Imperial College London

Model-Based Deep Learning for Inverse Problems in Imaging

Pier Luigi Dragotti, Imperial College London

15 December 2022

Imperial College London

Motivation: A Theory for DL

- Deep Neural Networks achieves state-of-the-art performance in many imaging tasks

- Fundamental questions:
  - Is there a systematic way to design the architecture of a deep neural networks?
  - Is there a systematic way to design interpretable neural networks

- Personal view: in inverse imaging problems interpretable deep neural networks with more predictable performances can only be achieved by combining model-based solvers with learned models.

Imperial College London

Outline

- Invertible Neural Networks and Wavelets
  - What are invertible Neural Networks (INN)?
  - Lifting Scheme and INN
  - Wavelet-like INN for denoising and deblurring
- Computational Imaging
  - Light-field Microscopy for neuroscience
  - Modelling of the image formation process
  - Model-based deep networks for volume reconstruction
- Conclusions

Imperial College London

Joint work with

Junjie Huang

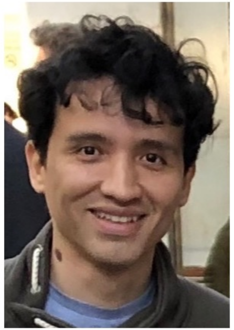

Herman Verinaz

Pingfan Song

Peter Quicke

Carmel Howe

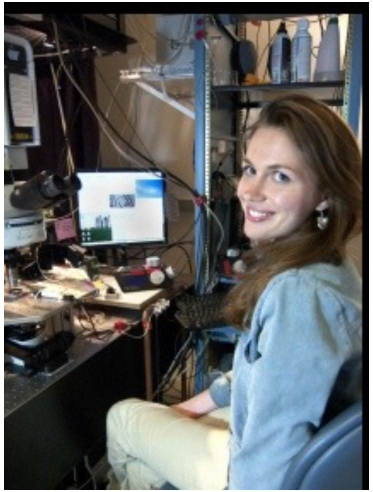

Amanda Foust

What are Invertible Neural Networks?

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

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

- and inverse mapping

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

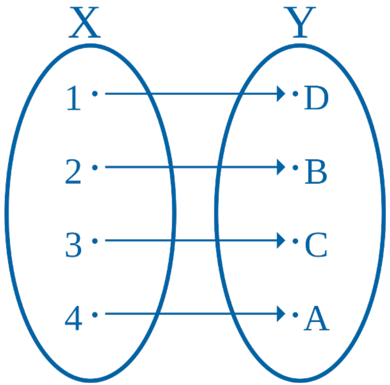

A bijective function (or invertible function)

What are Invertible Neural Networks?

- INNs are bijective function approximators

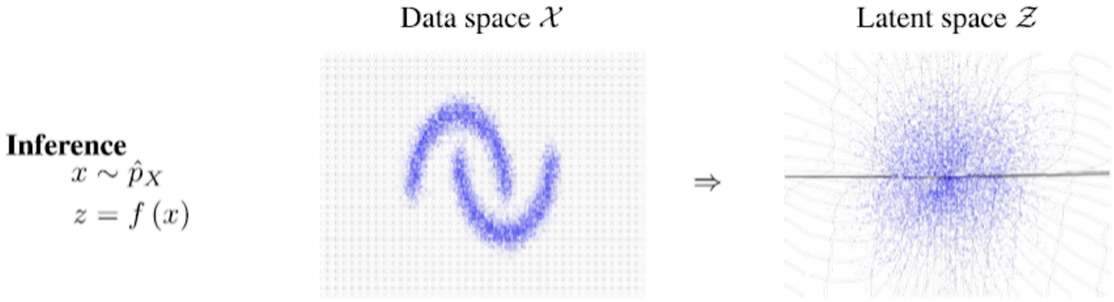

Imperial College London

Why Invertible Neural Networks?

- Generative modeling
  - Tractable Jacobian, allows explicit computation of posterior probability
  - Normalizing flows

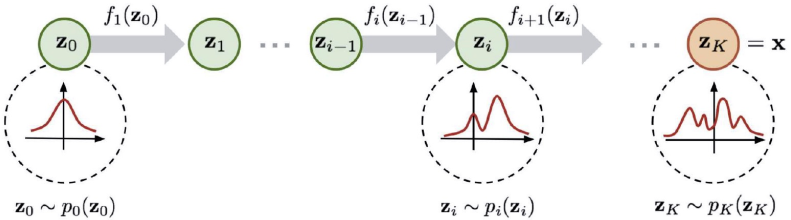

Kingma, Diederik P., and Prafulla Dhariwal. "Glow: Generative Flow with Invertible 1×1 Convolutions." arXiv preprint arXiv:1807.03039 (2018).

How to Achieve Invertibility?

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

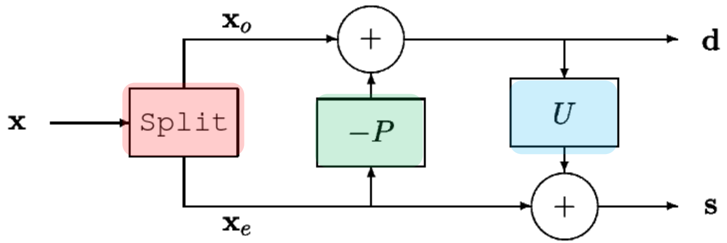

Split → $\begin{cases} d = x_o - P(x_e) \\ s = x_e + U(d) \end{cases}$

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 Invertible Neural Networks

- In the beginning there were Wavelets 😄
- Wavelets provide sparse representations of piecewise smooth images.
- This is why they have been successfully used in many application including denoising

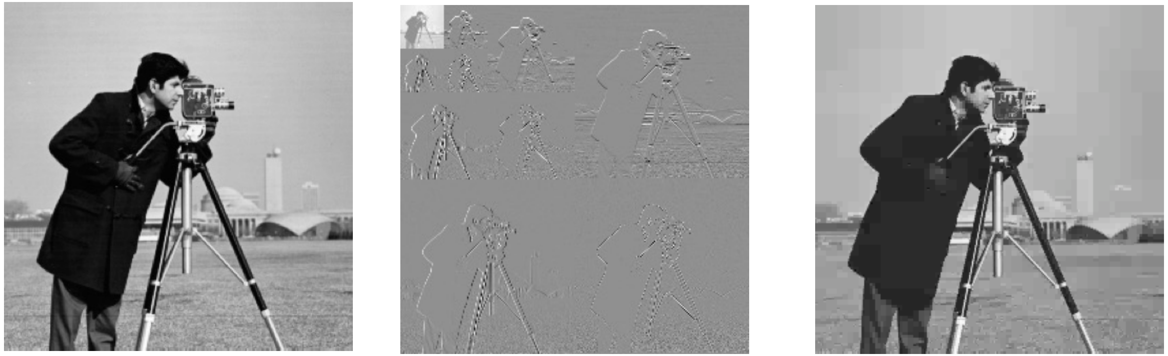

Figure: Cameraman is reconstructed using only 8% of the wavelet coefficients

Imperial College London

Wavelet-based Denoising

- Principles of wavelet denoising:

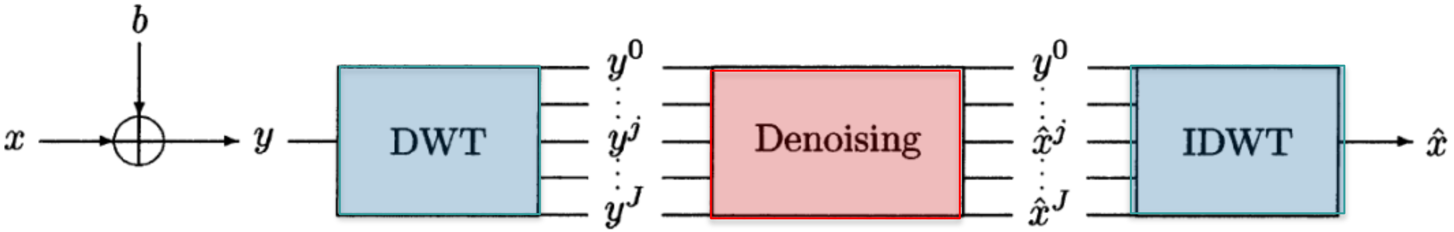

Wavelet transform

- 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

1-D Example

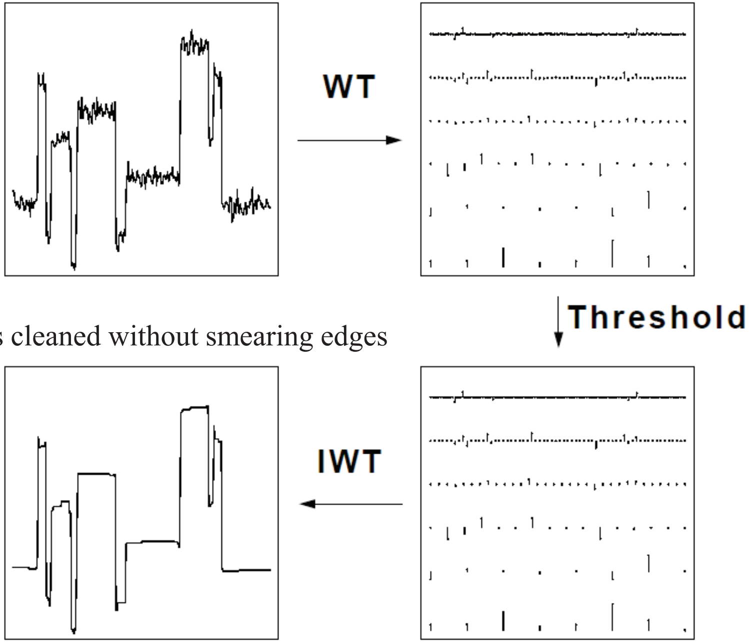

Smooth parts cleaned without smearing edges

Imperial College London

Wavelets for Deconvolution

- Sparsity constraints in the wavelet domain (or in another domain) can also be used as a regularizers for different applications, e.g., deconvolution

