Imperial College London

Inverse Problems in the Age of AI

Pier Luigi Dragotti, Imperial College London

21 May 2025

Imperial College London

What do I do for a living: Inverse Imaging Problems

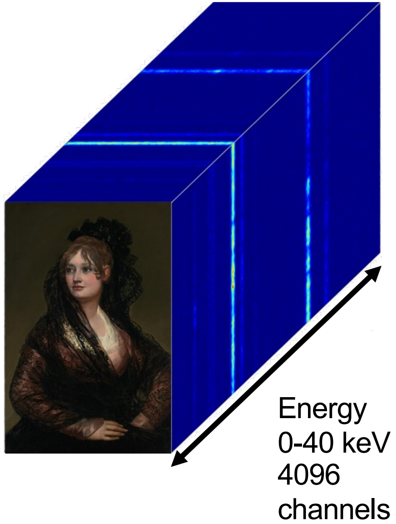

Technical study of Old Masters paintings

Image restoration problems

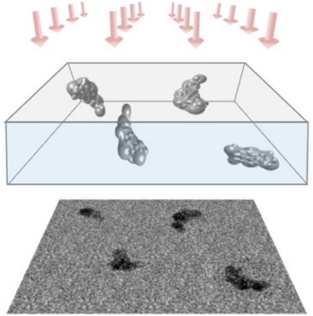

Cryo-EM for structural biology

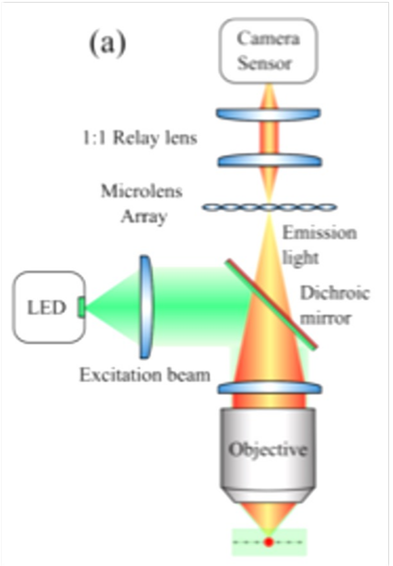

Light field microscopy for neuroscience

Imperial College London

Motivation

- Inverse problems involve reconstructing unknown physical quantities from indirect measurements.
- Recently we have witnessed a shift towards using data-driven methods to address inverse problems.
- In this talk, I cover three case studies to highlight advantages and pitfalls of data-driven solutions

Imperial College London

Three Case Studies in Imaging Science

Image restoration problems: Invertible neural networks and diffusion models

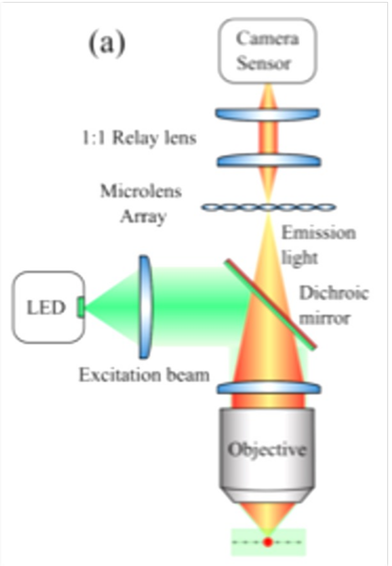

Light field microscopy for neuroscience

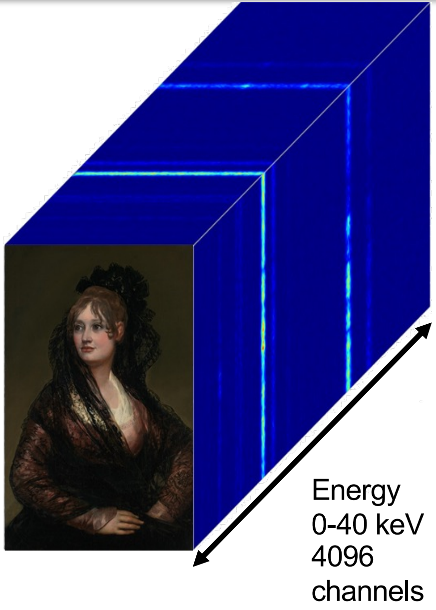

Technical study of Old Masters paintings

Imperial College London

Generative Models for Image Restoration

Image restoration problems: Invertible neural networks and diffusion models

INN and Diffusion Models for Inverse Problems

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

  consistency term prior

INN and Diffusion Models for Inverse Problems

- $\hat{x} = \min_{x} \; \mathcal{L}(x, y) + \quad \lambda\rho(x)$

  consistency term   prior

- $\hat{x} = \min_{x,v} \mathcal{L}(x, y) \quad + \lambda\rho(v)$ s.t $x = v$

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

- Solve by alternating between x and v

  - Consistency step: $\hat{x} = v + \eta\nabla_{\mathrm{x}}\mathcal{L}(v, \mathrm{y})$
  - 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

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

and inverse mapping

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

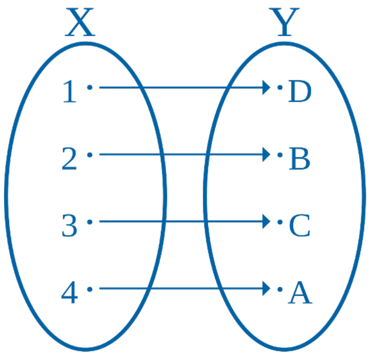

A bijective function (or invertible function)

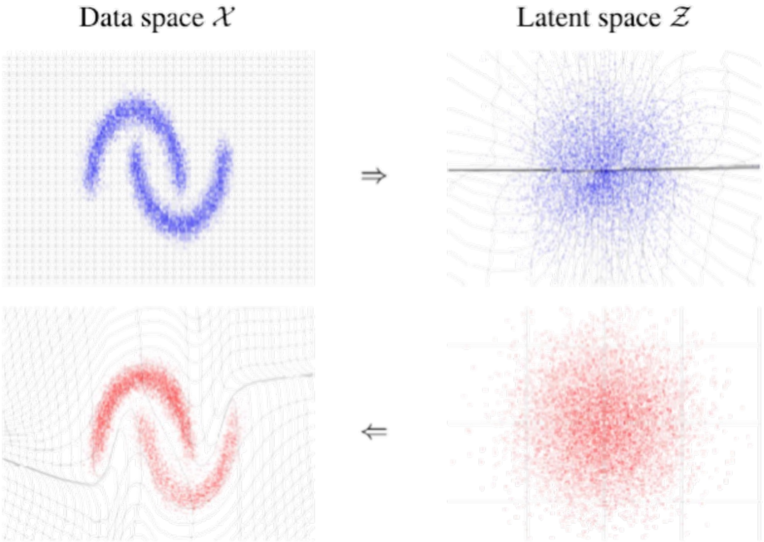

Kingma, Durk P., and Prafulla Dhariwal. "Glow: Generative flow with invertible 1x1 convolutions." in Proceedings of *Advances in Neural Information Processing Systems (NeurIPS)*, 2018.

How to Achieve Invertibility?

Invertible via lifting scheme

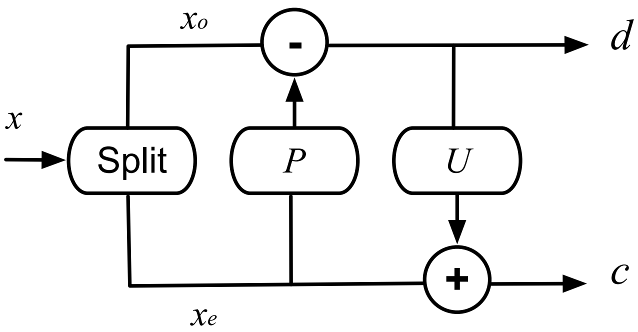

Split → $$\begin{cases} d = x_o - P(x_e) \\ c = x_e + U(d) \end{cases}$$

Forward pass

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

Invertible Neural Networks

How to Achieve Invertibility?

