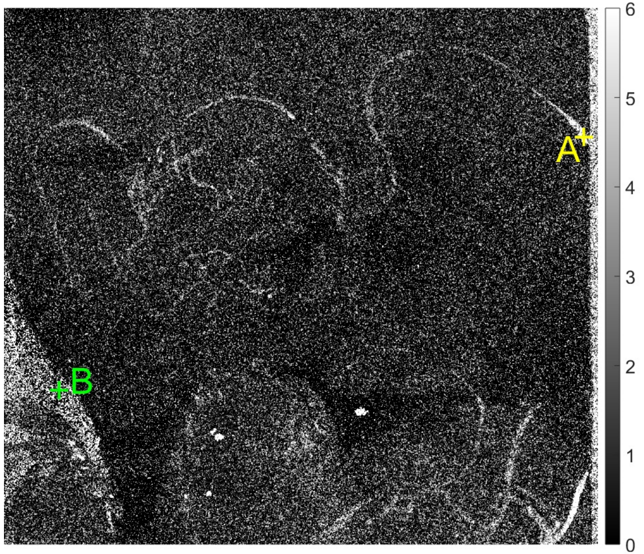

Zn quantity map

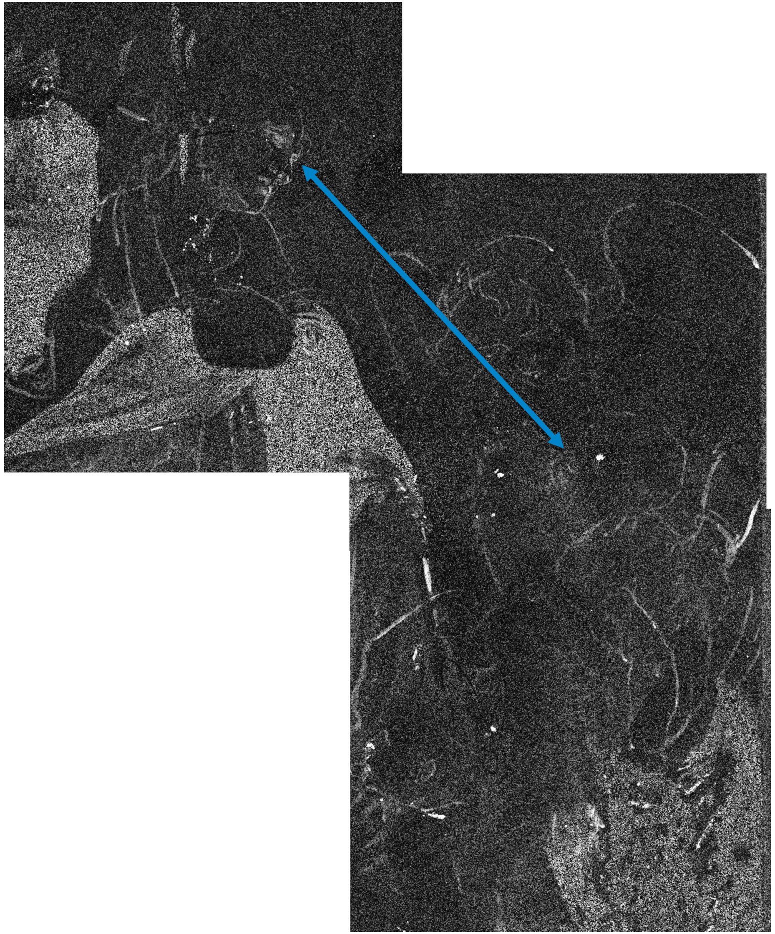

S. Yan, J.-J. Huang, N. Daly, C. Higgitt, and P. L. Dragotti, “When de Prony Met Leonardo: An Automatic Algorithm for Chemical Element Extraction in Macro X-ray Fluorescence Data”, IEEE Transactions on Computational Imaging, vol.7, 2021.

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

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

Point scanning (2PLSM)