- Iterative shrinkage:
  - $\min_{\alpha}(\|\boldsymbol{y} - \boldsymbol{H}\boldsymbol{W}^{-1}\boldsymbol{\alpha}\|^2 + \lambda\|\boldsymbol{\alpha}\|_1)$ where $\boldsymbol{y}$ is the blurred image and $\boldsymbol{x} = \boldsymbol{W}^{-1}\boldsymbol{\alpha}$ is the target image
  - $\boldsymbol{\alpha}_k = S_\lambda(\boldsymbol{\alpha}_{k-1} + \boldsymbol{W}\boldsymbol{H}^{\mathrm{T}}(\boldsymbol{y} - \boldsymbol{H}\boldsymbol{W}^{-1}\boldsymbol{\alpha}_{k-1}))$

Imperial College London

Wavelets and INN

- The wavelet transform can be implemented using the lifting scheme

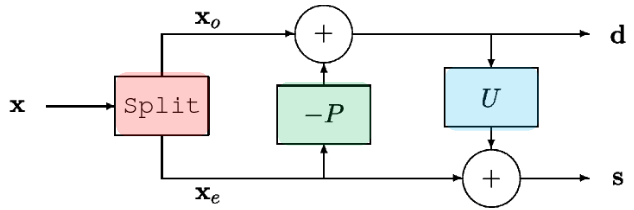

Split → $\begin{cases} d = x_o - P(x_e) \\ s = x_e + U(d) \end{cases}$ $\begin{cases} x_o = d + P(x_e) \\ x_e = s - U(d) \end{cases}$ → Merge

Forward pass 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?

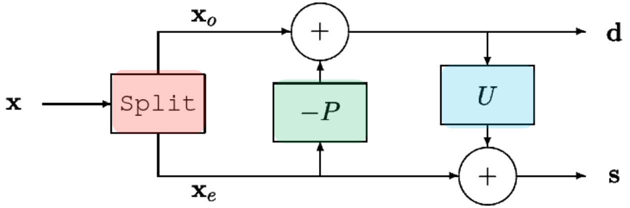

Split → $\begin{cases} d = x_o - P(x_e) \\ s = x_e + U(d) \end{cases}$

Forward pass

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

Backward pass

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

Wavelets and INN

Imperial College London

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

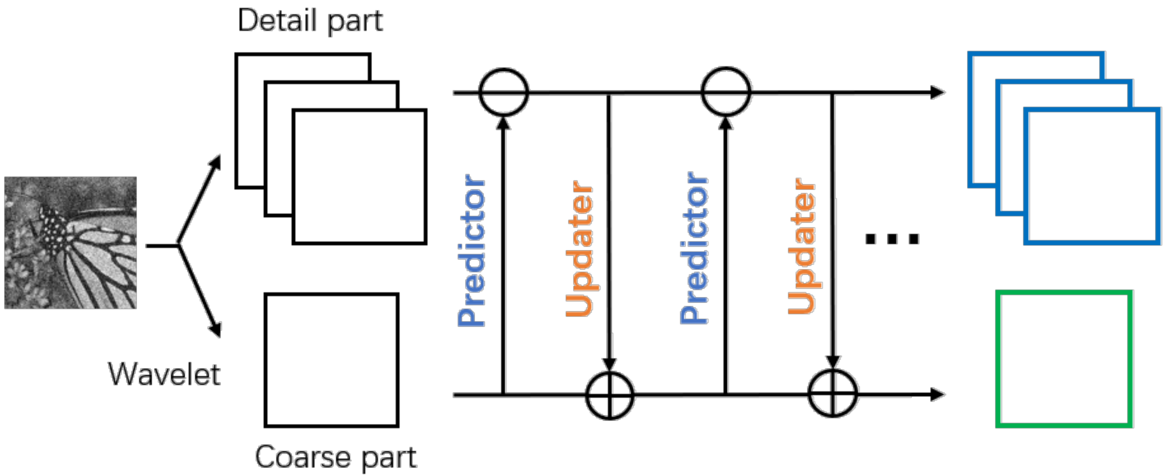

- Approach:
  - convert the P/U operators into two deep networks and learn them
  - Use denoising as the bottleneck to impose sparsity

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

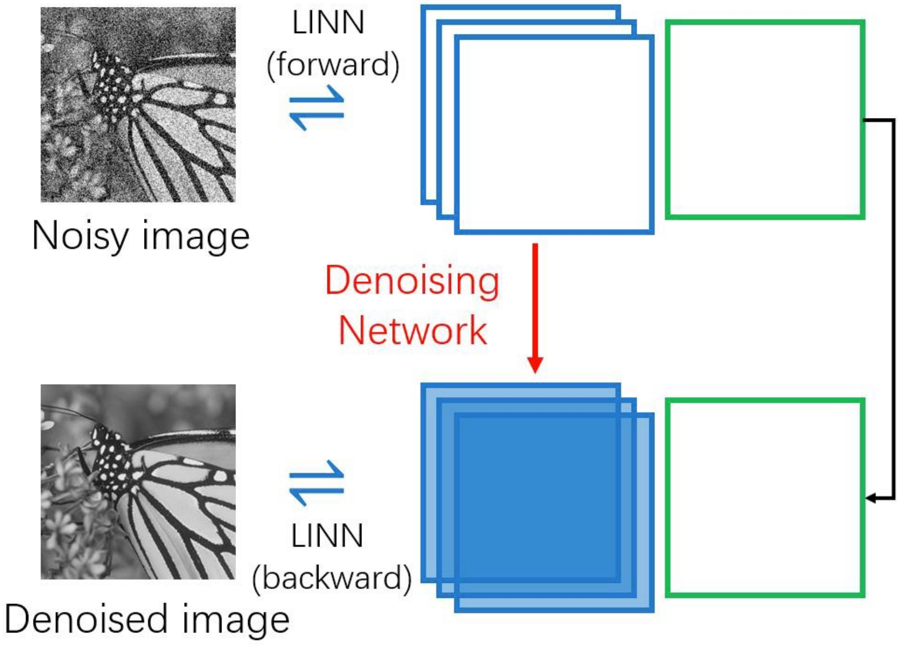

Imperial College London

Signal Decomposition

- Training with noiseless/noisy pairs leads to a sparsifying transform
- Each piece of the network is interpretable
- As for wavelets, we can now use the INN for e.g., denoising or deconvolution