Invertible via lifting scheme

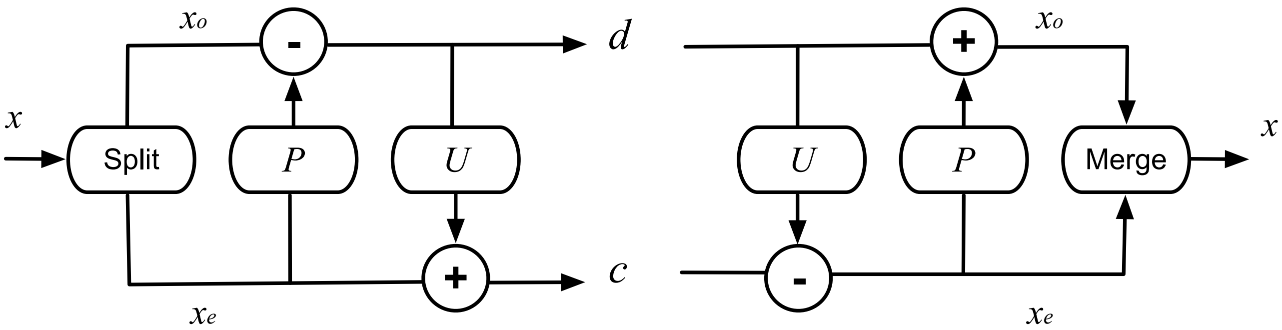

Split → $\begin{cases} d = x_o - P(x_e) \\ c = x_e + U(d) \end{cases}$

Forward pass

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

Backward pass

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

Invertible Neural Networks and Wavelets

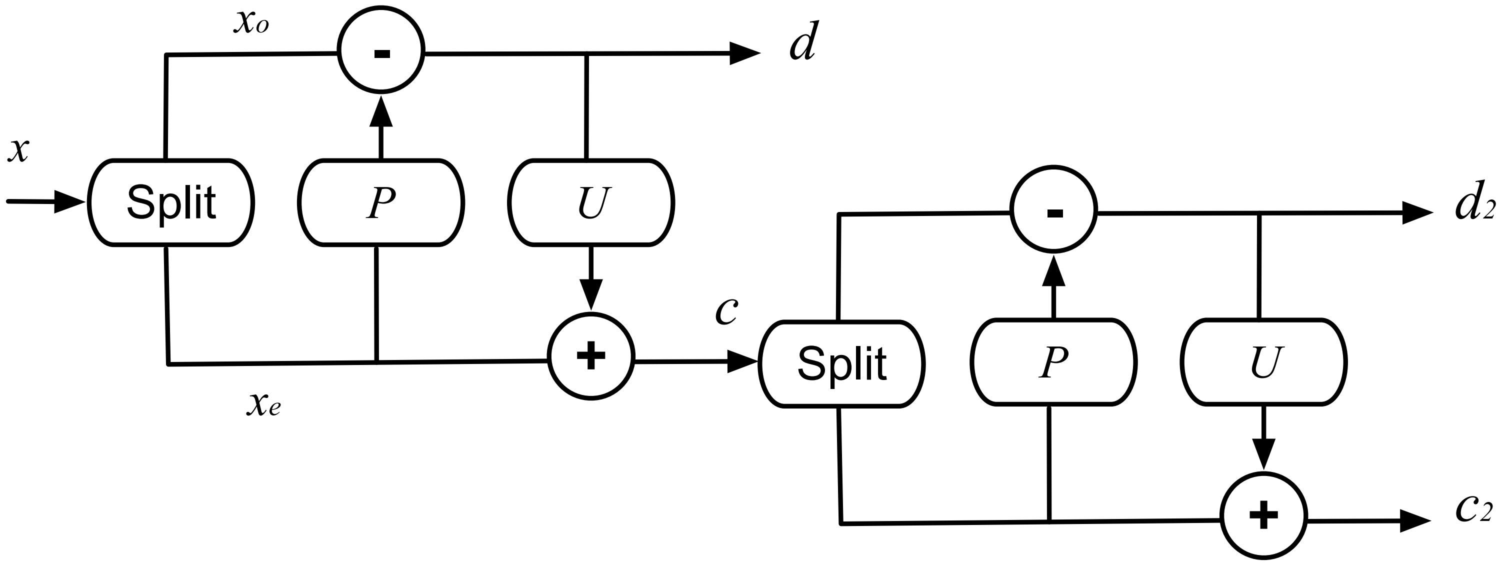

I. Daubechies and W. Sweldens, “Factoring Wavelet Transforms into lifting Steps” 1997

Imperial College London

INN and Diffusion Models for Inverse Problems

Invertible Neural Networks are ideal architectures to address inverse problems

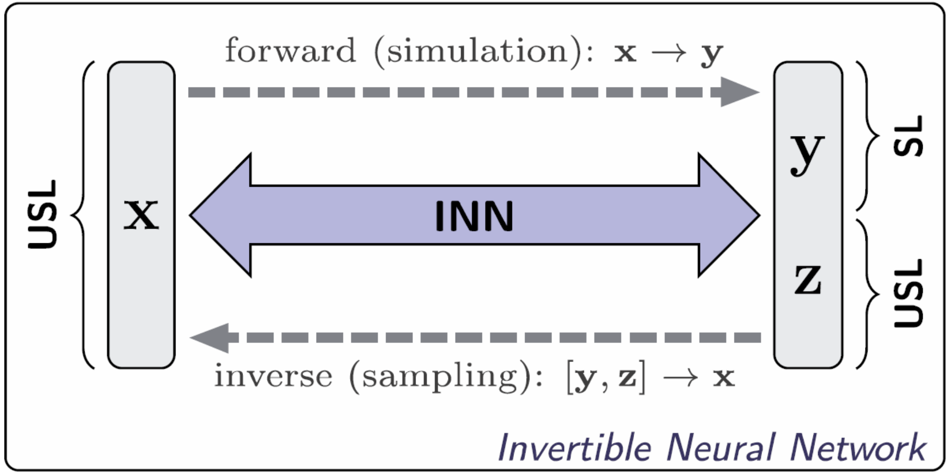

Figure from: Ardizzone, Lynton, Jakob Kruse, Sebastian Wirkert, Daniel Rahner, Eric W. Pellegrini, Ralf S. Klessen, Lena Maier-Hein, Carsten Rother, and Ullrich Köthe. "Analyzing inverse problems with invertible neural networks." in Proc. of *ICLR*, 2019.

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)

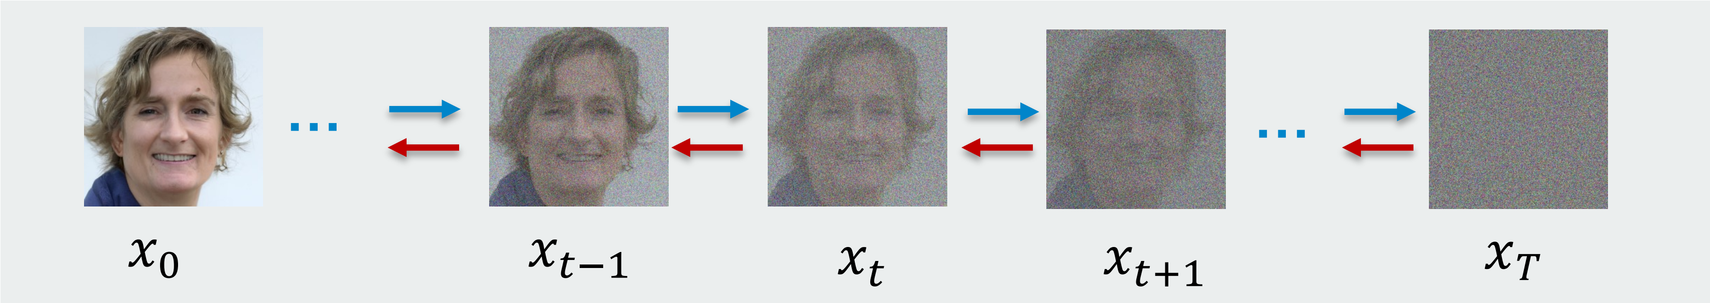

Motivation: Can we use a pretrained "unconditional" diffusion model for inverse problems?

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

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

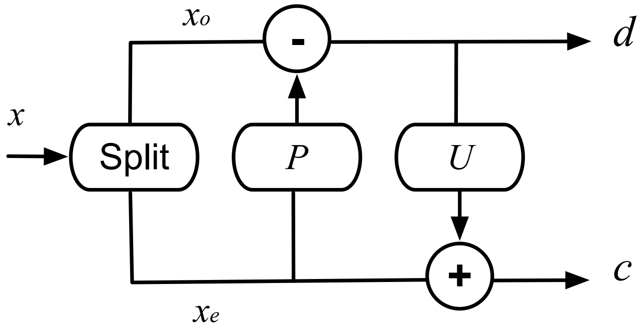

Imperial College London

INN and Diffusion Models for Inverse Problems

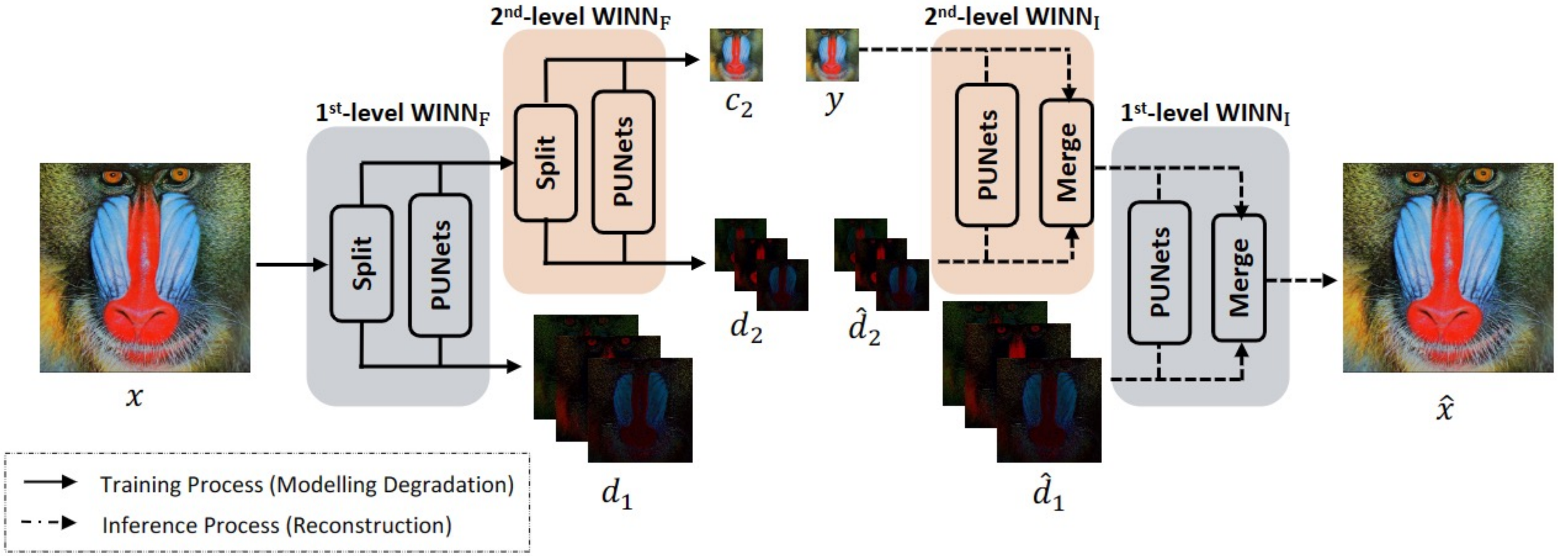

Imperial College London

INN and Diffusion Models for Inverse Problems

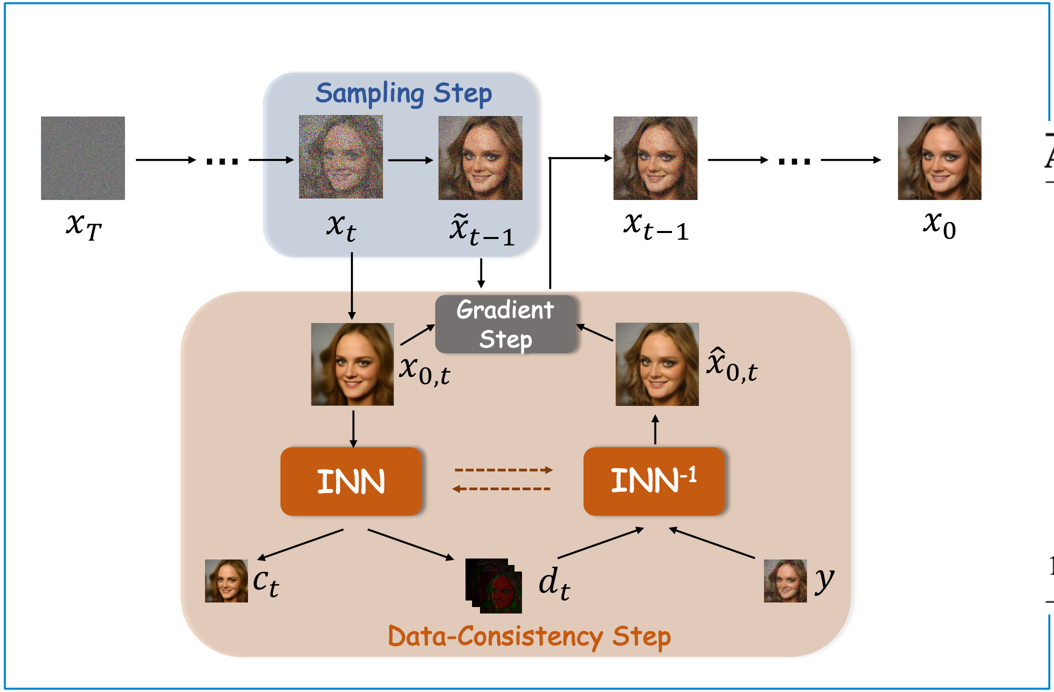

Algorithm 1 INDigo

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

INN and Diffusion Models for Inverse Problems

Ground Truth

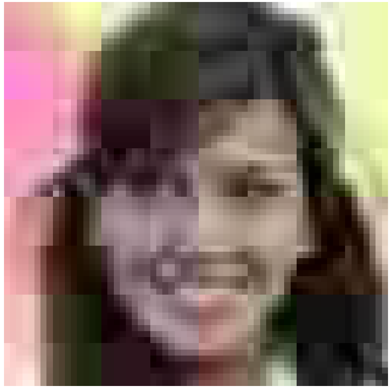

Degraded

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

Input | **Ours** | Ground-truth

INN and Diffusion Models for Inverse Problems

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

Imperial College London

Deep Joint Source Channel Coding (DeepJSCC)

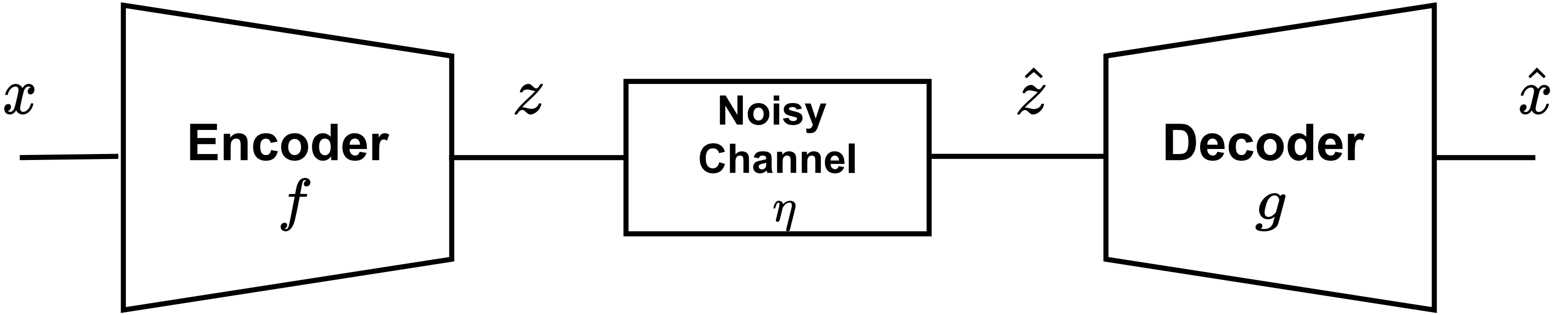

Key Feature of this approach:

- Implement a lossy communication system using a (pre-trained) autoencoder
- The encoder (DNN f) generates the latent variables which are transmitted through the channel
- The decoder (DNN g) receives a corrupted version of the latent variables and reconstruct the original image
- The system has attention modules to handle different levels of noise in the channel
- Works better than separate source and channel coding systems in challenging settings and is more flexible