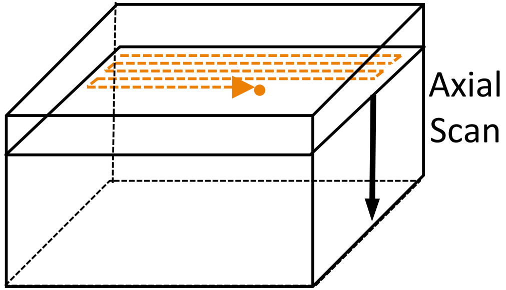

Two-Photon Microscopy

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

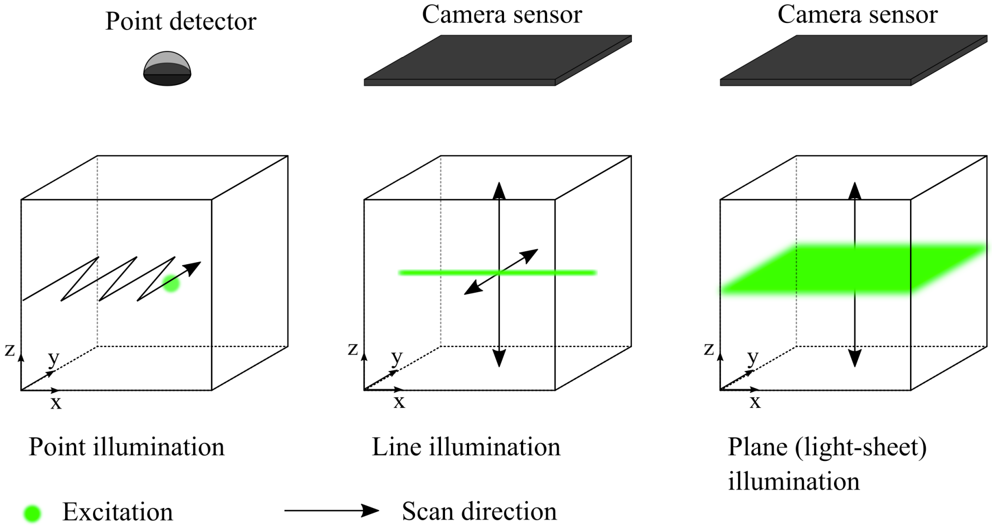

Imperial College London

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

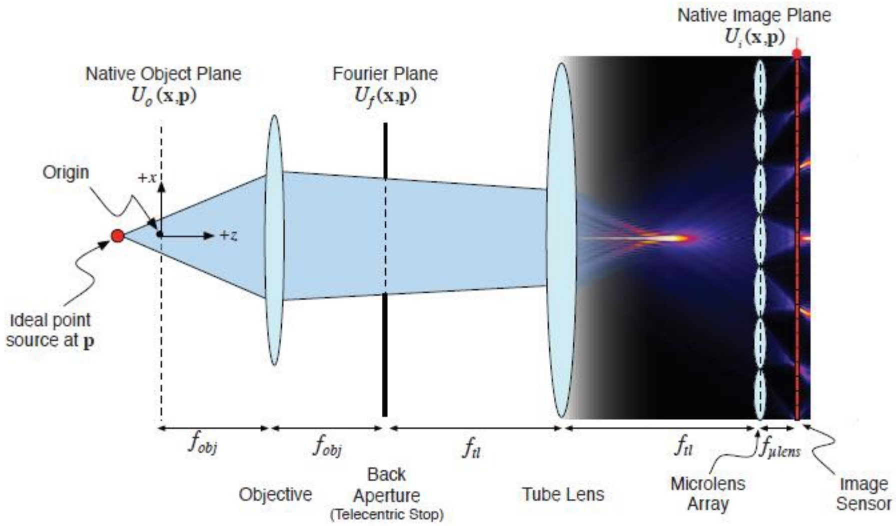

Imperial College London

Light-field Microscopy and EPI

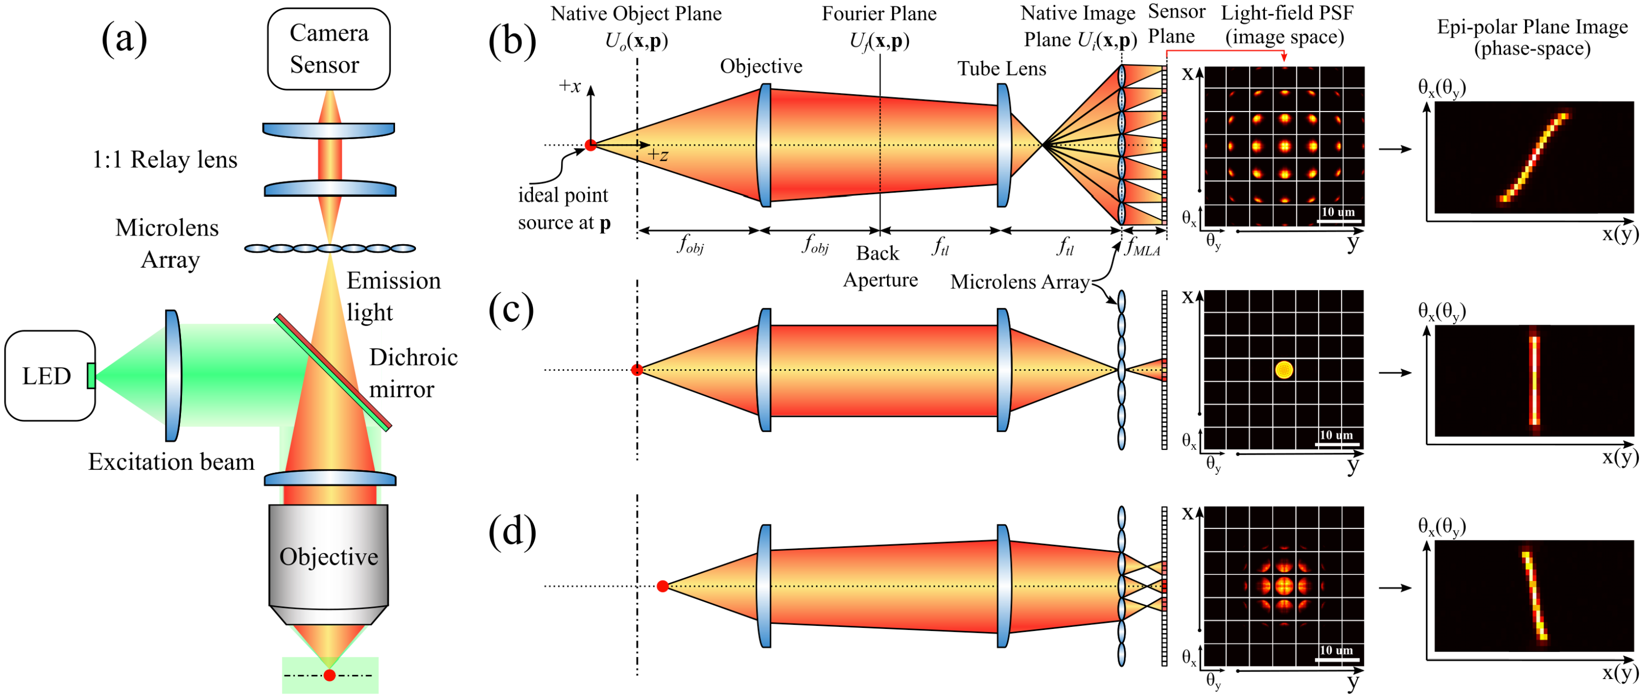