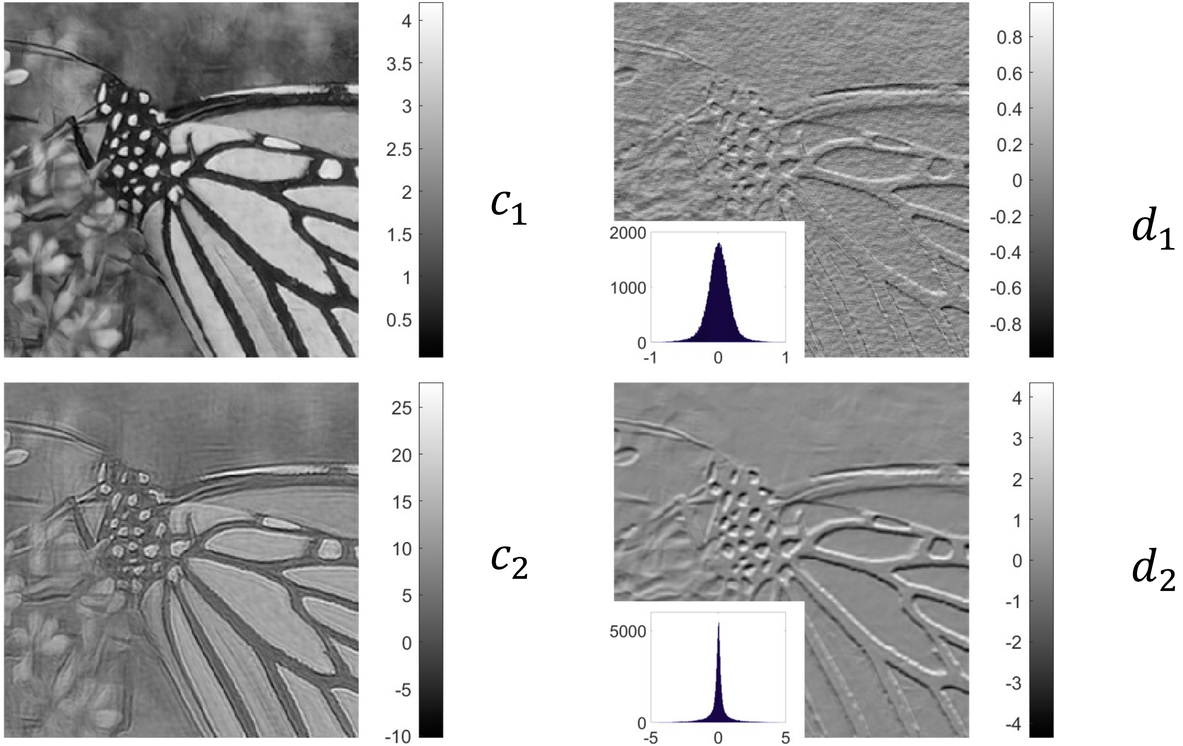

Imperial College London

Denoising - Overall Method

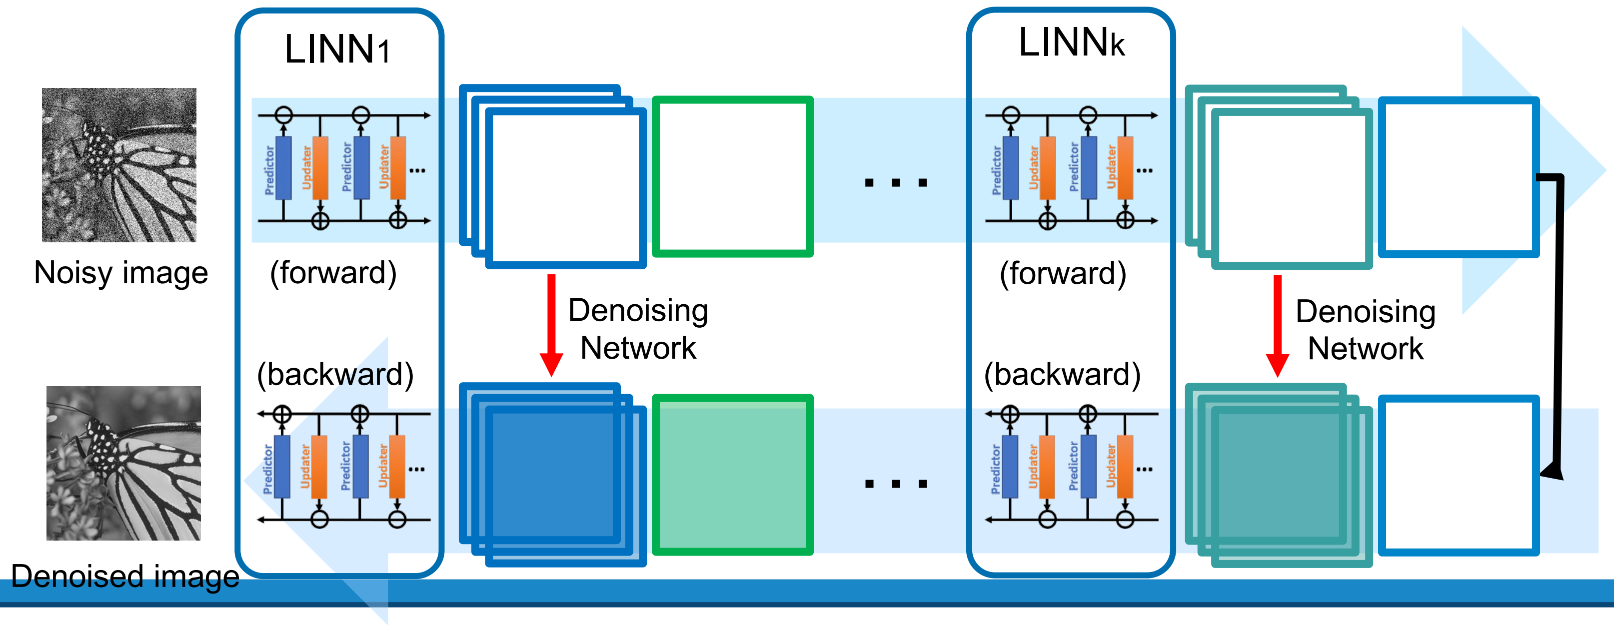

Denoising

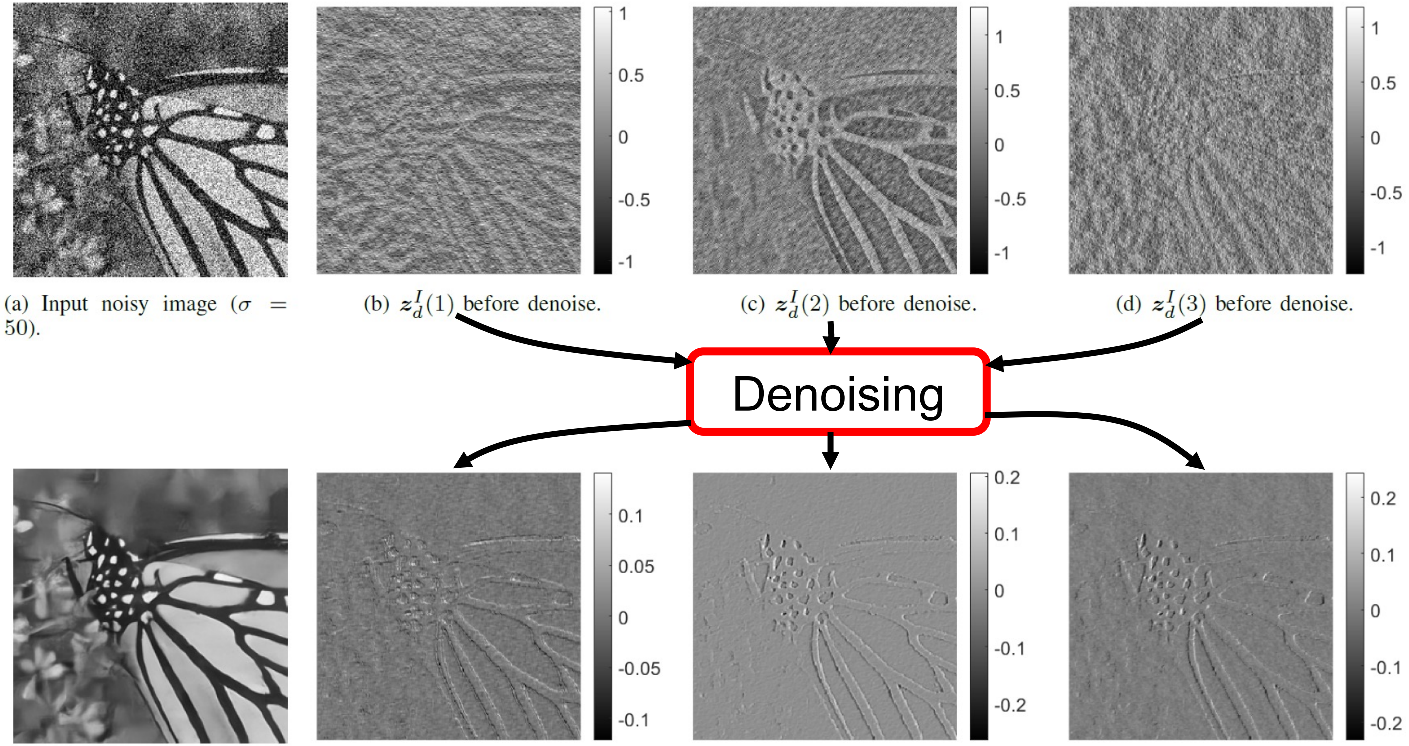

(a) Input noisy image ($\sigma = 50$).

(b) $z_d^I(1)$ before denoise.

(c) $z_d^I(2)$ before denoise.

(d) $z_d^I(3)$ before denoise.

(e) The denoised image (PSNR=26.75 dB).

(f) $z_d^I(1)$ after denoise.

(g) $z_d^I(2)$ after denoise.

(h) $z_d^I(3)$ after denoise.