DeepJSCC ad Inverse Problems

Key Insights:

- Treat the complete system as the degradation process in an imaging system
- Improve quality of reconstruction by solving an inverse problem
- Use (pre-trained) generative models to help with the reconstruction

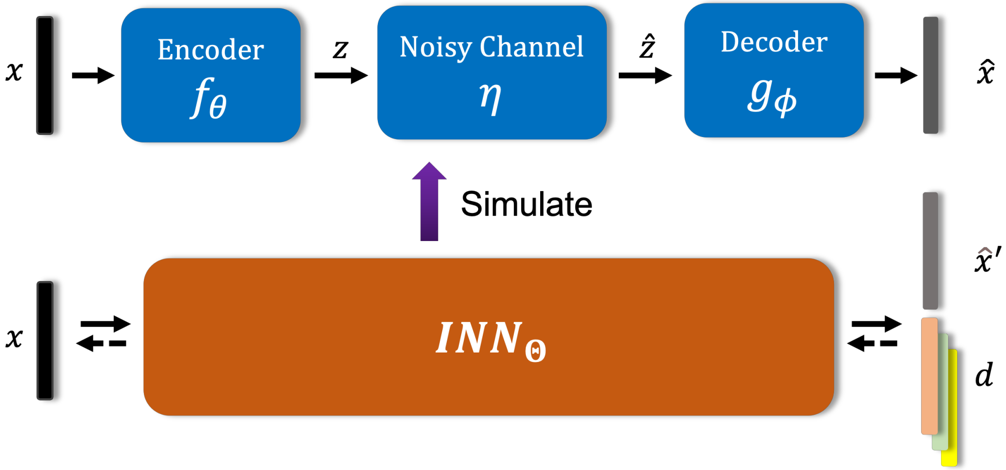

Imperial College London

DeepJSCC ad Inverse Problems: Numerical Results

DeepJSCC InverseJSCC Ours Ground Truth

SNR=1dB, bandwidth compression ratio $\lambda = 0.0013$

J. Chen et al. SING: Semantic image communications using null-space and INN-guided diffusion models arXiv:2503.12484

On going work: AI Hallucination

Imperial College London

A broad range of cases where AI model outputs are incorrect or inconsistent

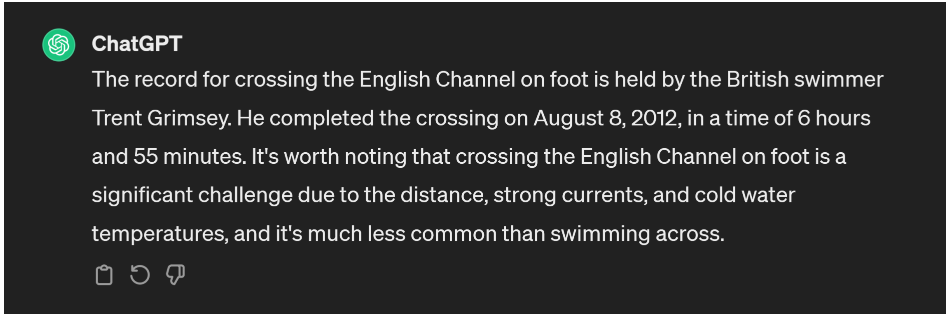

For image restoration problems

Where the restored output appears like a natural image but is semantically different from the ground truth.

Input | Ground truth | Reconstruction

Hallucination Detection

Imperial College London

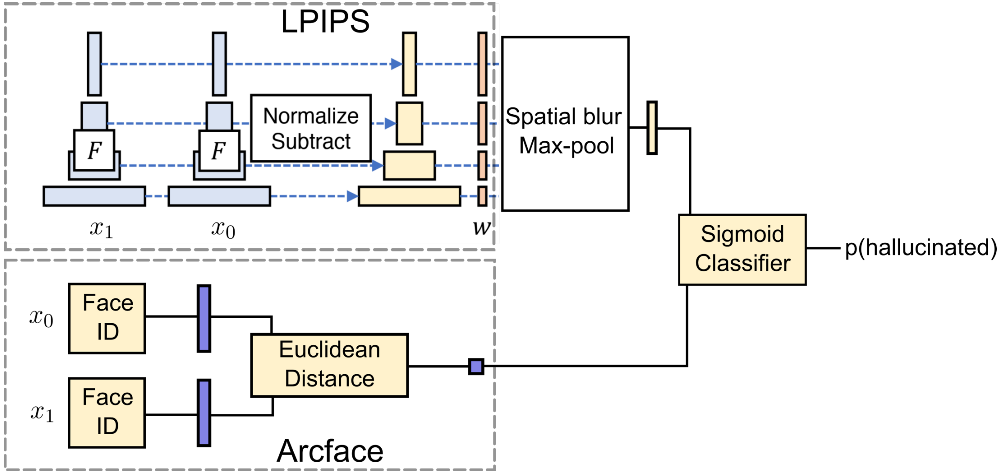

- Use LPIPS perceptual metric and Face ID (Arcface)
- Use max-pooling in place of averaging to detect localised hallucinations.
- Train weights for a single layer sigmoid classifier based on manually labelled degraded image – reconstruction pairs.

Hallucination Detection and Correction

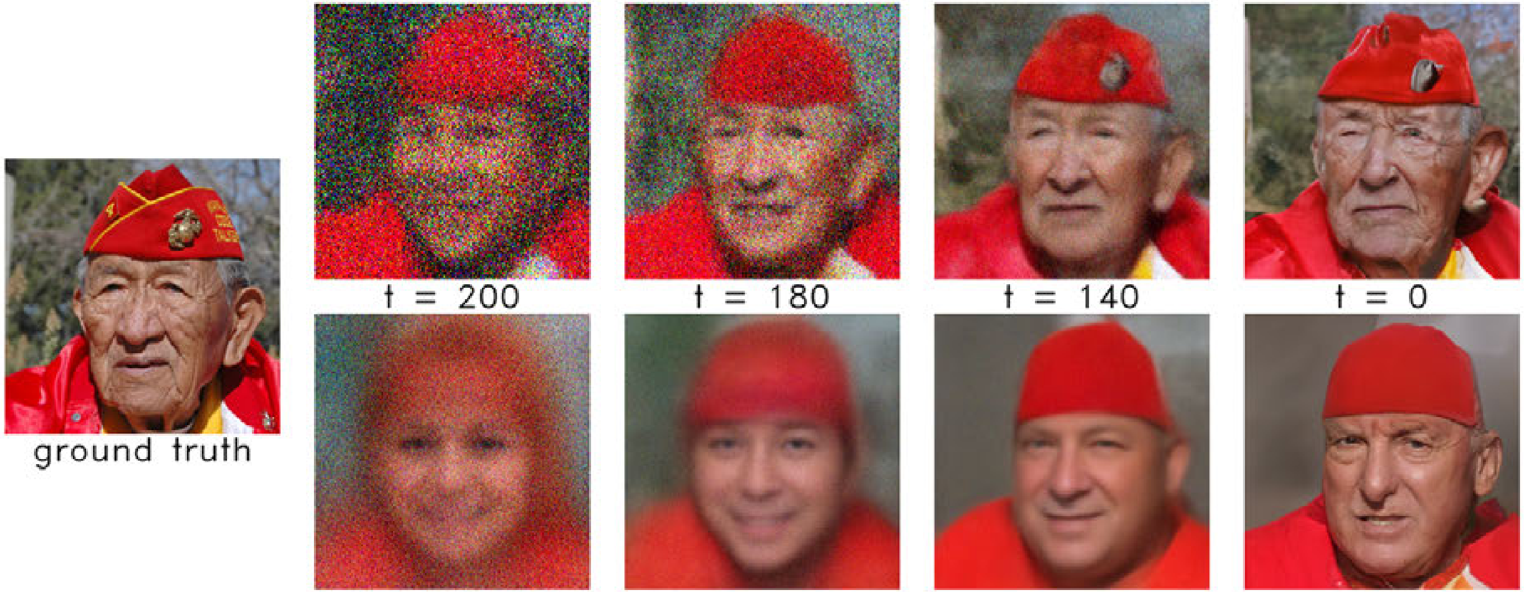

D Siromani, D You, PL Dragotti, “A Hallucination Metric and Correction Scheme for Diffusion-Based Image Restoration”, MLSP 2024

Imperial College London

First Set of Conclusions

- Invertible Neural Networks are an interesting new concept
- Designing INN and combining them with diffusion models leads to more interpretable and simpler architectures
- Good performance
- Potential for further developments (e.g., AI hallucination)

Imperial College London

Part 2: We now get real 😁

Imaging for scientific discovery:

- Very complex imaging workflows
- No ground truth data (but ever increasing size of data output)
- No silver bullet

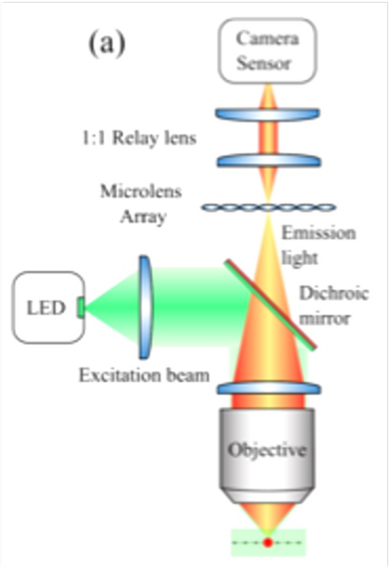

Light field microscopy for neuroscience

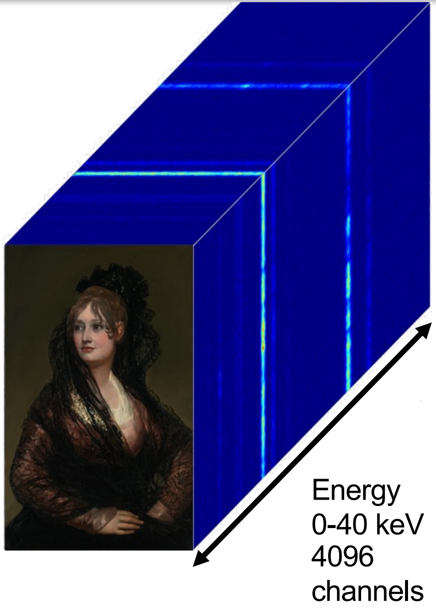

Technical study of Old Masters paintings

Model-Based Deep Learning

Plato: models, priors

Aristotle: data

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

- **Key in inverse problem** is to find the right balance between **data** and **priors**

Sparsity and Deep Unfolding Strategy

Explicit embedding of priors and constraints in deep networks

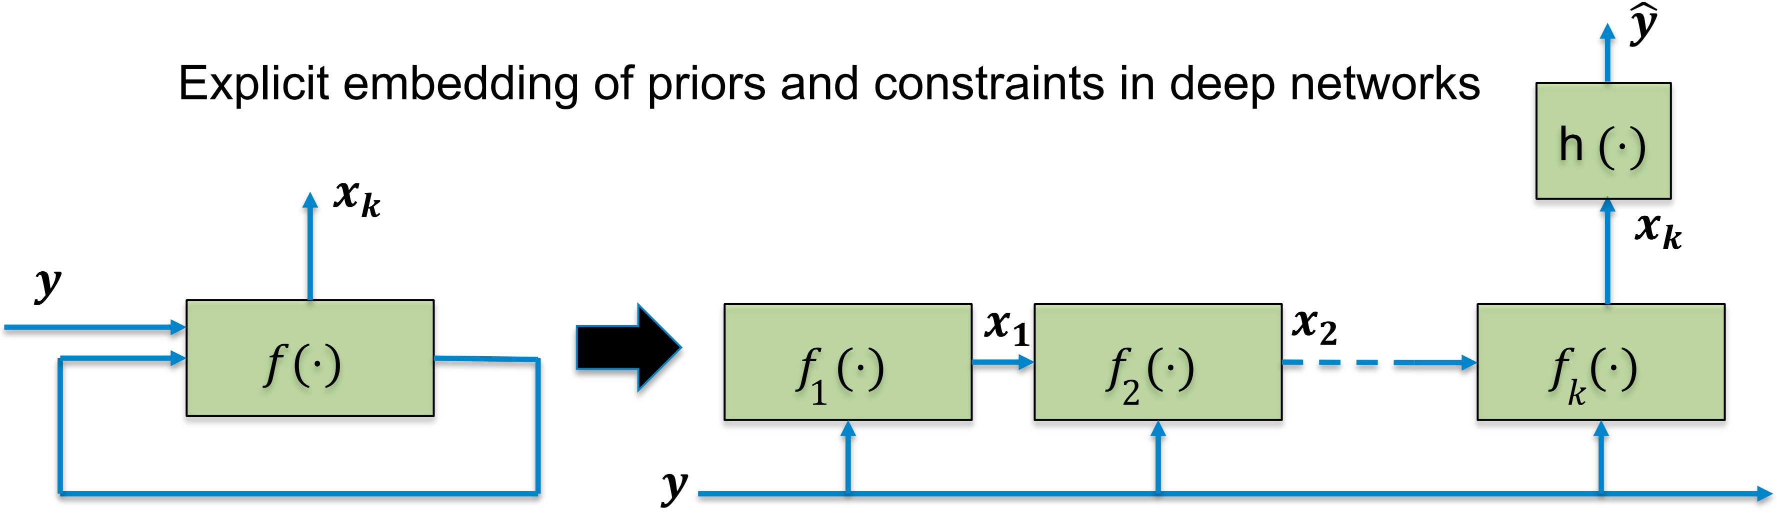

Iterative algorithm with $\boldsymbol{y}$ as input and $\boldsymbol{x}$ as output