Imperial College London

Light-field Microscopy and Illumination Strategies

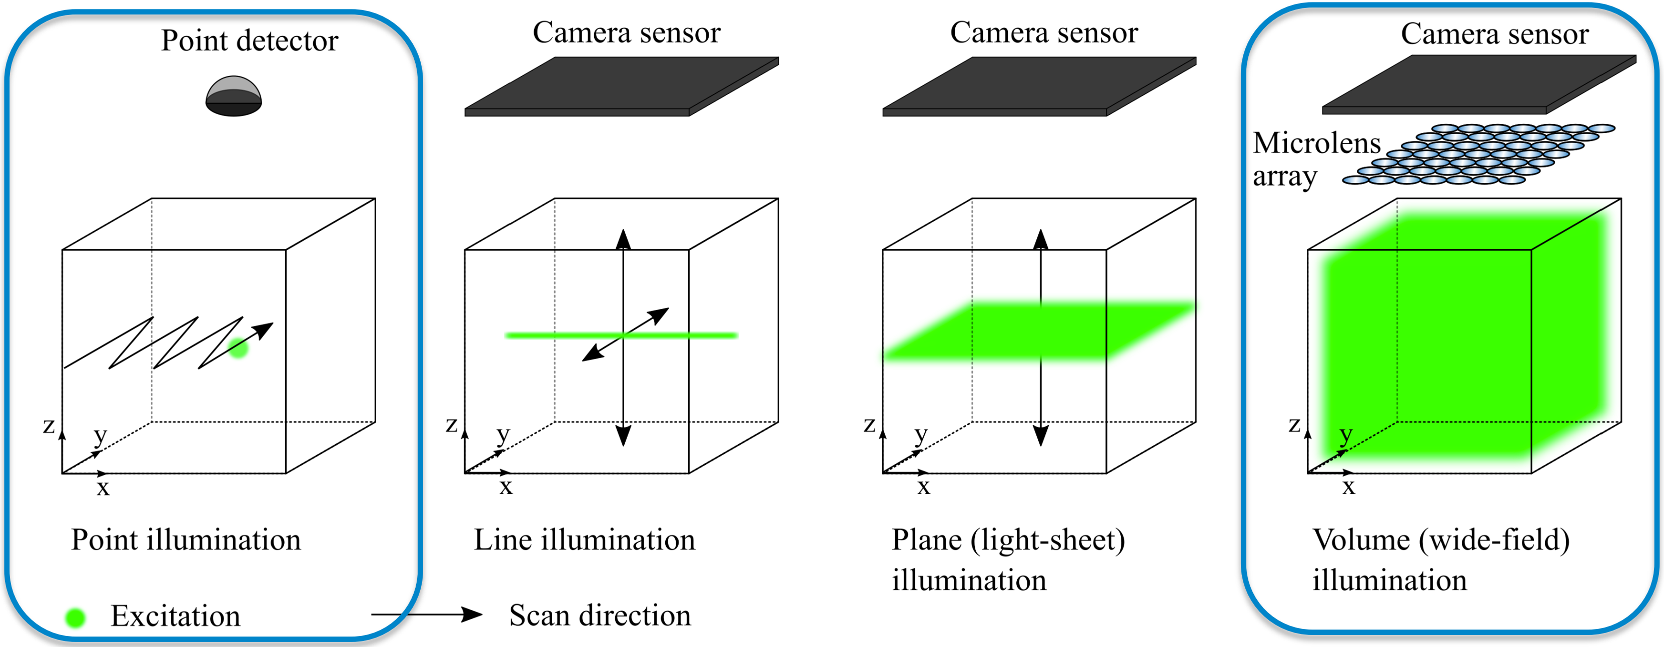

Key insight: use the 2P microscope for high-resolution structural information and the LFM for monitoring the activity of neurons.

Imperial College London

Our Solution: Scattering-robust structural volumes + high-bandwidth, scanless functional volumes