Imperial College London

Results

Denoising:

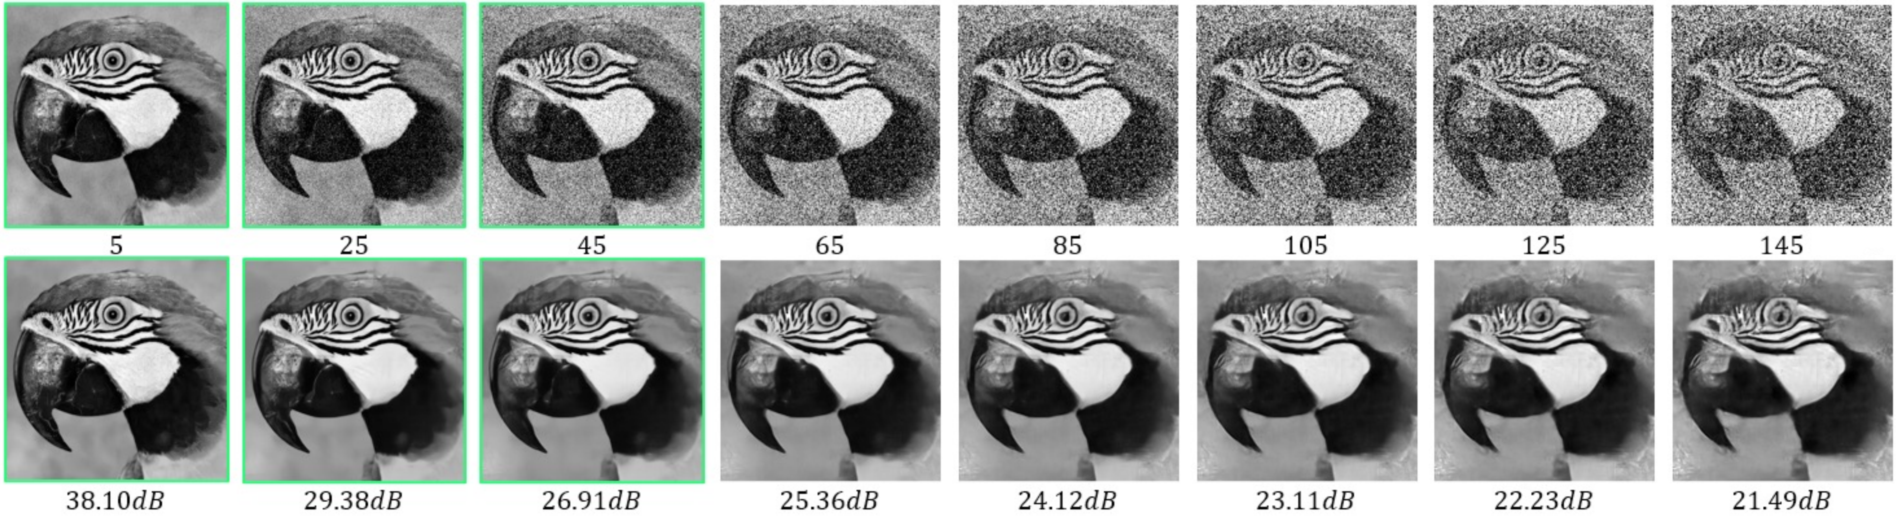

Image Deblurring

Results

Deconvolution:

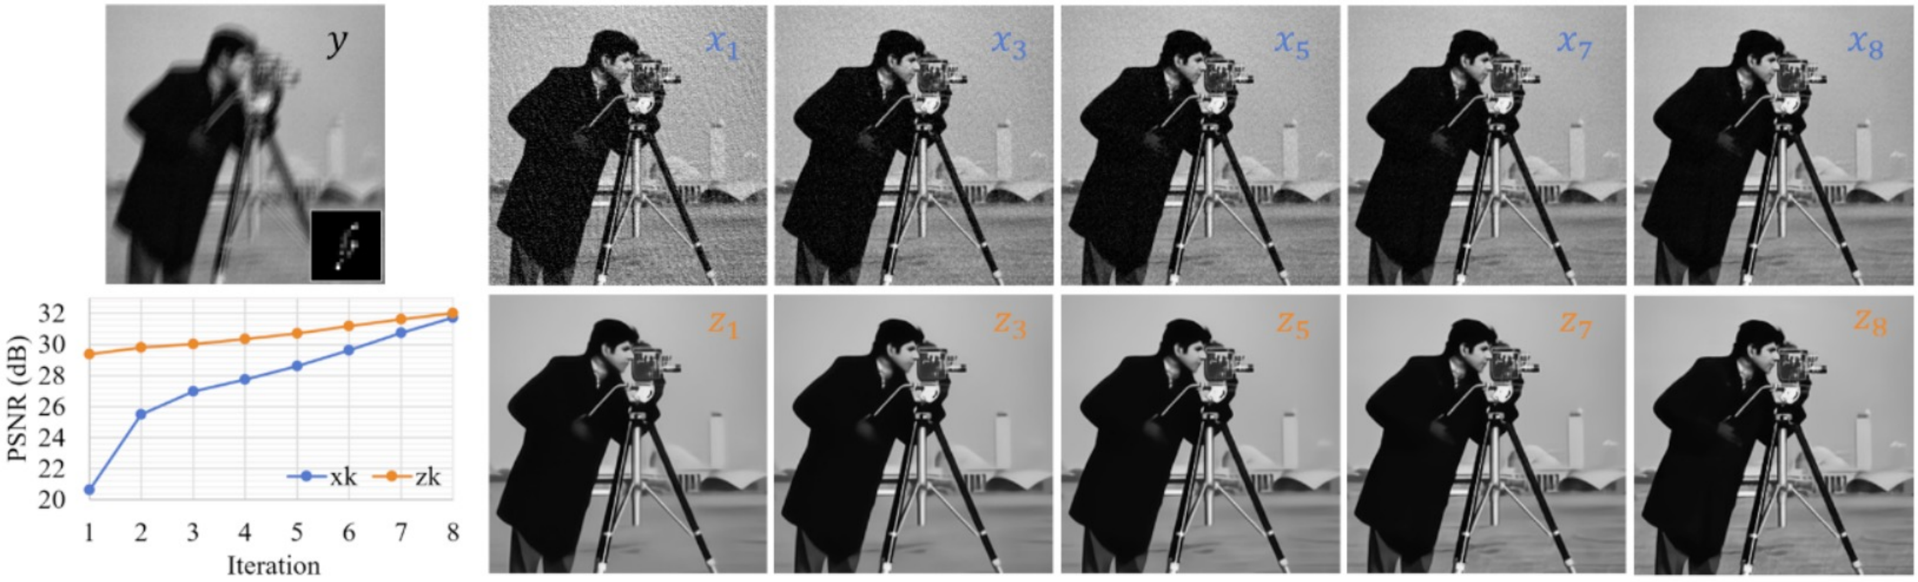

First Set of Conclusions

- Invertible Neural Networks is an interesting new concept
- Designing INN using wavelets/lifting leads to a more interpretable network
- Good generalization ability

Imperial College London

Computational Imaging

- The revolution in sensing, with the emergence of many new sensing and imaging techniques, offers the possibility of gaining unprecedented access to the physical world
- In order to fully exploit these advances, it is necessary to rethink imaging as an integrated sensing and inference model
- Integration of physical and learned models is key

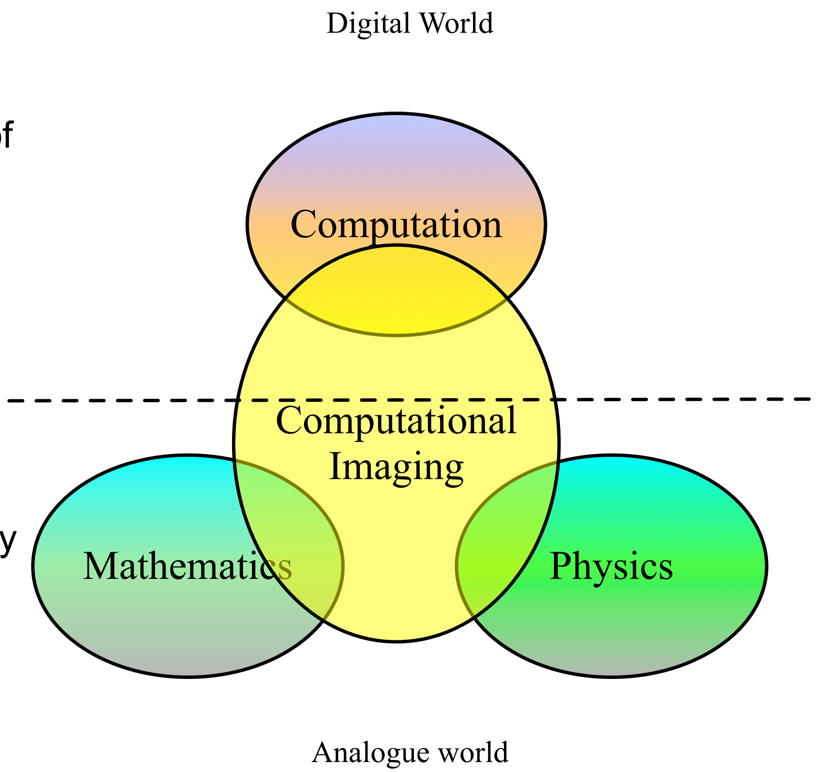

Imperial College London

Two-Photon Microscopy for Neuroscience

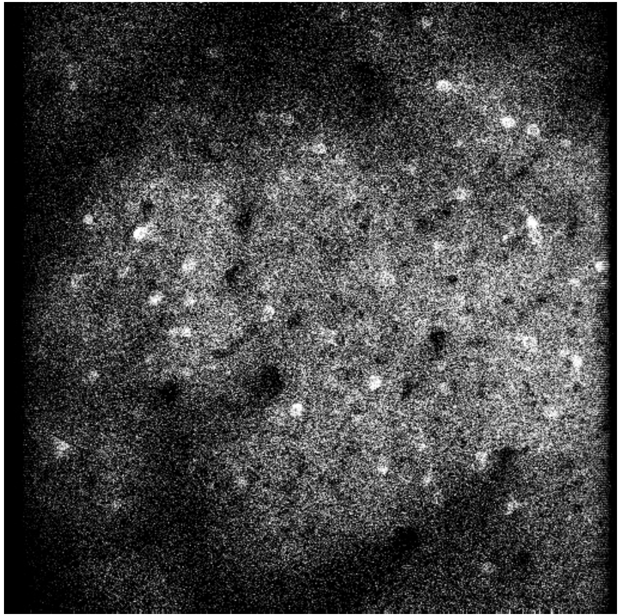

- Goal of Neuroscience: to study how information is processed in the brain
- Neurons communicate through pulses called Action Potentials (AP)
- Need to measure in-vivo the activity of large populations of neurons at cellular level resolution
- Two-photon microscopy combined with right indicators is the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