Unfolded version of the iterative algorithm with learnable parameters

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

Y. Eldar et al, "Algorithm Unrolling: Interpretable, Efficient Deep Learning for Signal and Image Processing", IEEE Signal Processing Magazine, 2021

Imperial College London

Sparsity as the model for deep unfolding

- The dictionary is usually learned

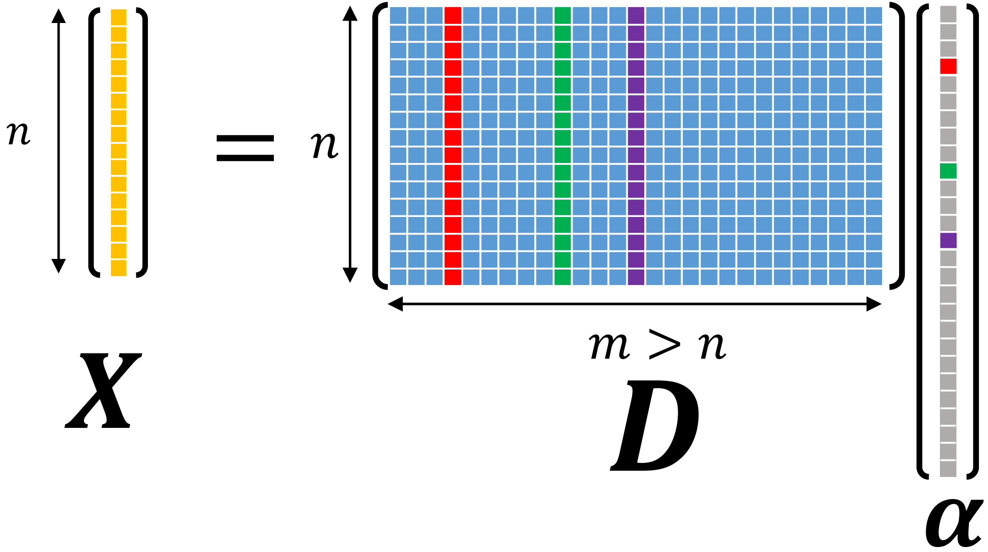

Deep Unfolding Strategy

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

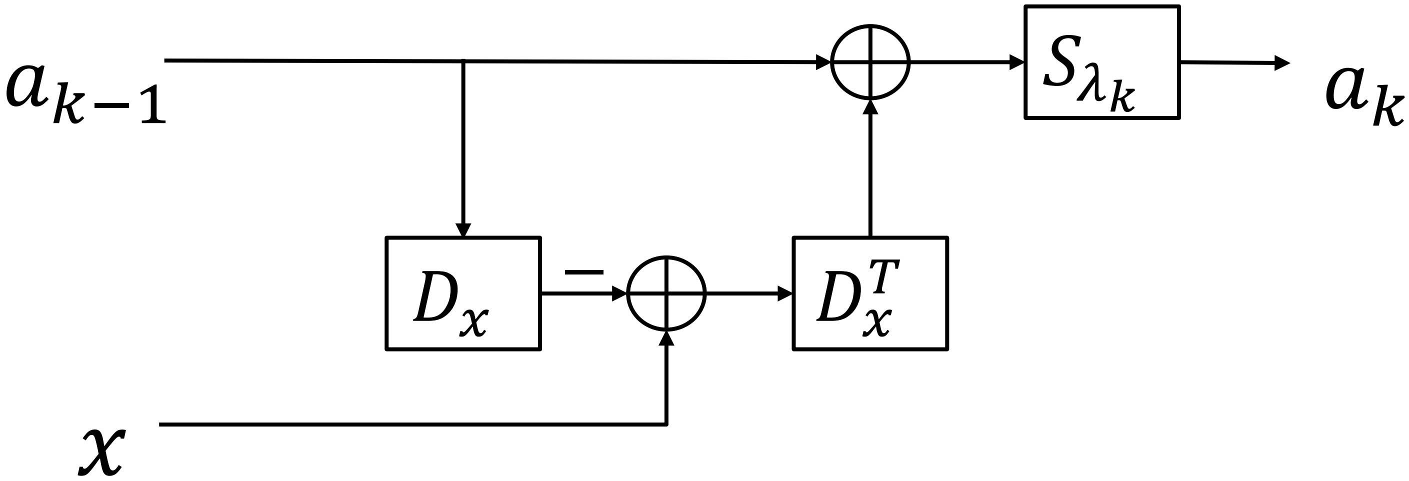

Imperial College London

ISTA network

- Solving by ISTA algorithm through unfolding:

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

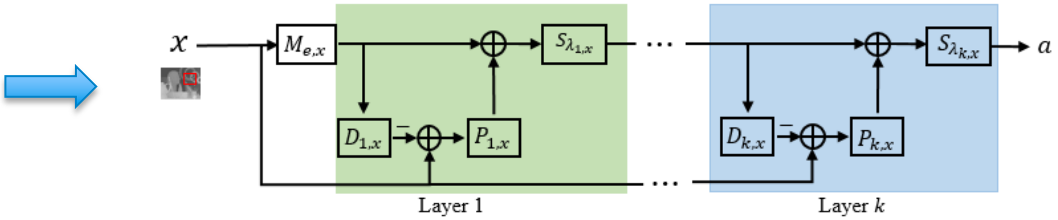

- Gregor Karol and LeCunYann, "Learning fast approximations of sparse coding ", Proceedings of the 27th International Conference on International Conference on Machine Learning, 2010
- Y. Eldar et al, "Algorithm Unrolling: Interpretable, Efficient Deep Learning for Signal and Image Processing", IEEE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

Imperial College London

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
- Multy-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