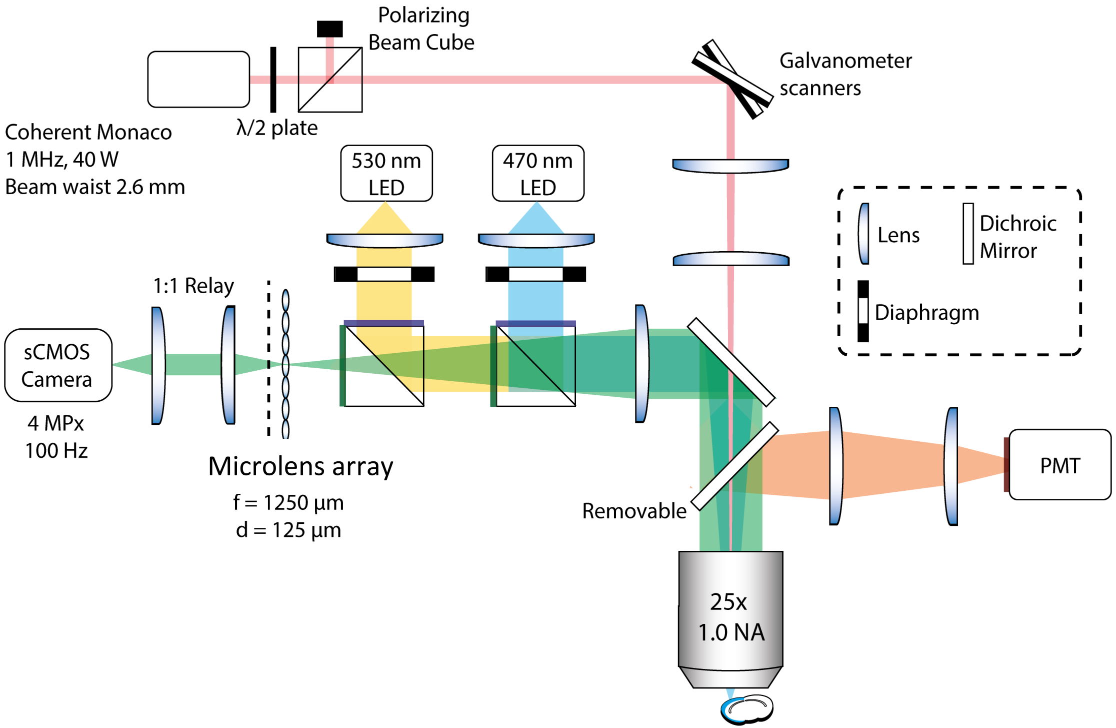

Light-field Microscopy

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

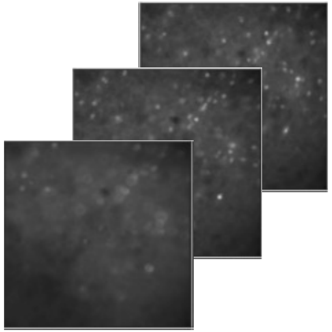

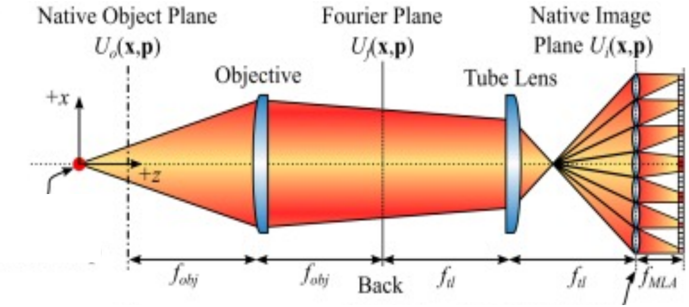

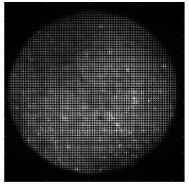

2D Measured LF image

3D Input

Computational
Algorithm

Volume Reconstruction from Light-field 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

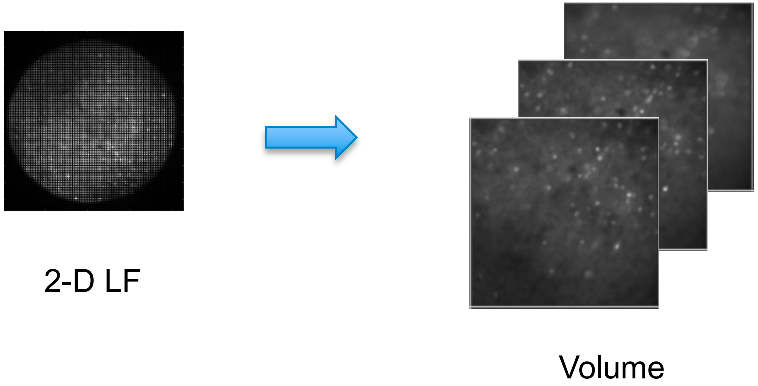

Imperial College London

Forward Model

- Forward model is linear which means $\boldsymbol{y} = \boldsymbol{Hx}$
  - $\boldsymbol{H}$ 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)