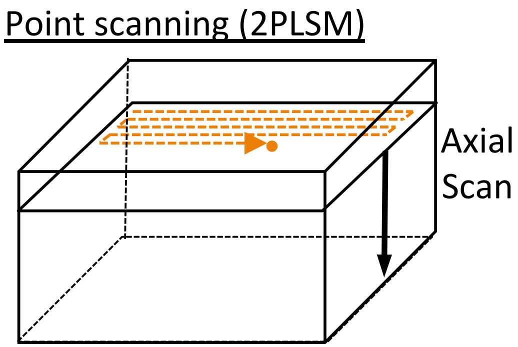

Imperial College London

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

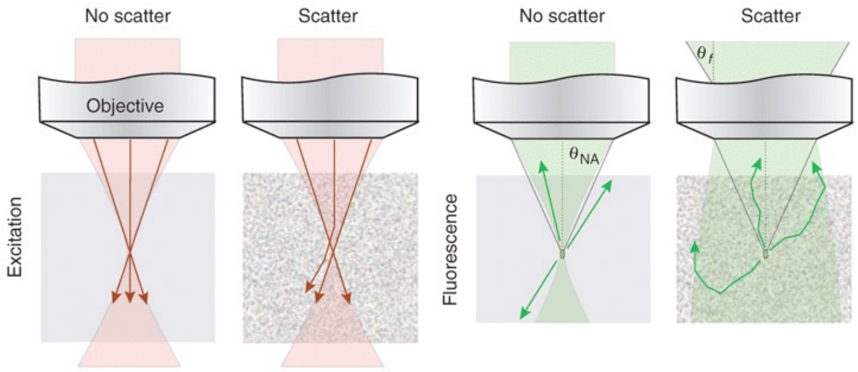

Two-Photon Microscopy

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

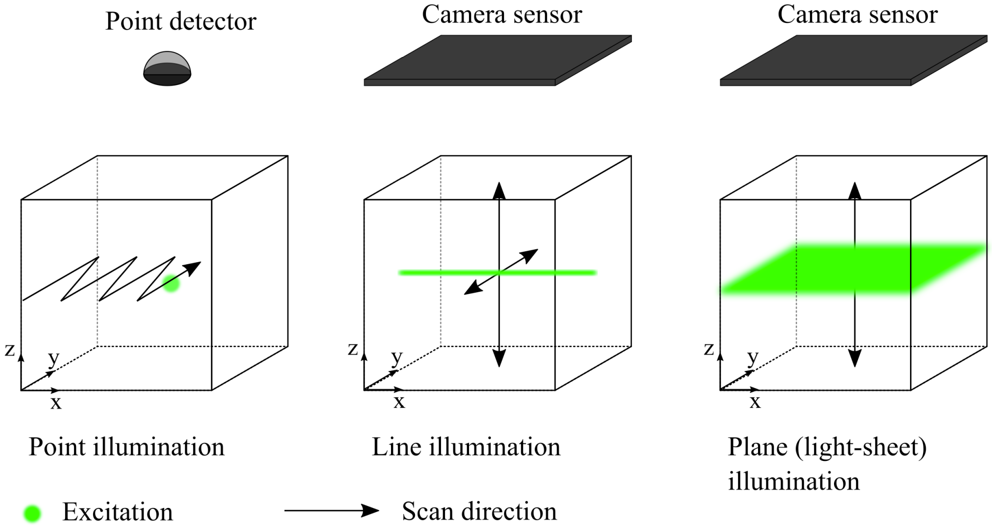

Light-field Microscopy

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

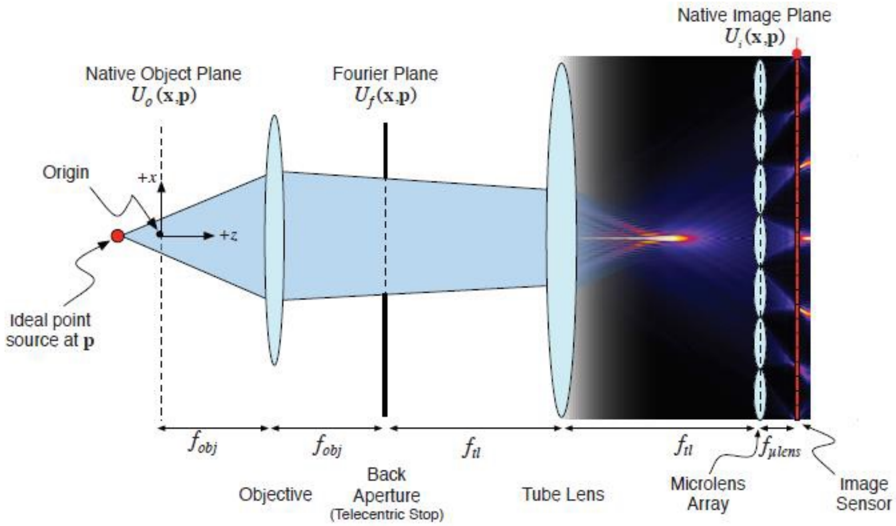

Light-field Microscopy and EPI

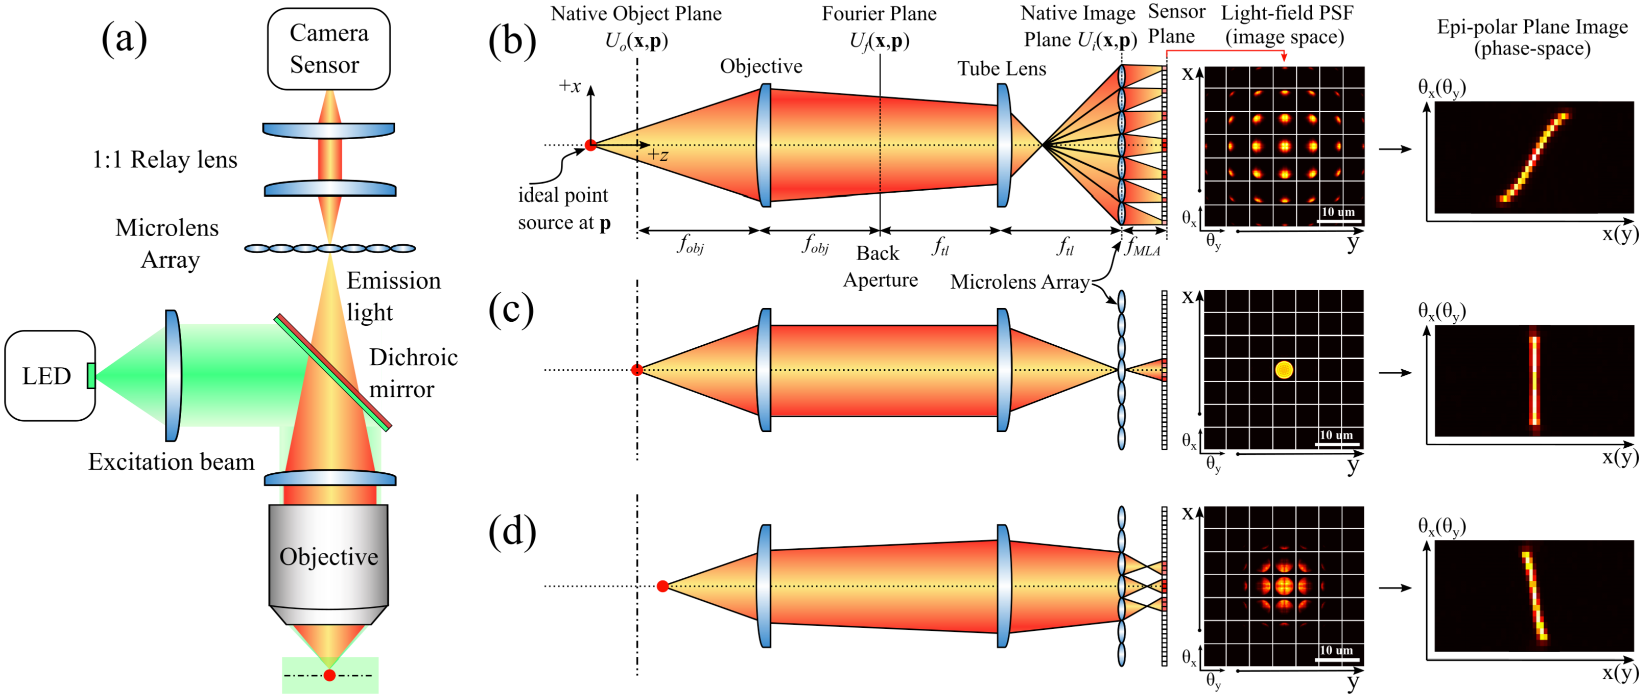

Imperial College London

Light-field Microscopy and Illumination Strategies

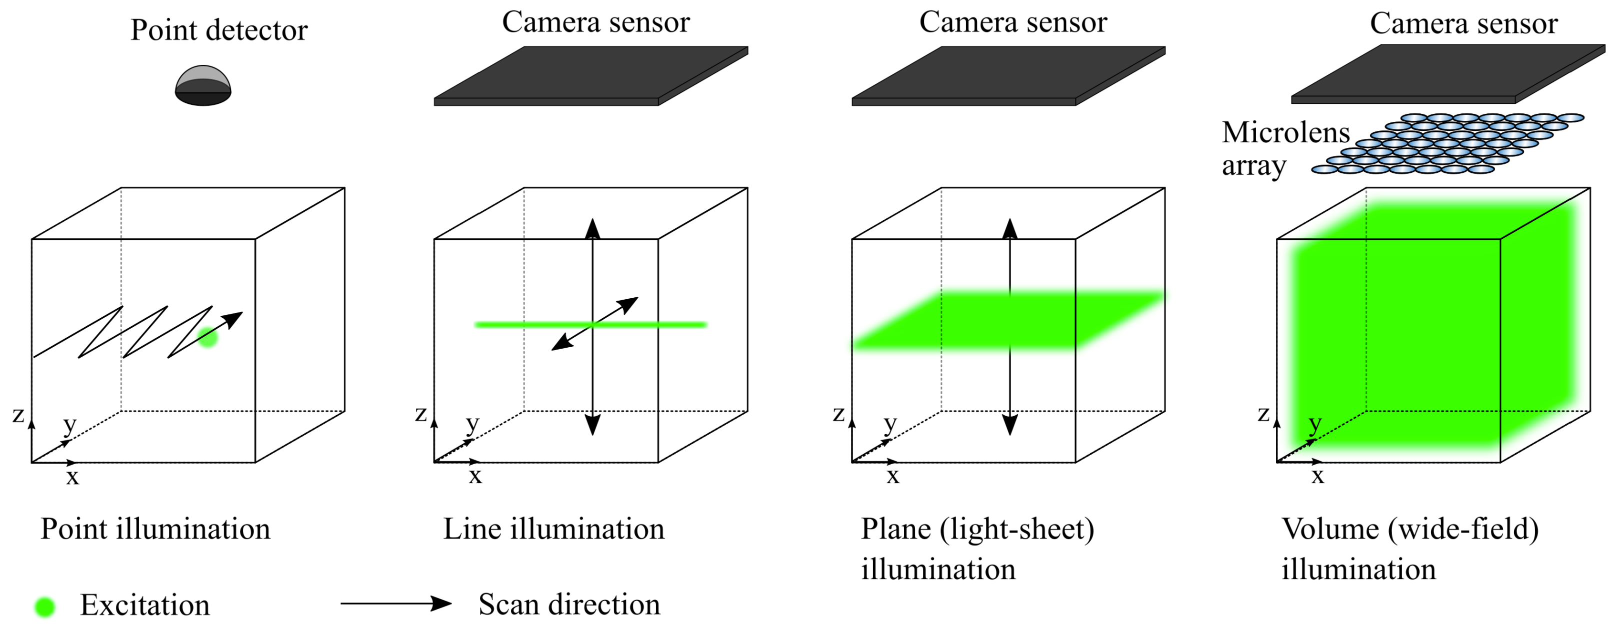

Imperial College London

Light-field Microscopy

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