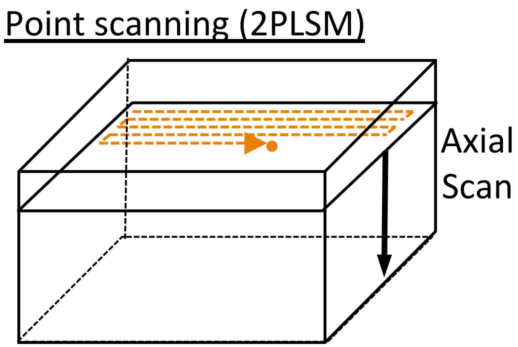

Imperial College London

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

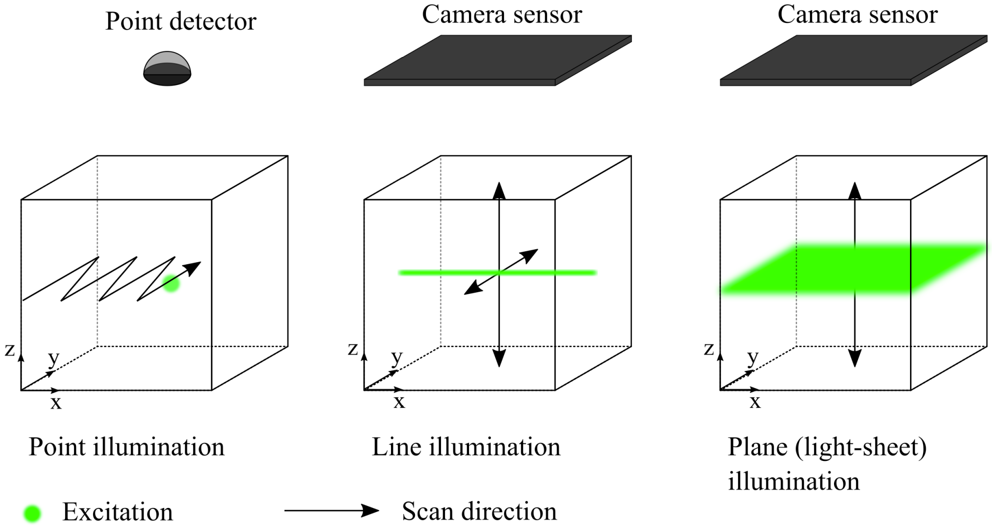

Imperial College London

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

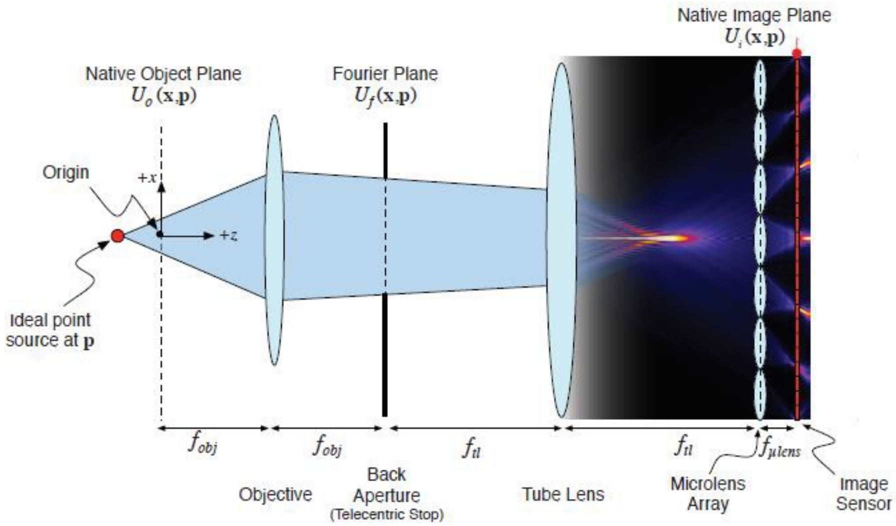

Imperial College London

Light-field Microscopy and EPI

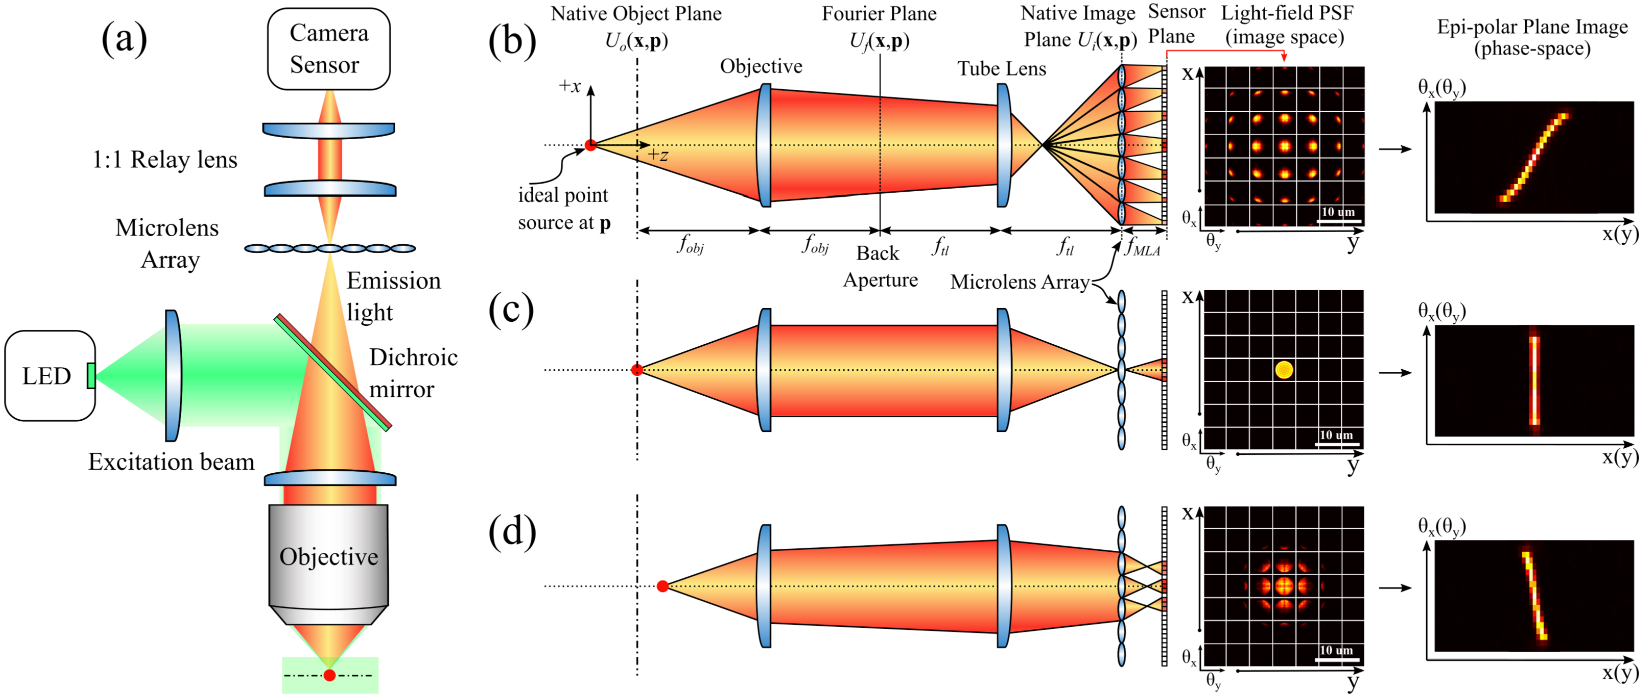

Imperial College London

Light-field Microscopy and Illumination Strategies

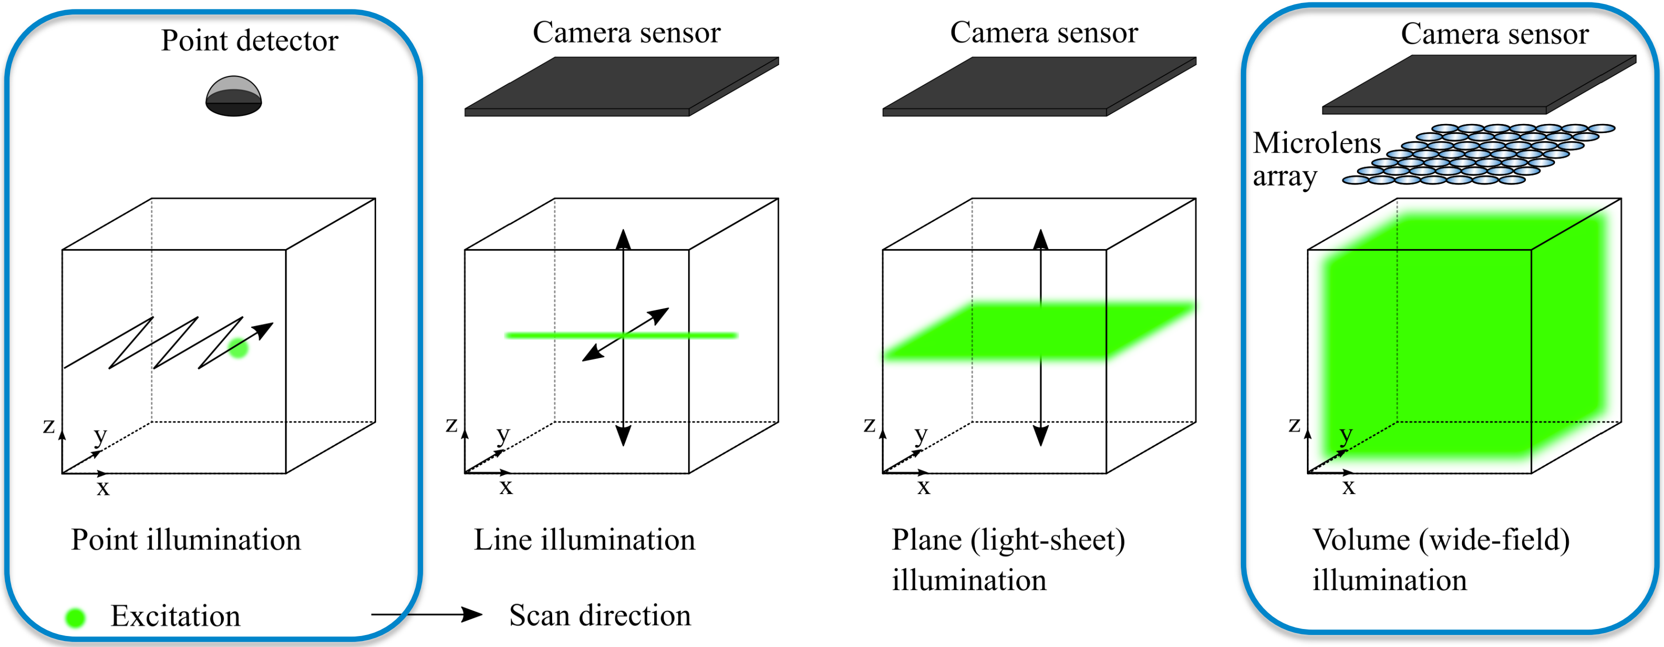

Key insight: use the 2P microscope for high-resolution structural information and the LFM for monitoring the activity of neurons.

Imperial College London

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

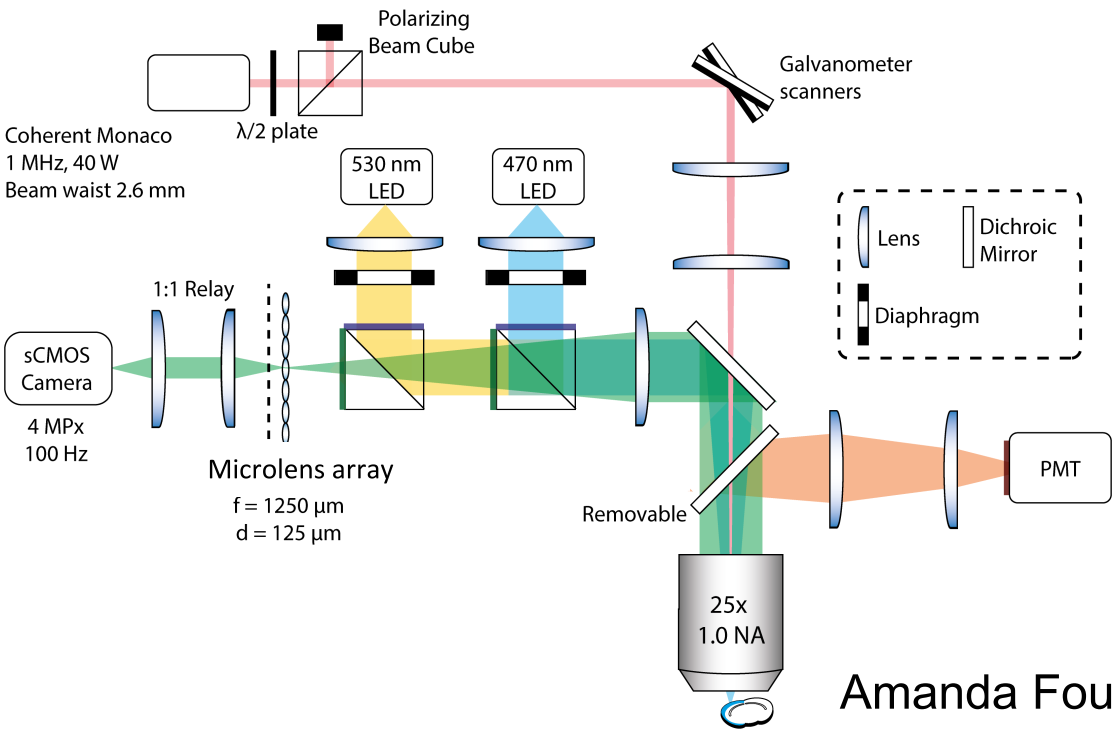

Amanda Foust

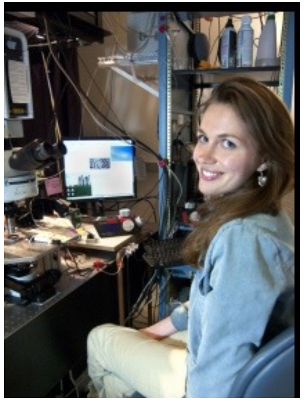

Light-field Microscopy

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

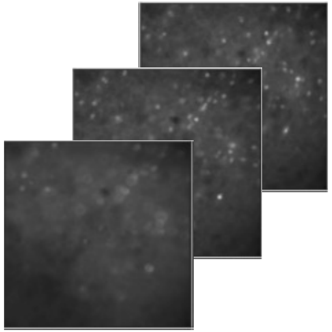

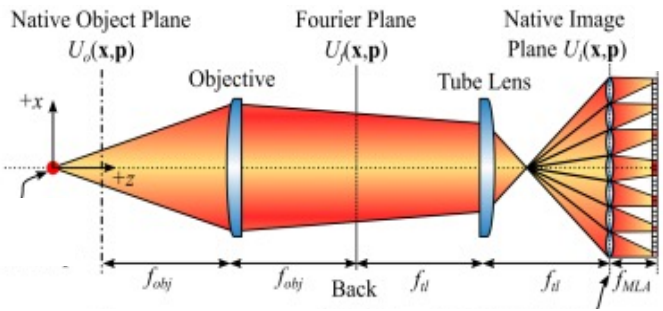

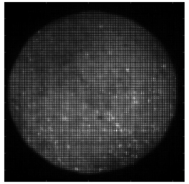

2D Measured LF image

3D Input

Computational
Algorithm

Imperial College London

Volume Reconstruction from 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

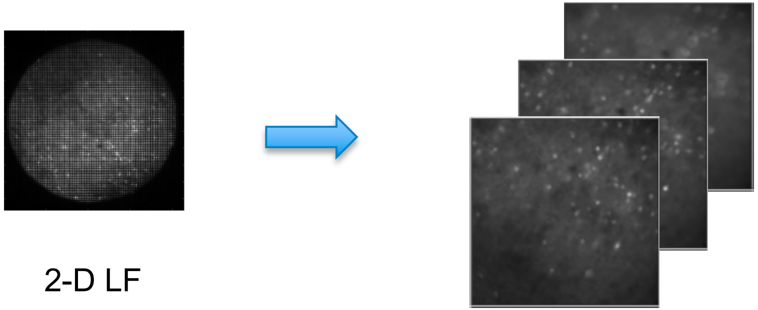

2-D LF

Volume

Imperial College London

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)