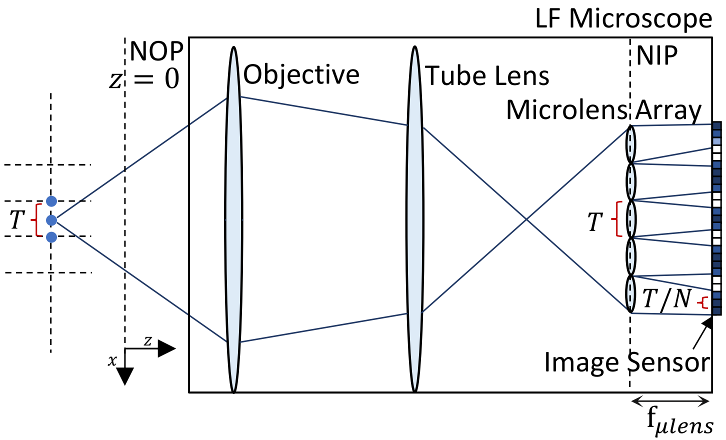

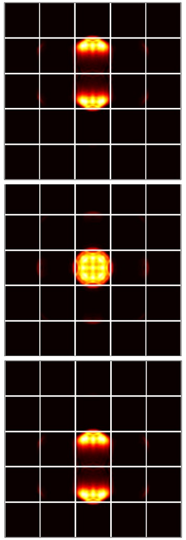

Neural Network for Volume Reconstruction

- 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

Neural Network for Volume Reconstruction

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

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

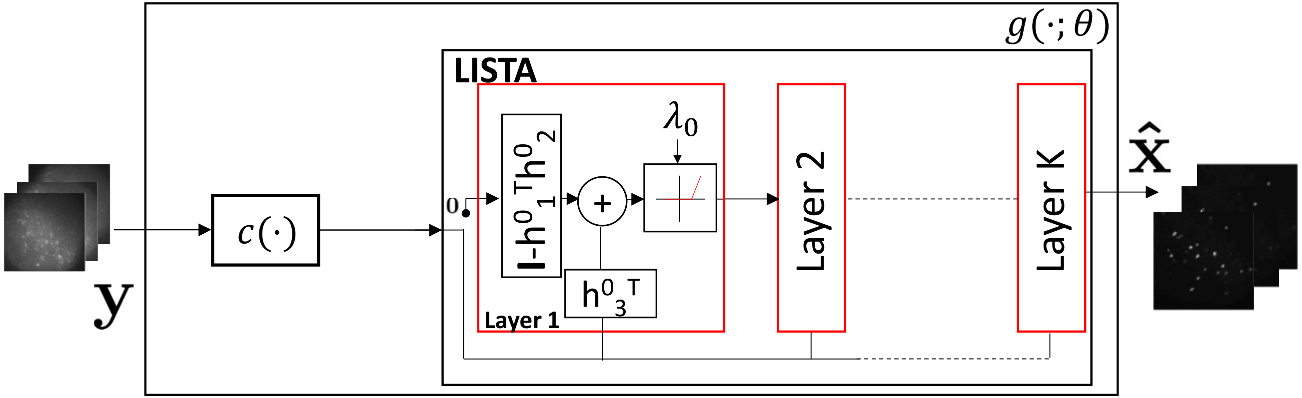

Imperial College London

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
    - Light-field loss based on re-synthesizing light-field from reconstructed volume
    - Adversarial network for adversarial loss

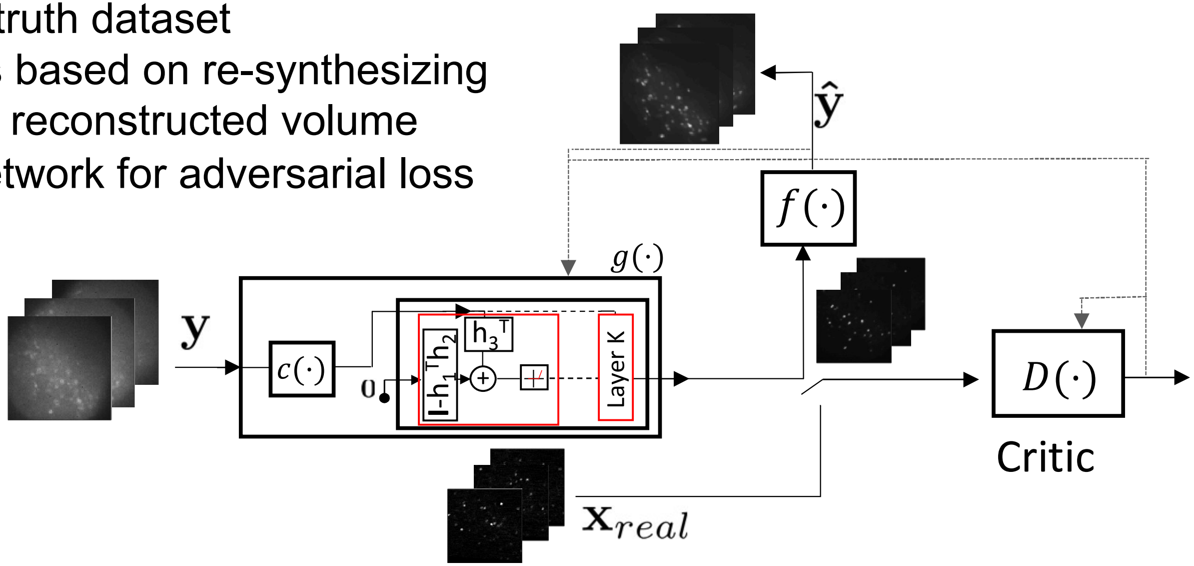

Imperial College London

Results – Structural Data

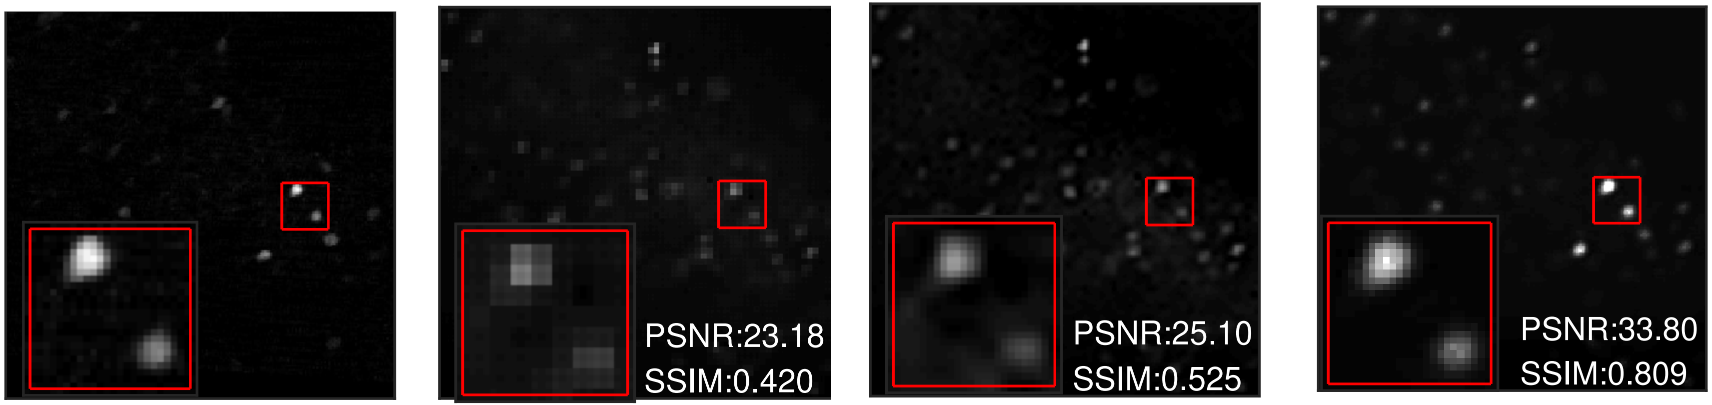

Ground-truth ISRA ADMM

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

Fast volumetric jGCaMP8f time-series extraction

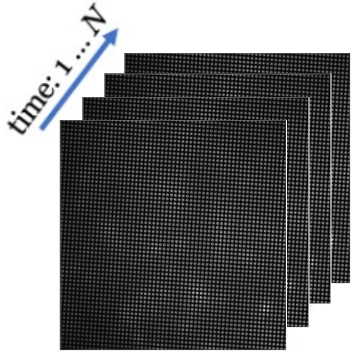

Input F(t):
LF video acquired in brain slice cortical layer 2/3

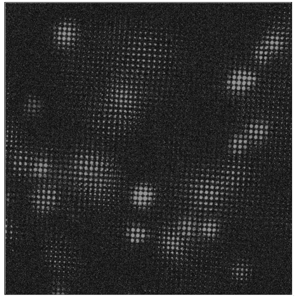

LF activity map: max(ΔF)/σ(F)