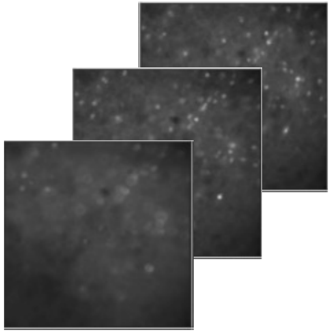

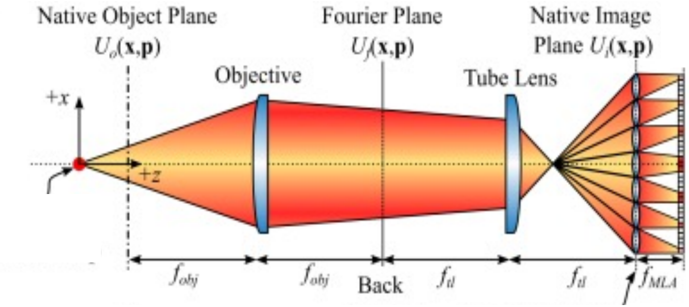

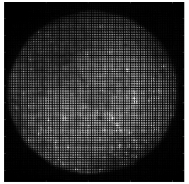

2D Measured LF image

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

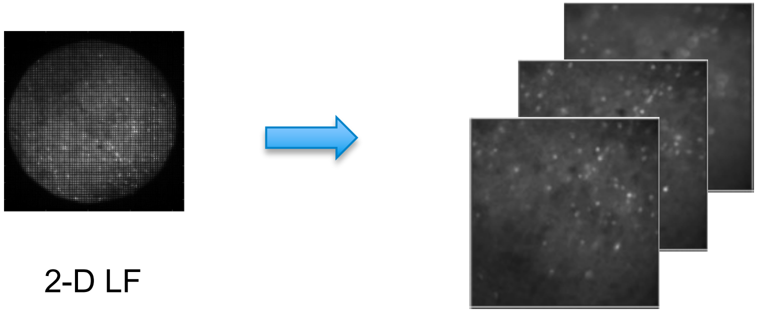

Imperial College London

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)

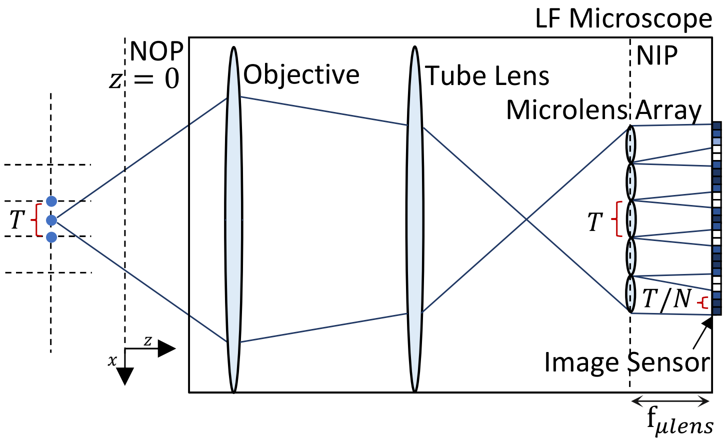

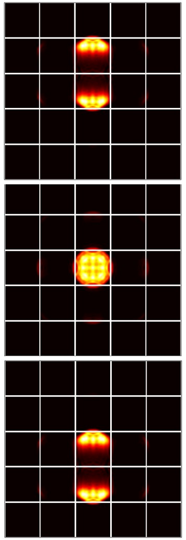

Imperial College London

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

Gregor Karol and LeCunYann, “Learning fast approximations of sparse coding ”, Proceedings of the 27th International Conference on International Conference on Machine Learning, 2010