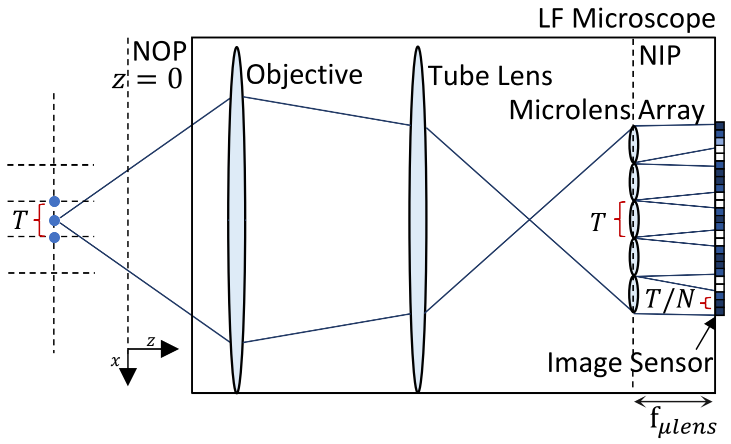

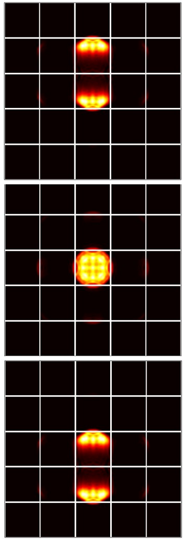

Forward Model

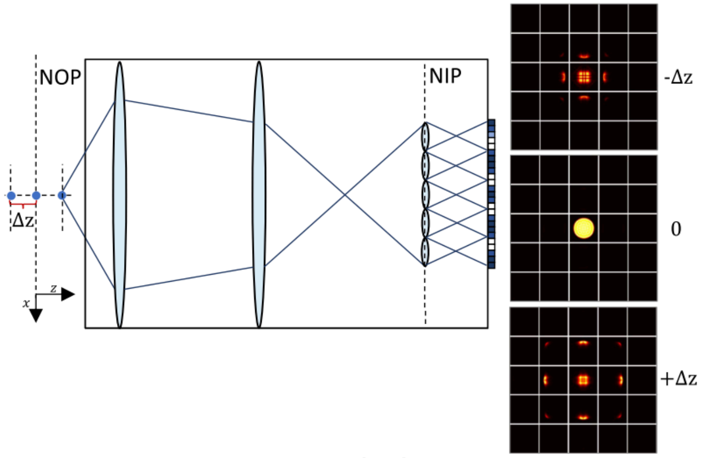

- The discrete LFM system: $\mathbf{y} = \mathbf{H}\mathbf{x}$
- $\mathbf{H} = \sum_{i=1}^{D} \mathbf{H}_i \mathbf{S}_i$

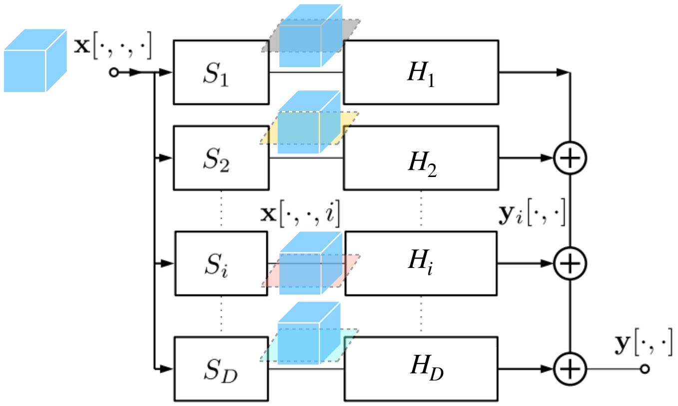

Neural Network for Volume Reconstruction

- 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

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

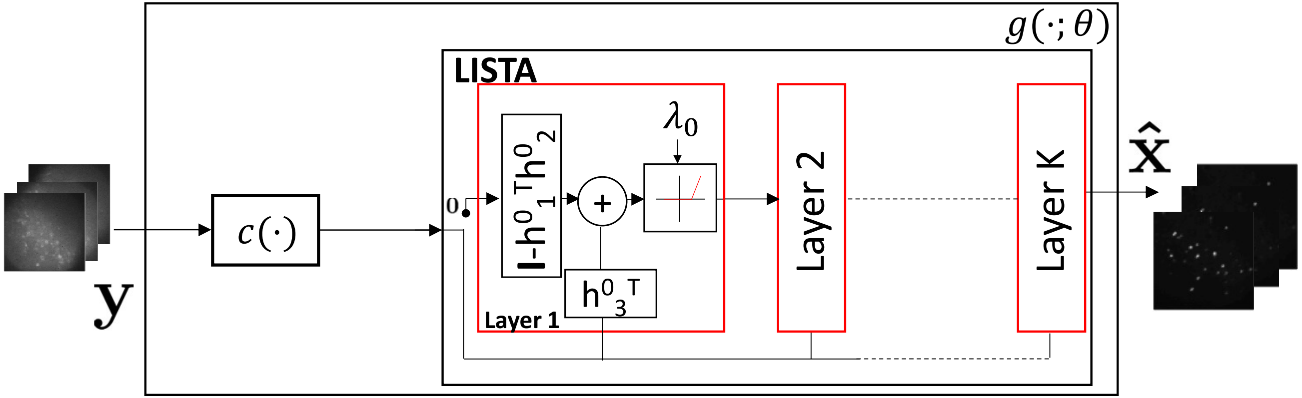

Imperial College London

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

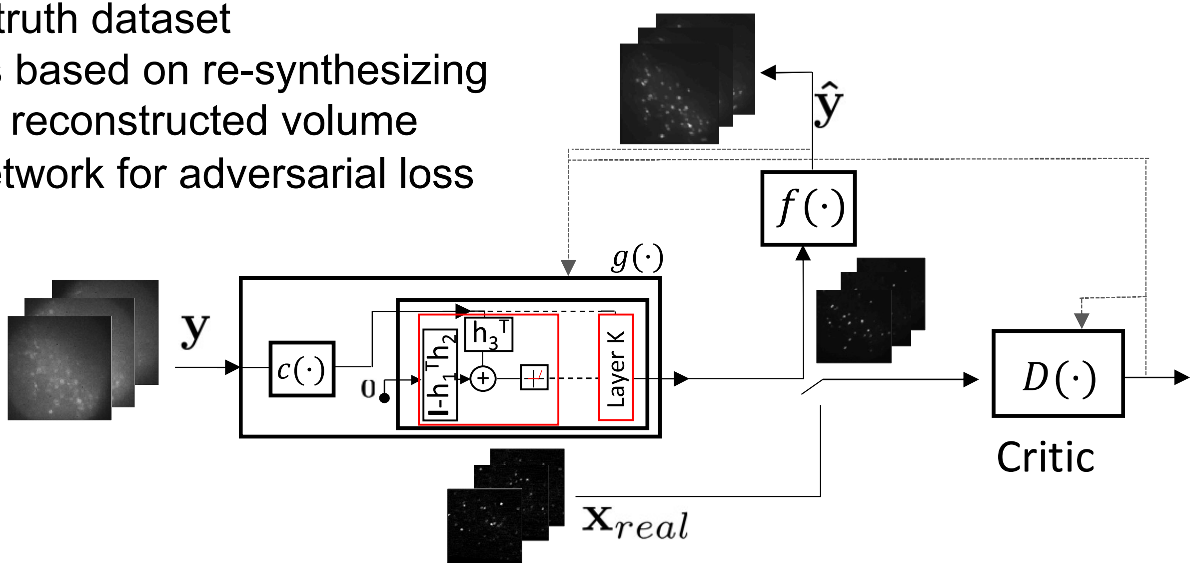

Imperial College London

Results – Structural Data

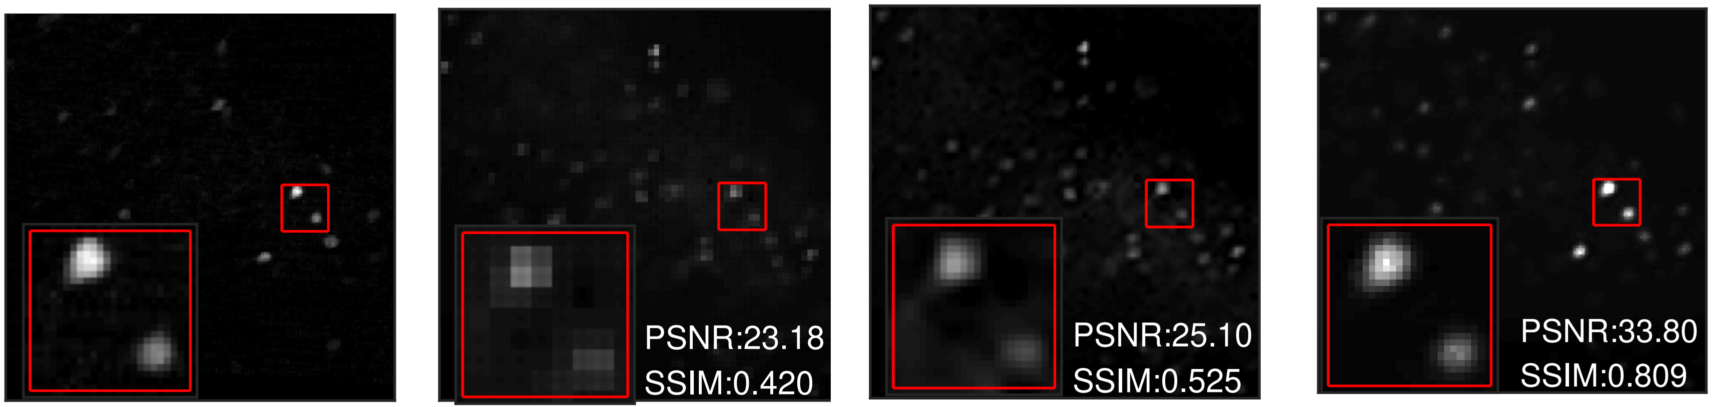

Ground-truth ISRA ADMM

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

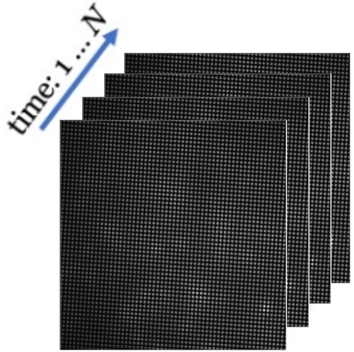

Input F(t):
LF video acquired in
brain slice
cortical layer 2/3

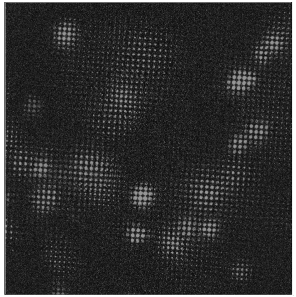

LF activity map:
max(ΔF)/σ(F)

LISTA-based net
Volume Estimation

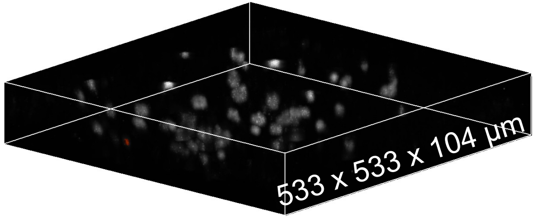

Imperial College London

Fast volumetric jGCaMP8f time-series extraction

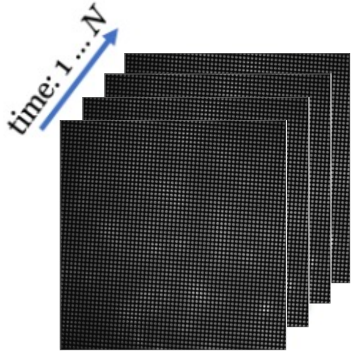

Input F(t):
LF video acquired in
brain slice
cortical layer 2/3

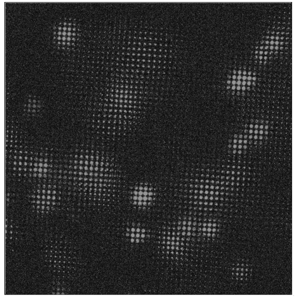

LF activity map:
max(ΔF)/σ(F)