LISTA-based net Volume Estimation

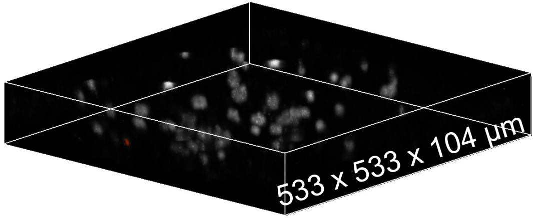

8-iteration Richardson-Lucy Deconvolution

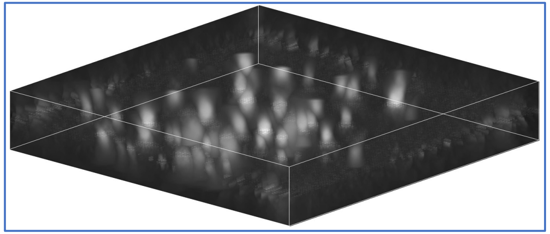

Imperial College London

Fast volumetric jGCaMP8f time-series extraction

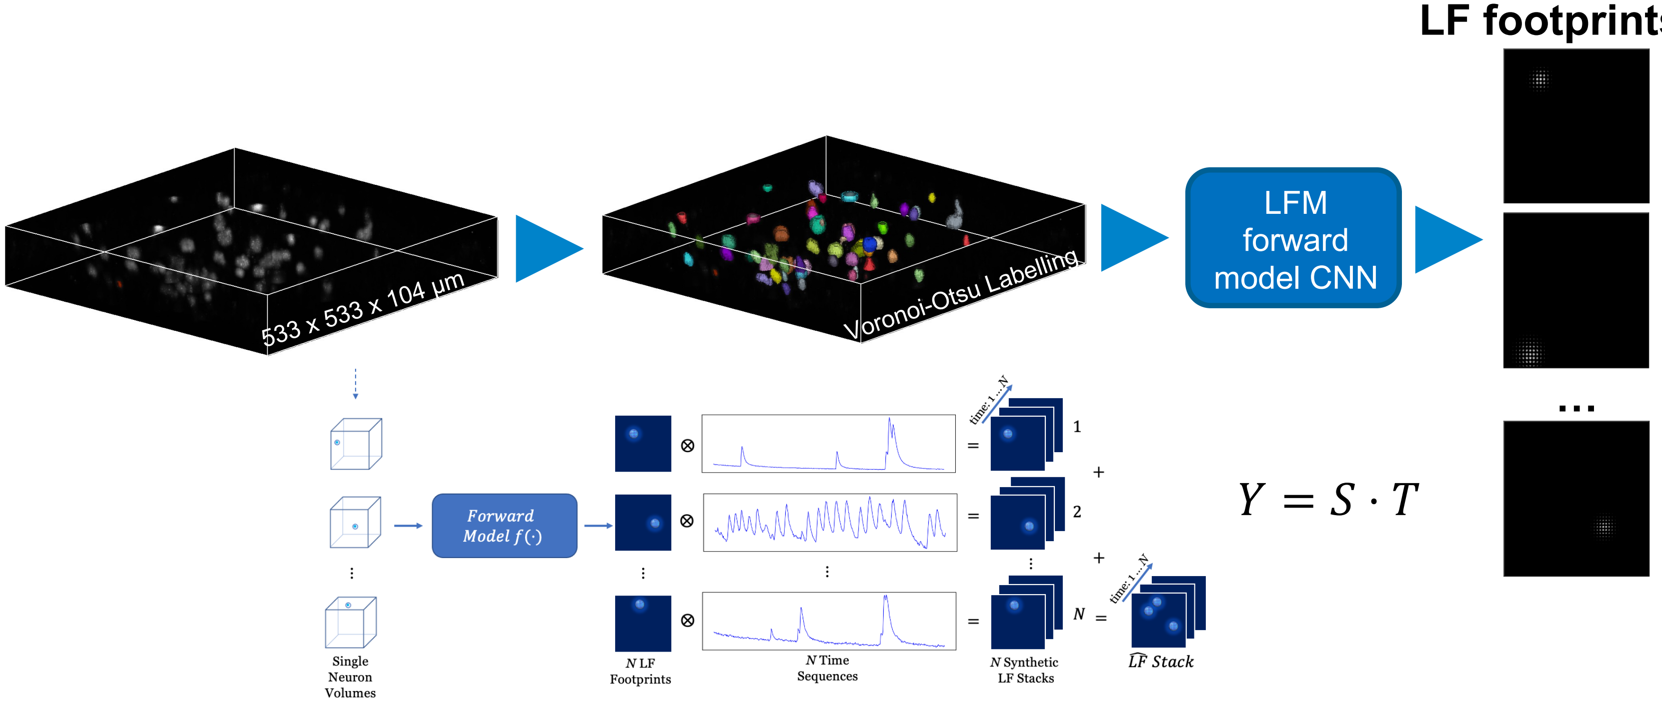

Imperial College London

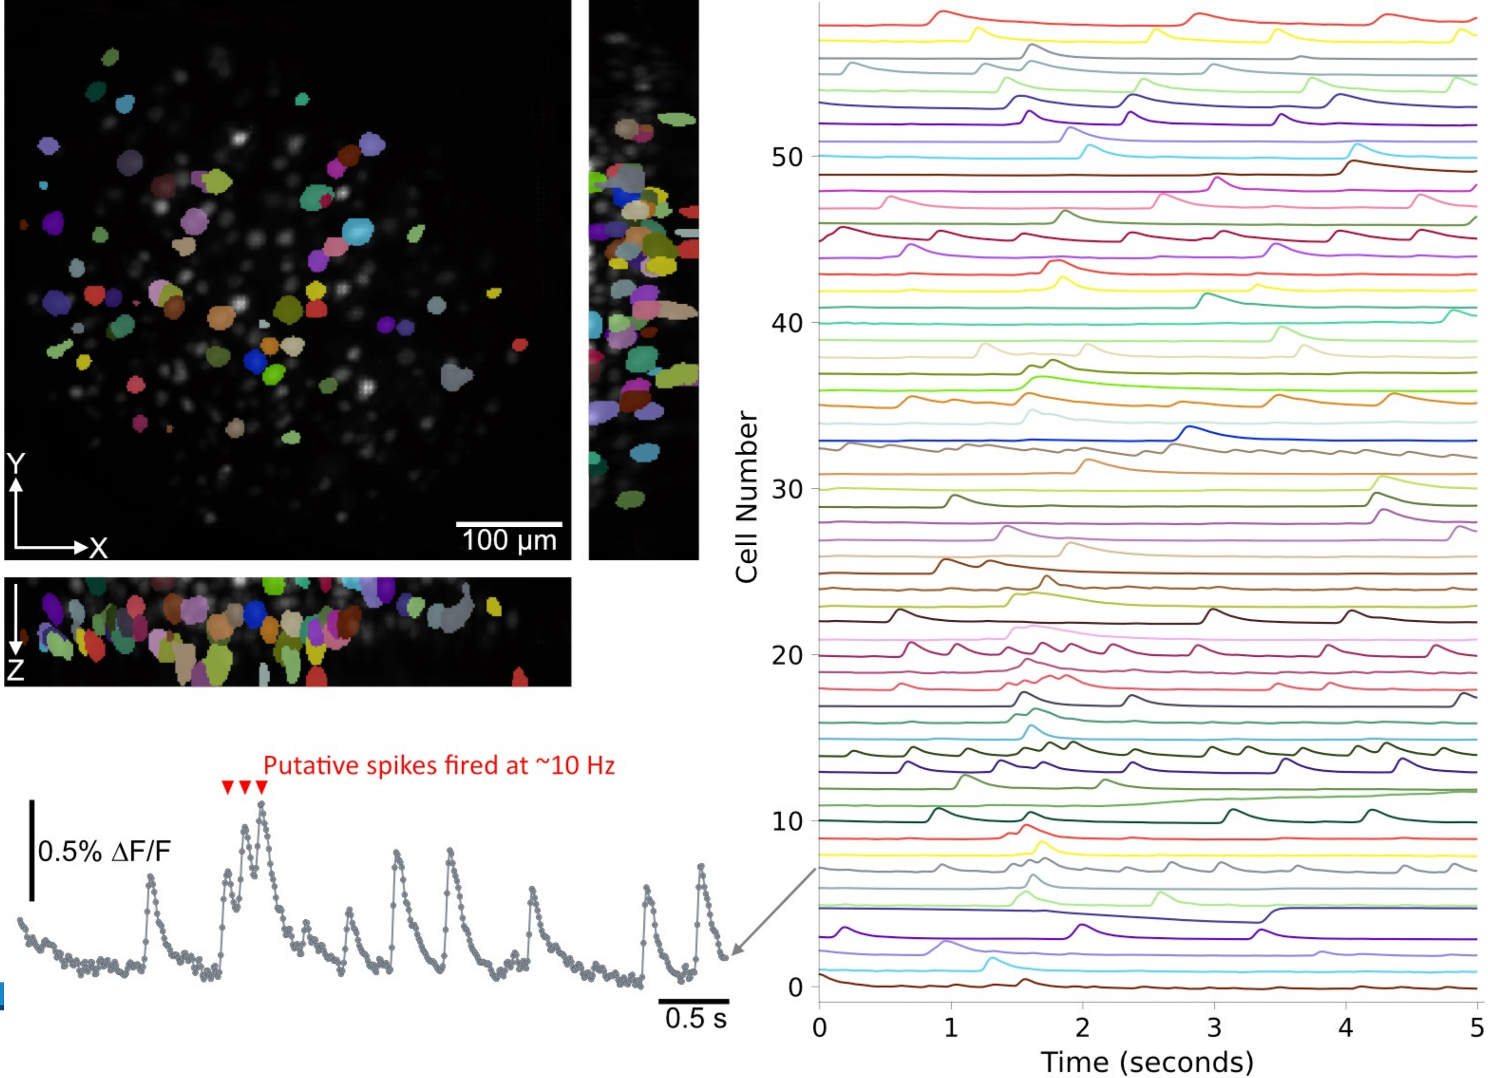

LISTA-based net decreases crosstalk between neighbouring neurons

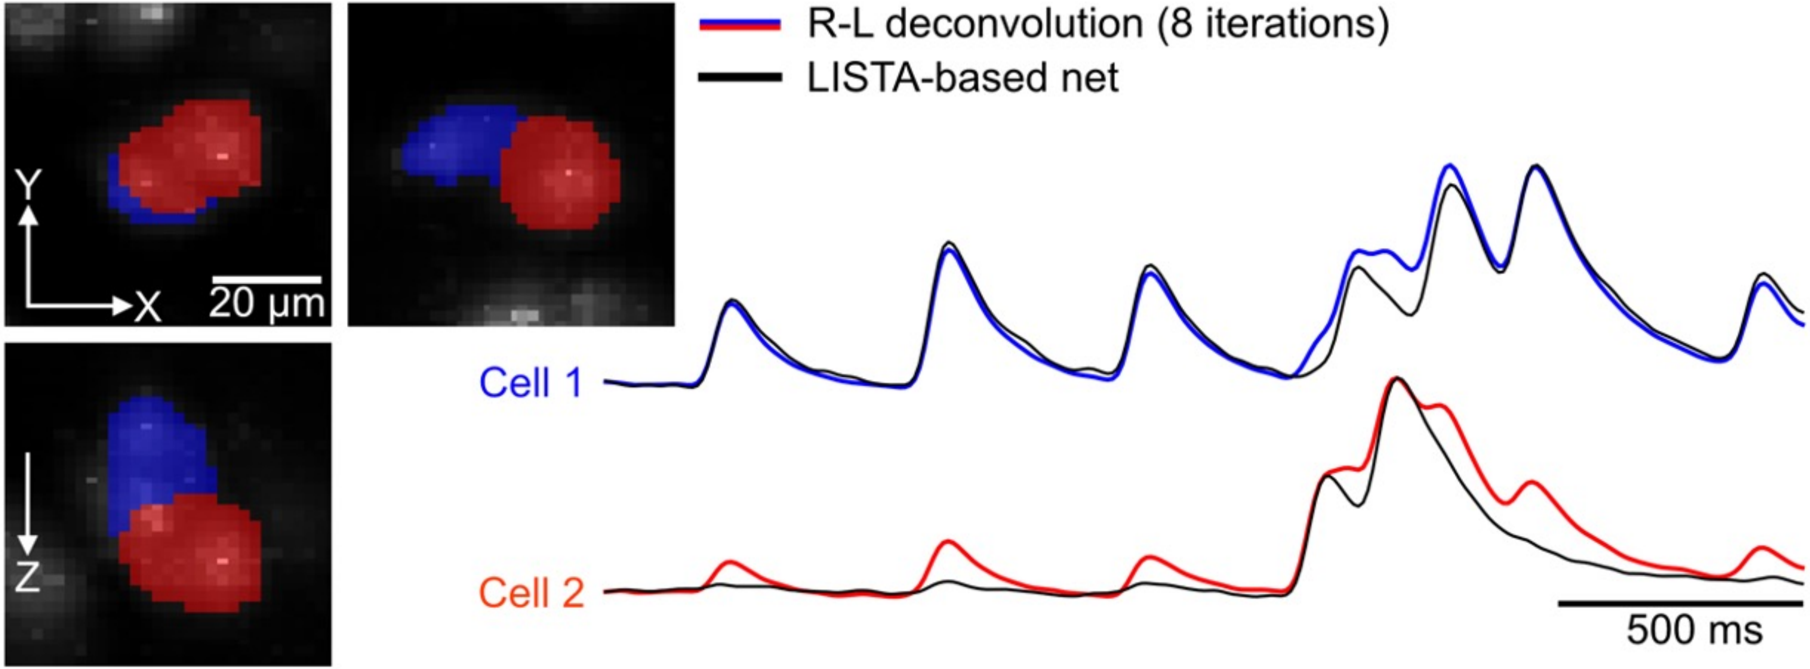

Conclusions

- In imaging problems:
  - operating at the interface between physics and computation is essential
  - Cross fertilization between model-based approaches and deep learning is fruitful
  - Some computational approaches are transferable