Imperial College London

Unfolding Strategy

Explicit embedding of priors and constraints in deep networks

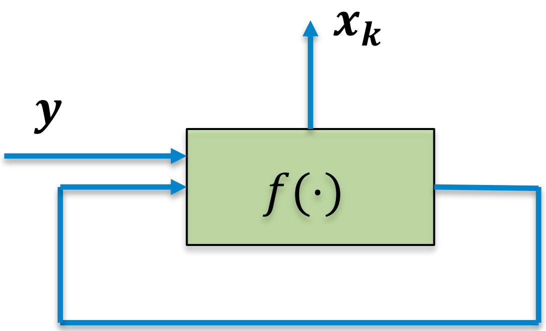

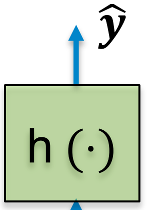

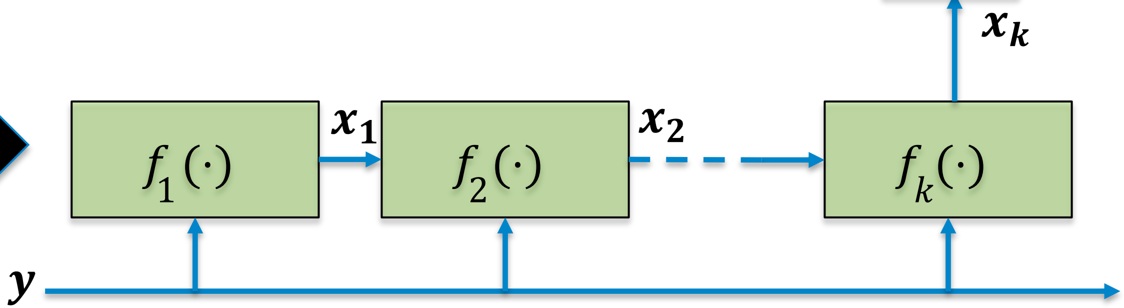

Iterative algorithm with y as input and x as output

Unfolded version of the iterative algorithm with learnable parameters

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

Neural network for volume reconstruction

Imperial College London

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

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

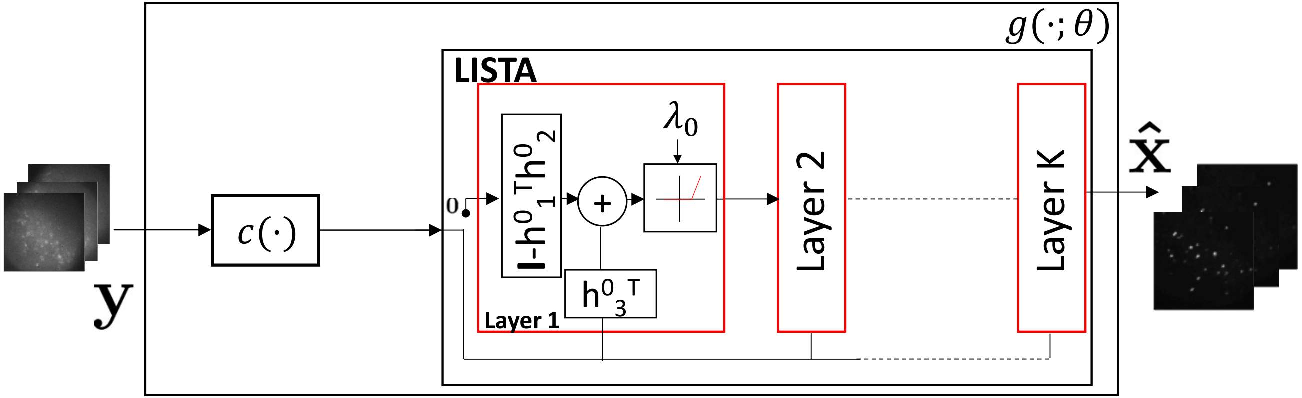

Imperial College London

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

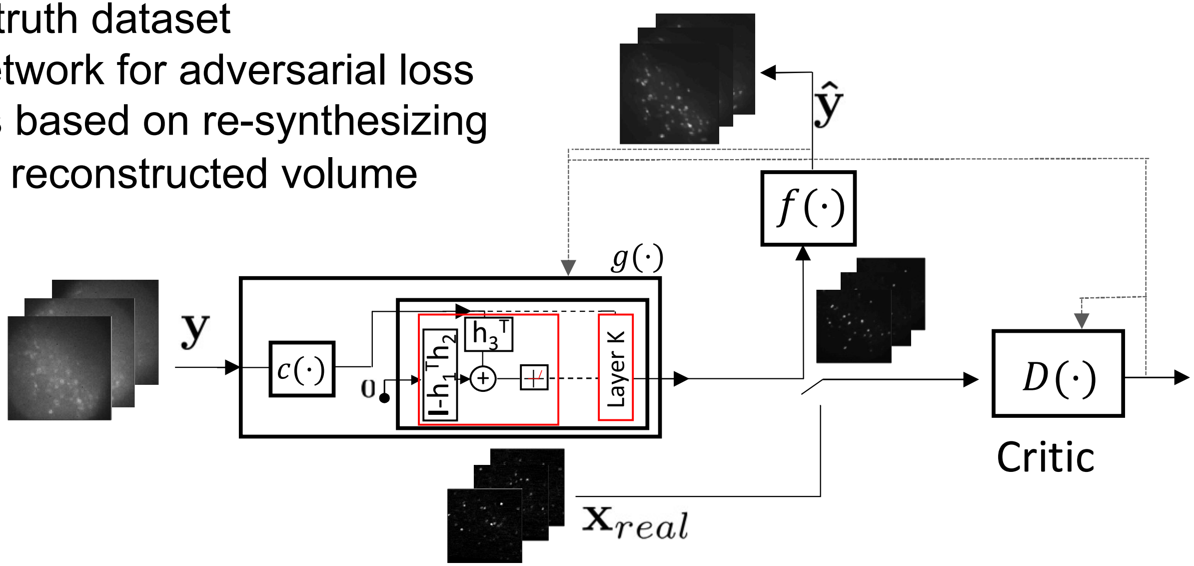

Imperial College London

Training of the neural network

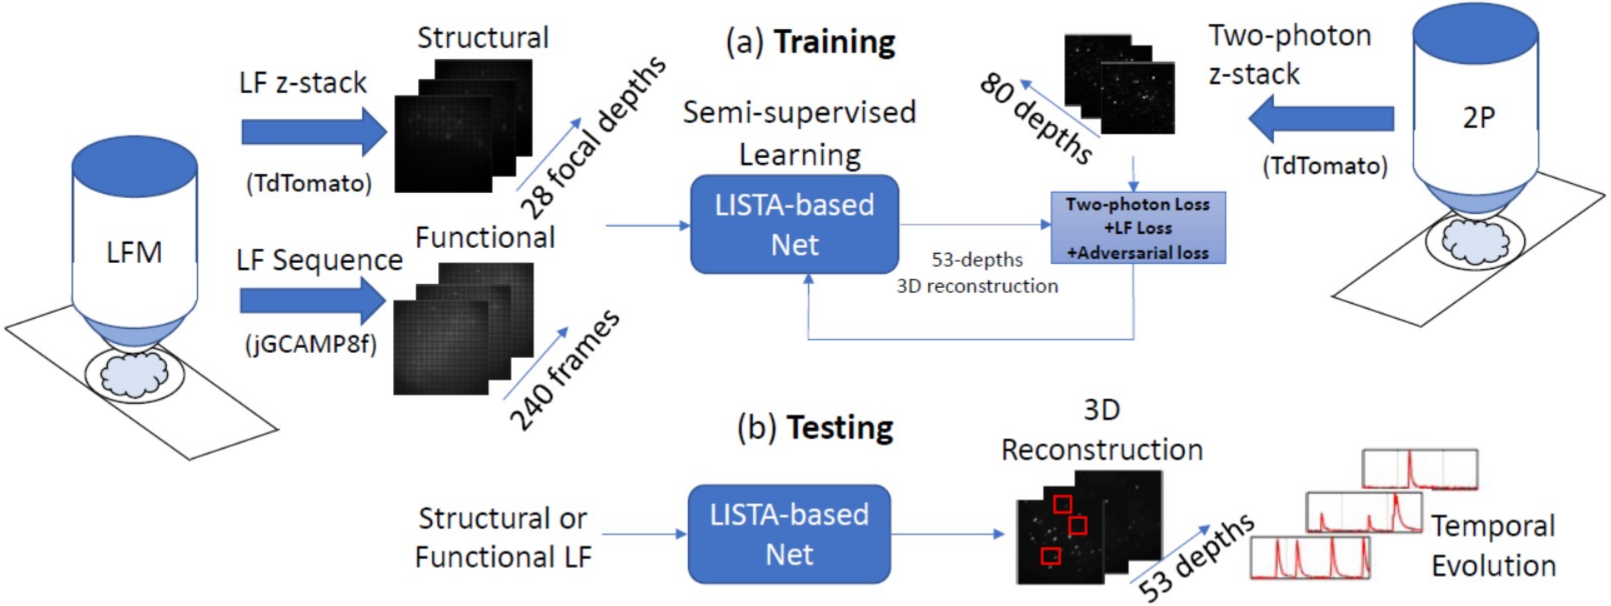

The picture can't be displayed.

n't be displayed.