LISTA-based net
Volume Estimation

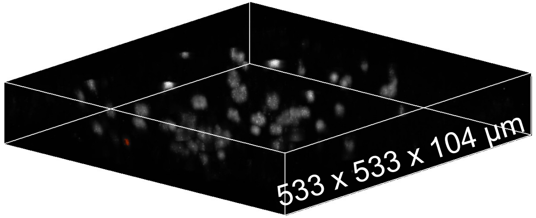

8-iteration Richardson-Lucy Deconvolution

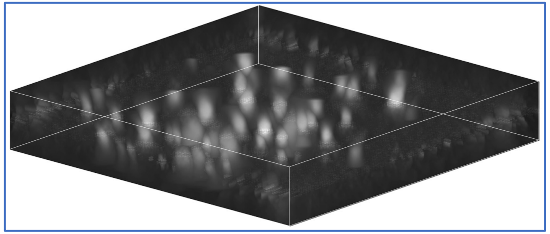

Imperial College London

Fast volumetric jGCaMP8f time-series extraction

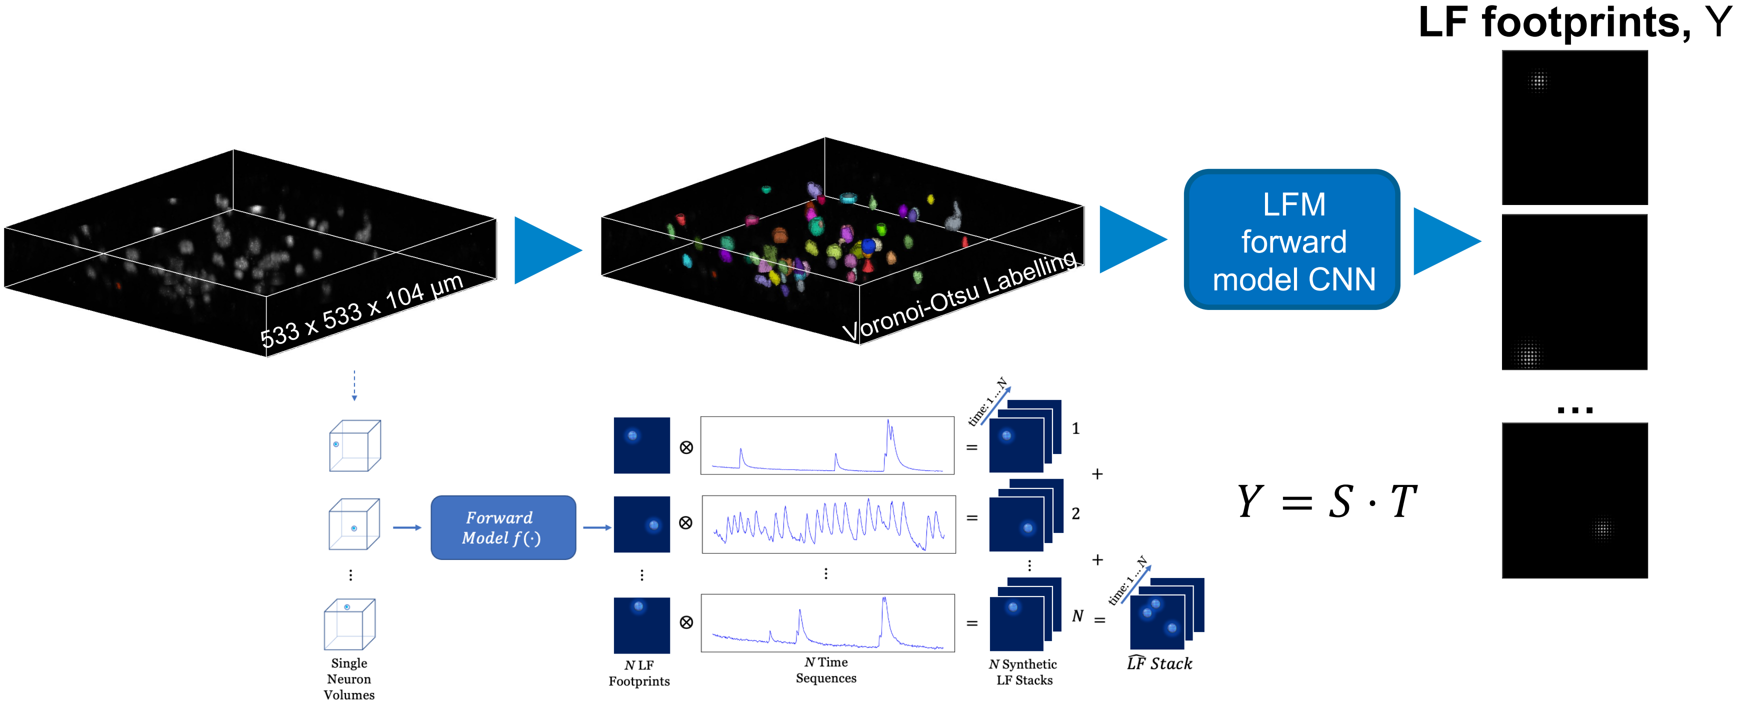

Imperial College London

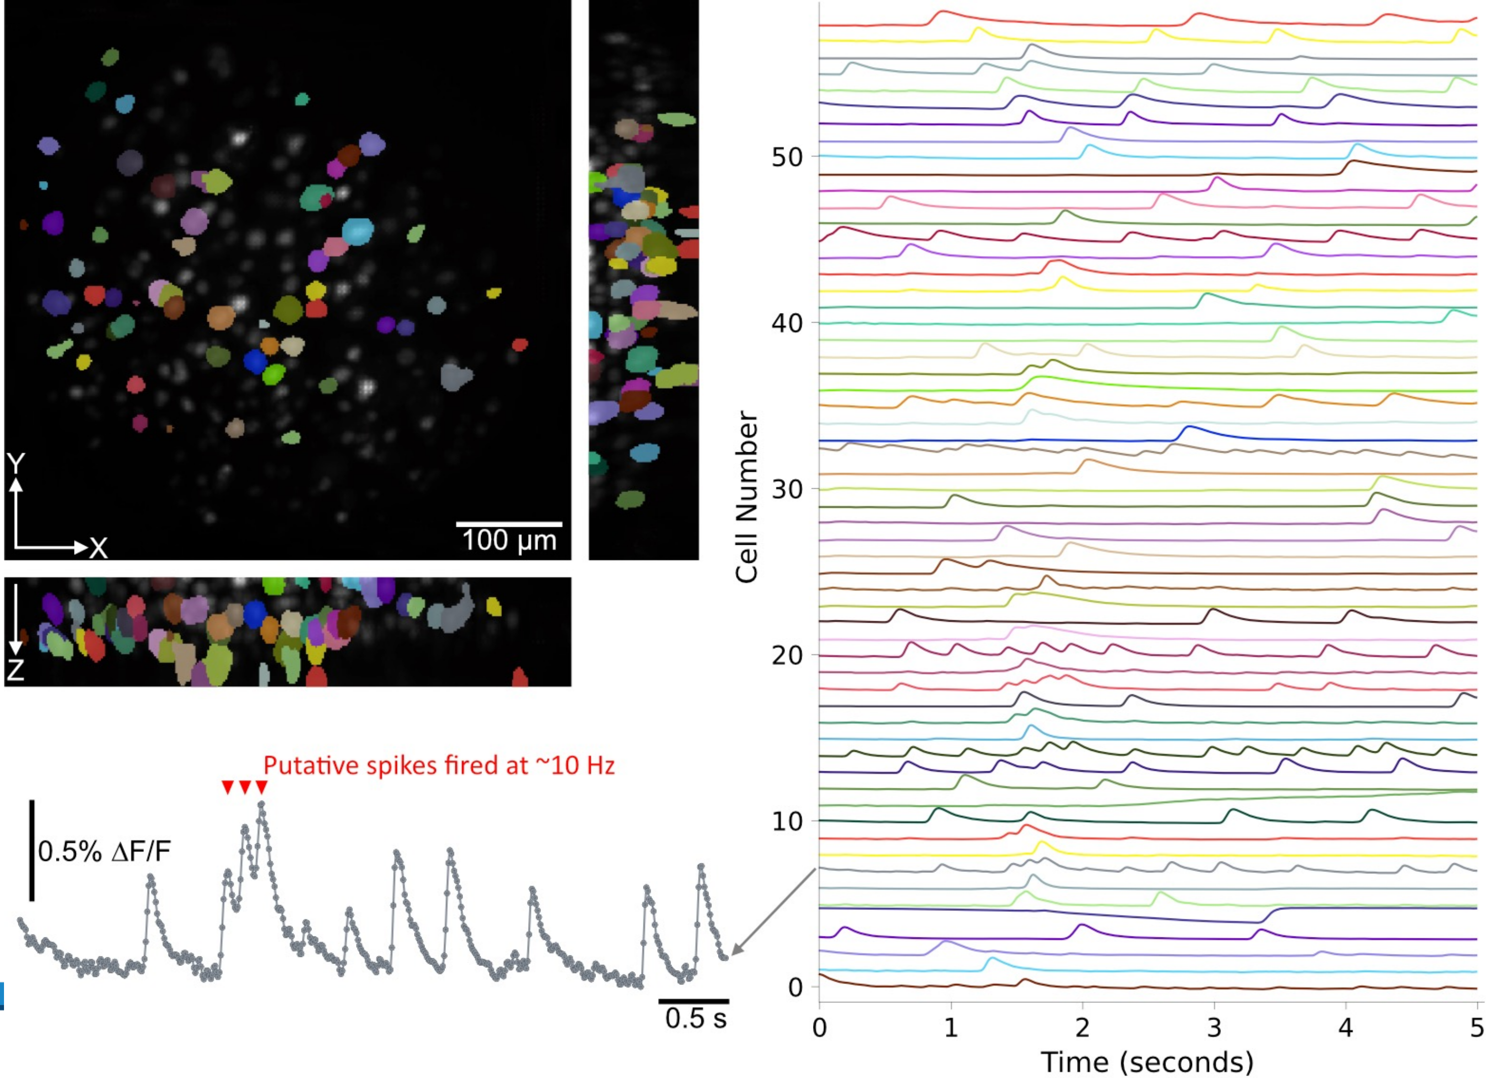

Imperial College London

LISTA-based net decreases crosstalk between neighbouring neurons

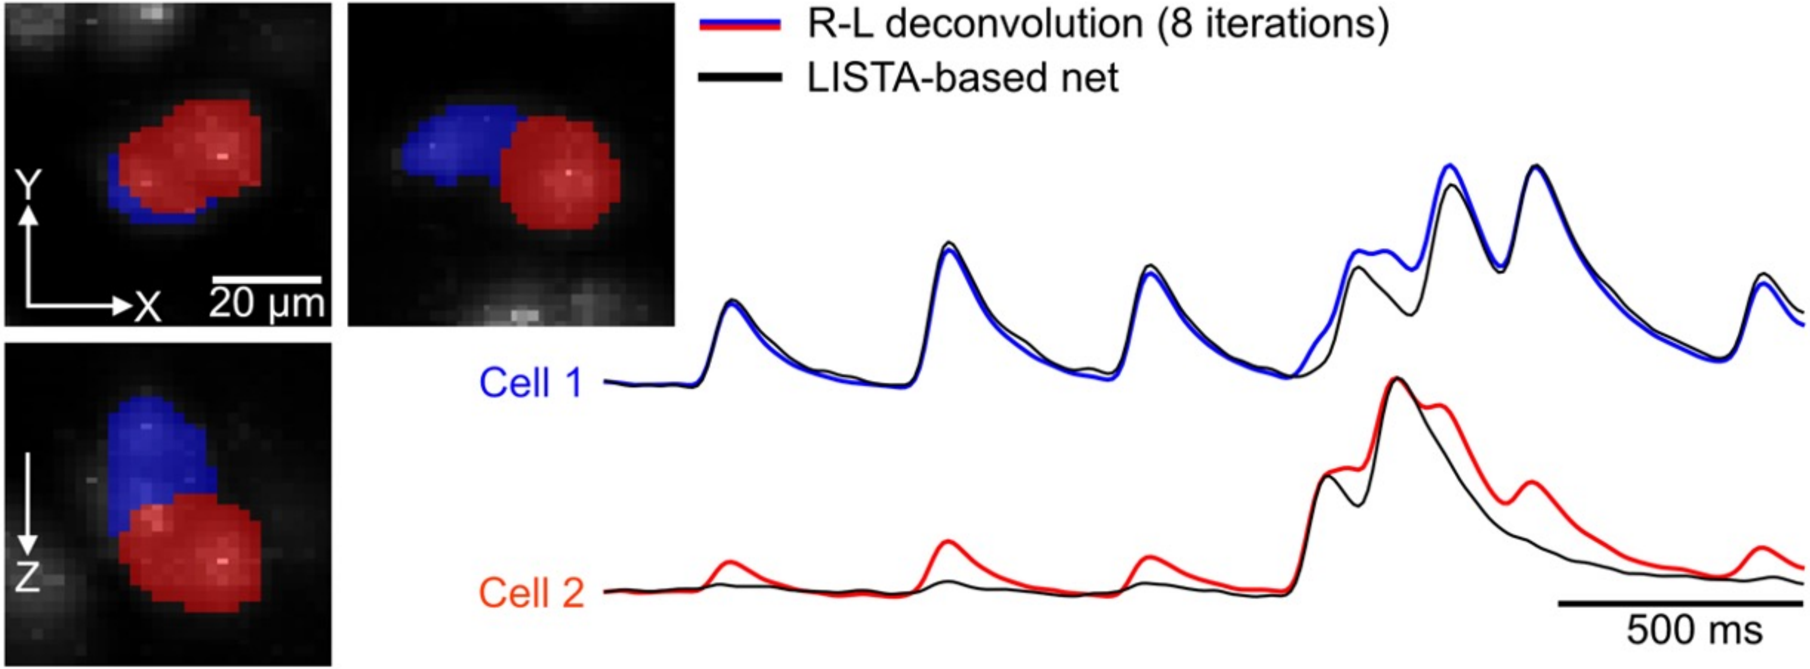

Imperial College London

Art-Investigation

- Technical Studies of Old Masters paintings :
  - For conservation and to understand the creative process of the artist
  - To date paintings and address questions of attribution
  - To present the collection in new ways