- Inverse imaging problems:
  - are fun 🙂
  - and inter-disciplinary

Imperial College London

A special thank to:


Junjie Huang

Pingfan Song

Herman Verinaz

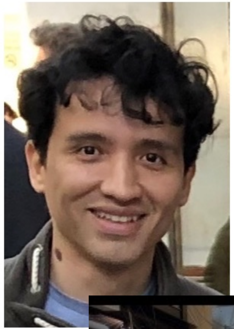

Kate Zhao

Peter Quicke

Carmel Howe

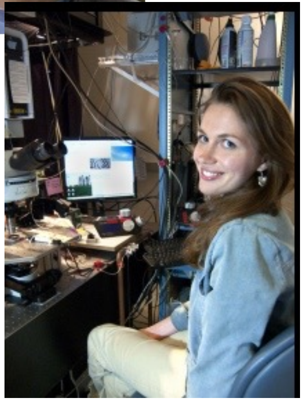

Amanda Foust

Consortium involving: UCL, ICL, Duke and National Gallery

Imperial College London

Thank you!

Imperial College London

Related Publications

- Wavelet-inspired INN and Diffusion Models:
  - J. Huang and P.L. Dragotti, “WINNet: Wavelet-inspired Invertible Network for Image Denoising”, IEEE Transactions on Image Processing, 2022, **software**: https://github.com/pld-group/WINNet
  - D.You and P.L. Dragotti, “INDIGO+: A Unified INN-Guided Probabilistic Diffusion Algorithm for Blind and Non-Blind Image Restoration”, IEEE Journal of Selected Topics in Signal Processing, 2024, **software**: https://github.com/pld-group/indigo_plus

- Light-field Microscopy:
  - 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.

- Art Investigation
  - 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.
  - S Yan, JJ Huang, H Verinaz-Jadan, N Daly, C Higgitt, PL Dragotti, “A fast automatic method for deconvoluting macro X-ray fluorescence data collected from easel paintings”, IEEE Transactions on Computational Imaging, 2023, **software**: https://github.com/pld-group/XRF_fast_deconvolution