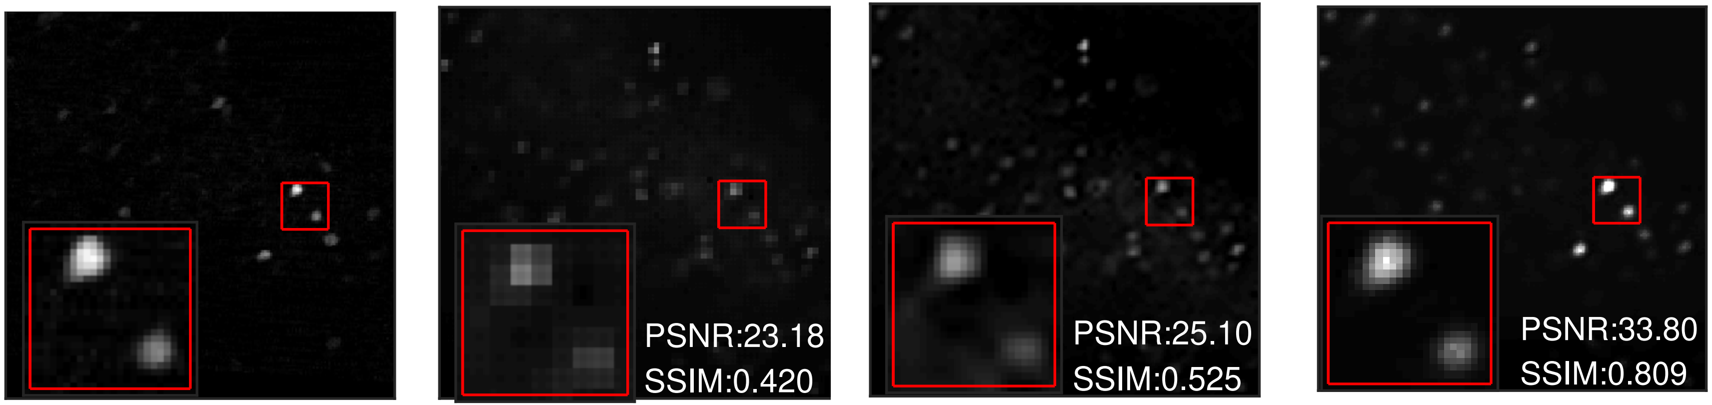

Ground-truth | ISRA | ADMM | New method (0.3s to reconstruct one volume)

Imperial College London

Results – Functional Data

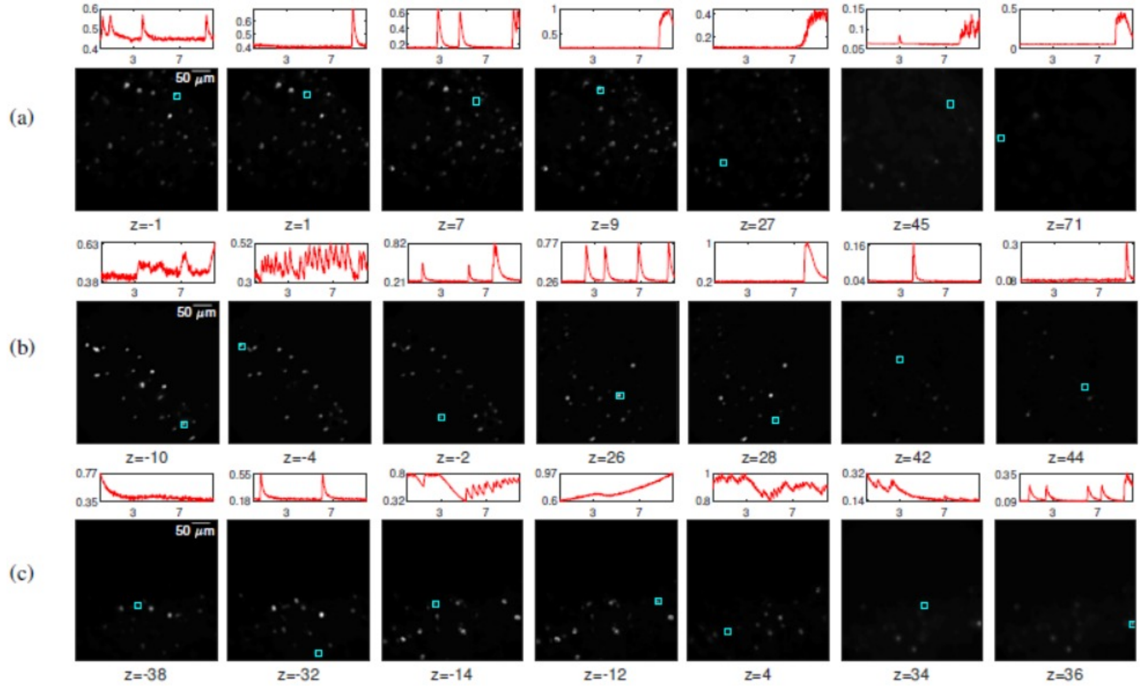

Three brain samples are shown in parts (a), (b), and (c)

Sample 2 (S2A3)

Imperial College London

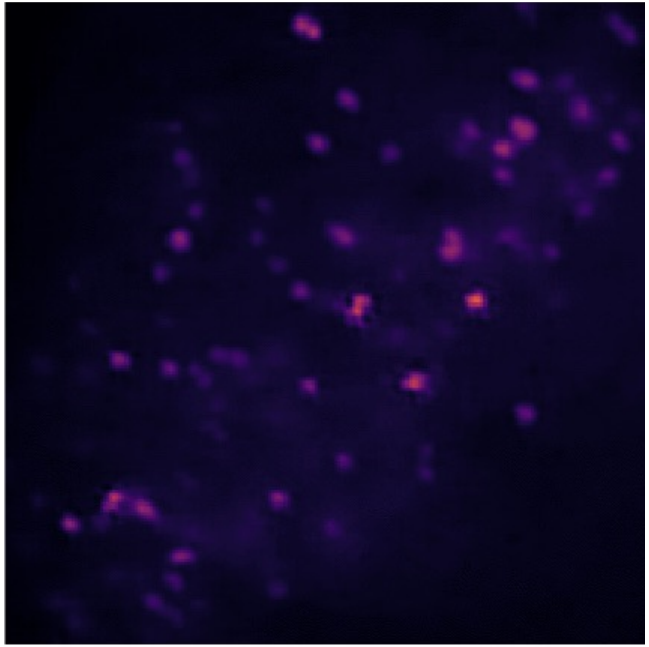

Z Average

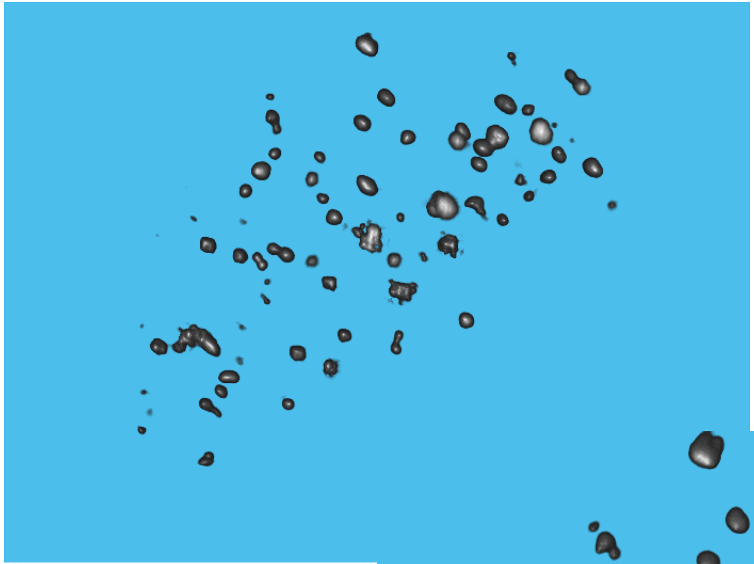

3D (view 1)

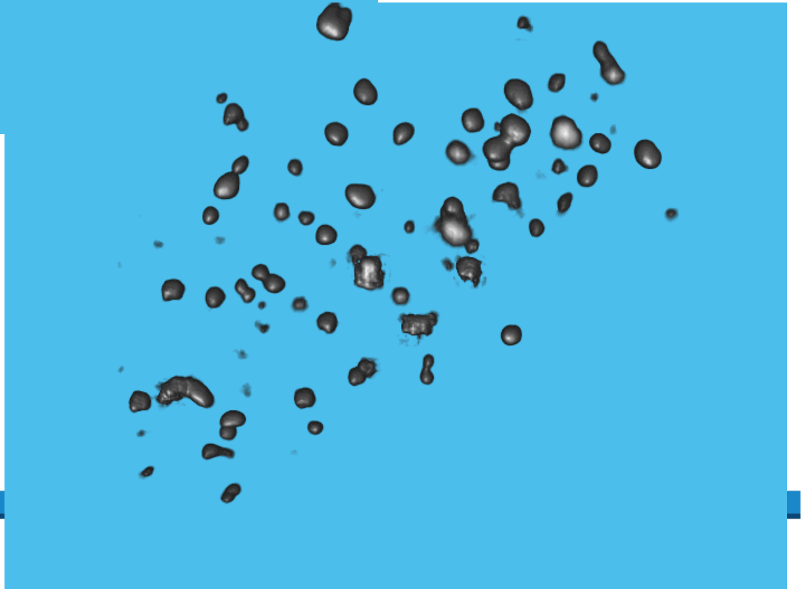

3D (view 2)

Imperial College London

Conclusions

- Cross fertilization between sparse representation and deep learning is fruitful
- Computational Imaging:
    - Light field microscopy can have an impact in neuroscience because of the crucial trade-off between resolution in time and space
    - Understanding the physics of the problem is crucial
    - Learning will labelled data is challenging

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

- **P. Song, H. Verinaz Jadan**, C. Howe, P. Quicke, A. Foust and P.L. Dragotti, Light-field microscopy for optical imaging of neuronal activity, IEEE Signal Processing Magazine, 2022.

- **H. Verinaz** et al. "Physics-based Deep Learning for Imaging Neuronal Activity via Two-photon and Light-field Microscopy, submitted to IEEE Trans. on Computational Imaging, 2022,