Visible

X-ray

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

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

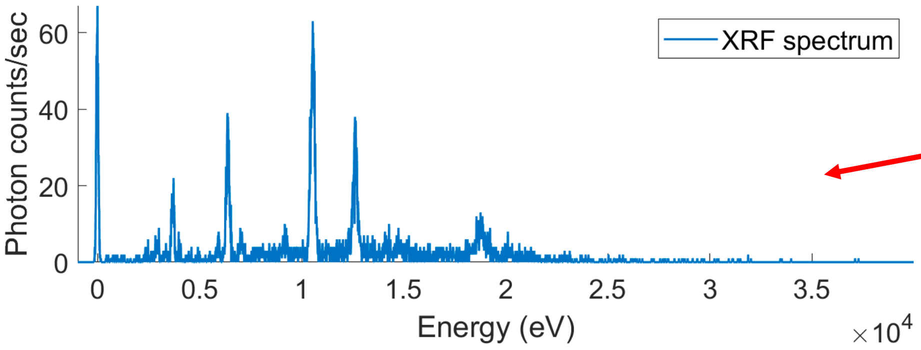

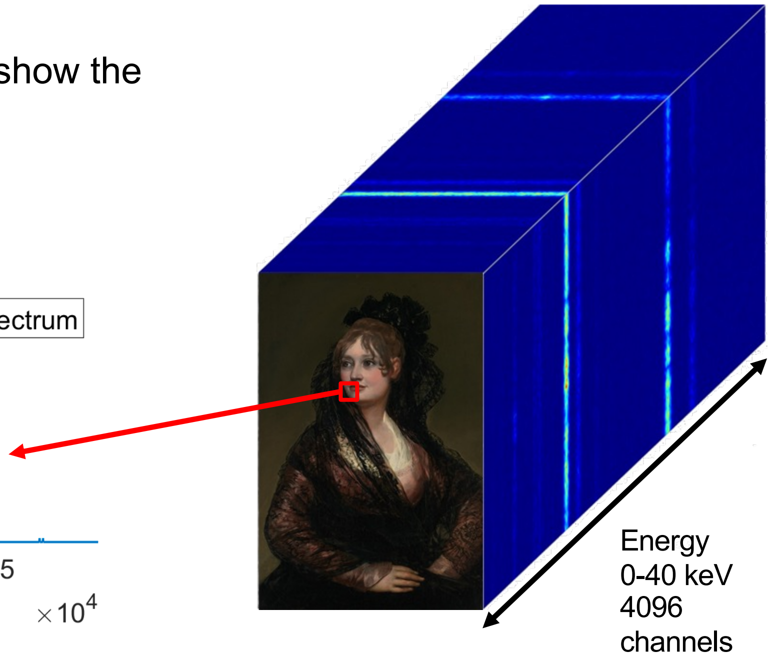

Images © The National Gallery, London

Imperial College London

Extraction of Elemental Maps

Our XRF Deconvolution Algorithm

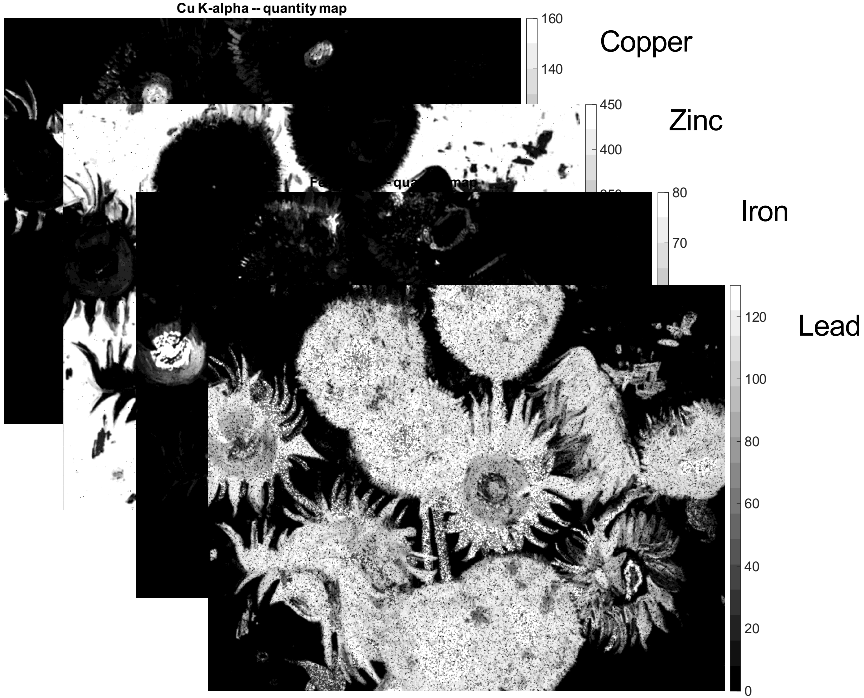

Vincent van Gogh, “Sunflowers (NG3863)”, © The National Gallery, London.

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

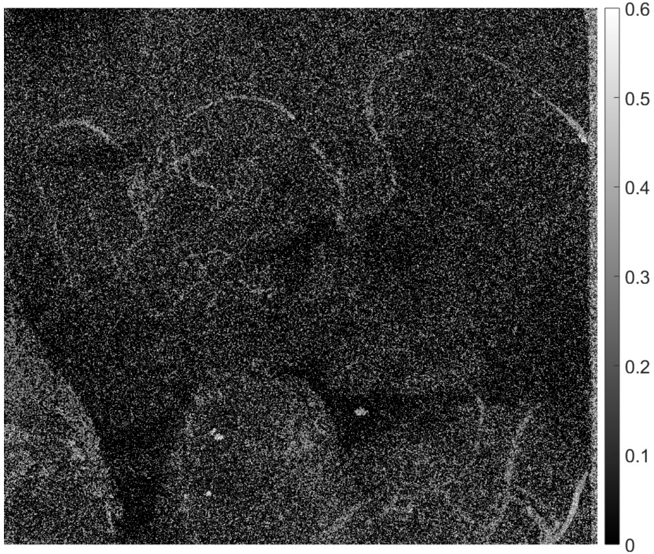

Zn confidence map

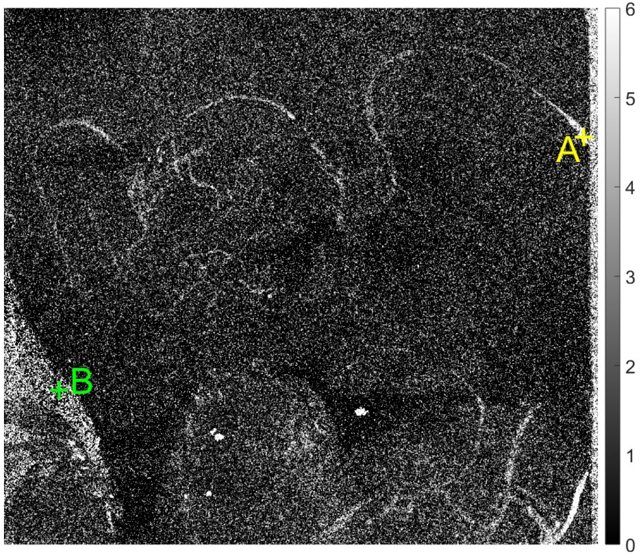

Zn quantity map

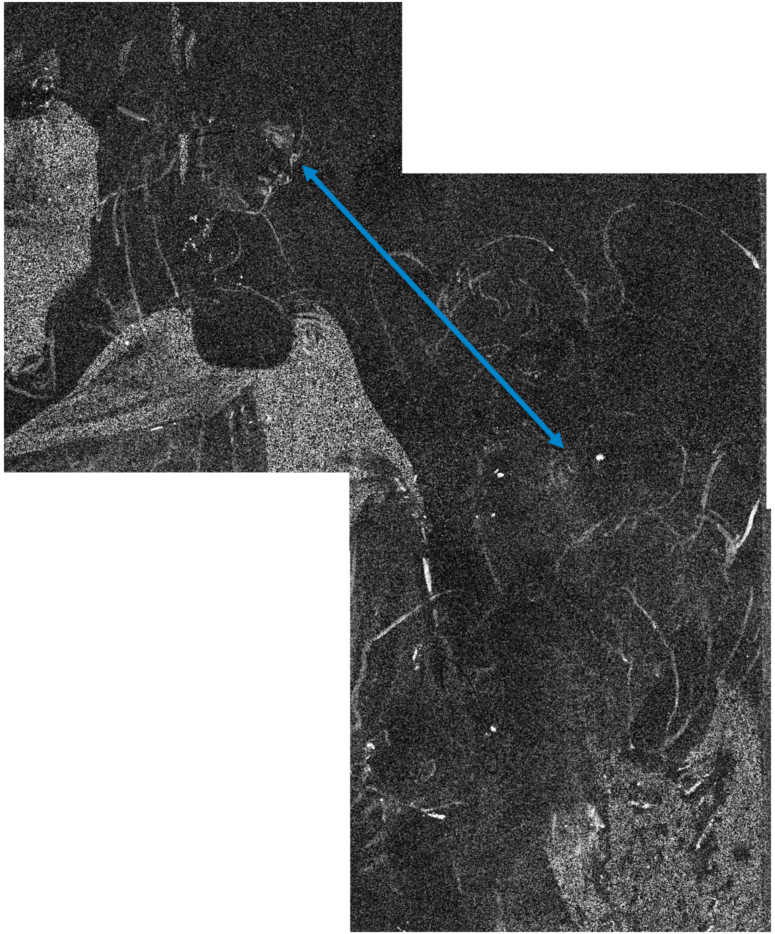

S. Yan, J.-J. Huang, N. Daly, C. Higgitt, and P. L. Dragotti, “When de Prony Met Leonardo: An Automatic Algorithm for Chemical Element Extraction in Macro X-ray Fluorescence Data”, IEEE Transactions on Computational Imaging, vol.7, 2021.

Imperial College London

Conclusions

- In imaging problems:
  - operating at the interface between maths/physics and computation is essential
  - Cross fertilization between model-based approaches and deep learning is fruitful

- My message:
  - embrace inverse imaging problems because they:
    - are fun 🙂
    - and inter-disciplinary

Imperial College London

A special thank to:

Di You

Pingfan Song

Herman Verinaz

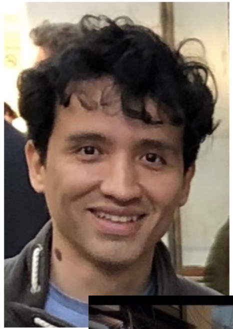

Kate Zhao

Peter Quicke

Carmel Howe

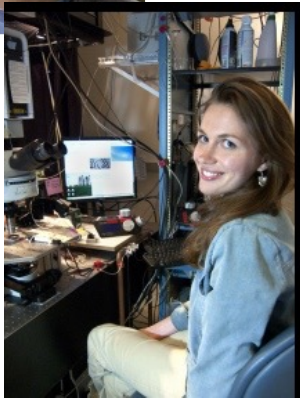

Amanda Foust

Consortium involving: UCL, ICL, Duke and National Gallery

Imperial College London

Thank you!

Imperial College London

Related Publications

- INN and Diffusion Models:
  - J. Huang and P.L. Dragotti, “WINNet: Wavelet-inspired Invertible Network for Image Denoising”, IEEE Transactions on Image Processing, 2022
  - 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
  - J. Chen et al. SING: Semantic image communications using null-space and INN-guided diffusion models arXiv:2503.12484
  - D Siromani, D You, PL Dragotti, “A Hallucination Metric and Correction Scheme for Diffusion-Based Image Restoration”, IEEE 34th International Workshop on Machine Learning for Signal Processing (MLSP), 2024

- 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.