Imperial College London

7 September 2023

Computational Imaging: Theory and Applications

Pier Luigi Dragotti

Imperial College London

Computational Imaging

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

Seeing imaging as a whole is the domain of Computational Imaging

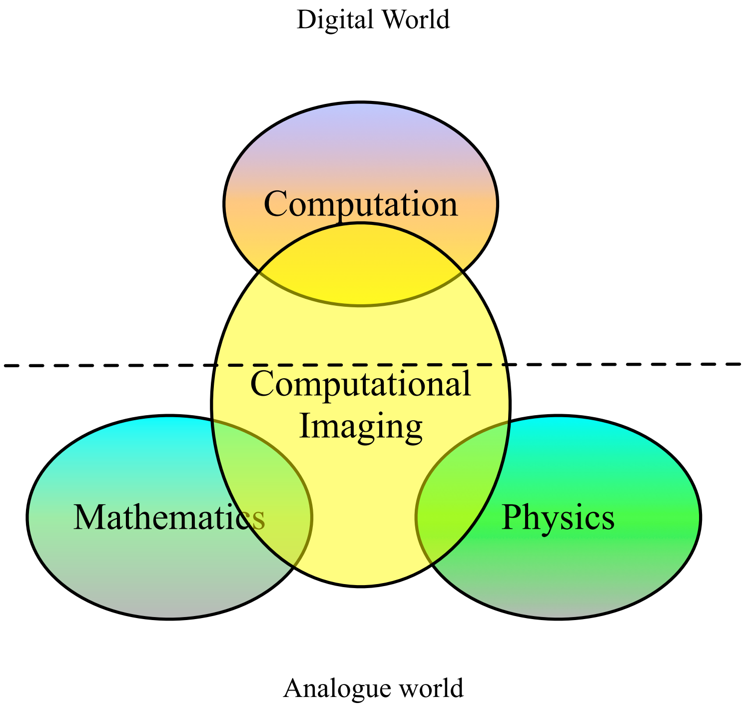

Central Topics in Computational Imaging

Sampling theory

- Sampling theory provides
  - the bridge between the analogue and digital domains
  - appropriate models for the signals and the acquisition devices
  - constructive reconstruction algorithms that can inspire the design of deep neural networks
  - performance bounds

Interplay between physical and learned models

- How to embed priors about the signal and the device in neural networks
  - Deep Unfolding

Three Case Studies in Computational Imaging

Imperial College London

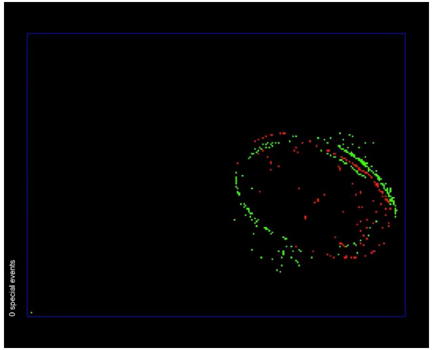

Event-Driven Cameras for energy efficient fast sensing[1]
→ ***Time-based*** Sampling
→ ***End-to-End learning*** using sampling and deep unfolding

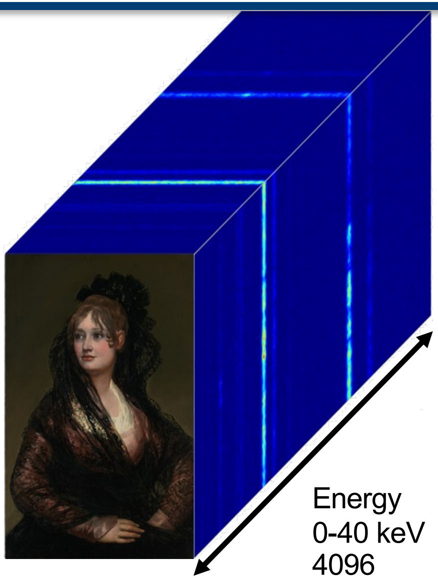

Technical study of Old Masters paintings
→ ***Sparse*** sampling methods

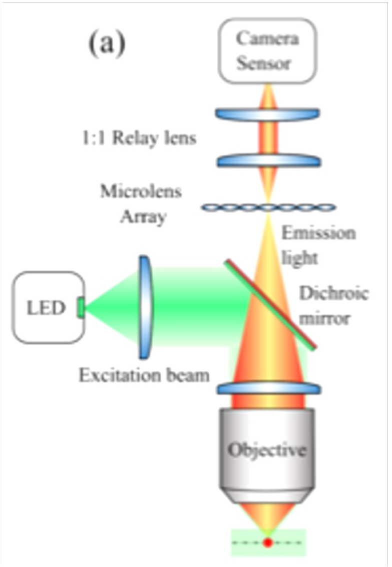

Light field microscopy
→ Traditional sampling theory
→ Deep unfolding

[1]Note: video taken from Inivation.com

Imperial College London

Sensing based on Timing Information

- Energy-efficient sensing inspired by nature raises a fundamental representation question:
  - How can we embed information related to complex signals into the timing information of spikes?
  - Besides its theoretical implications, addressing this question will lead to new neuromorphic sensing devices

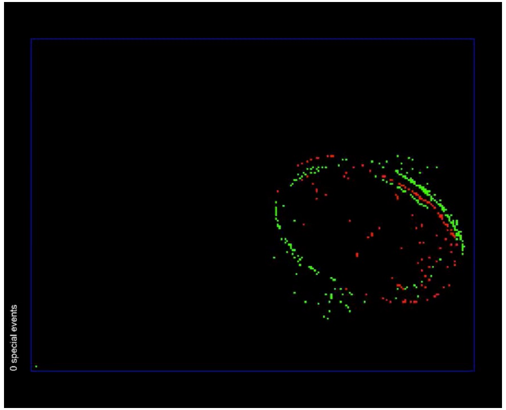

Videos from Inivation.com

Imperial College London

Bio-Inspired Energy Efficient Sensing

- Current sensing methods are energy inefficient especially when low-latency is needed.
- Example: Rainfall estimation

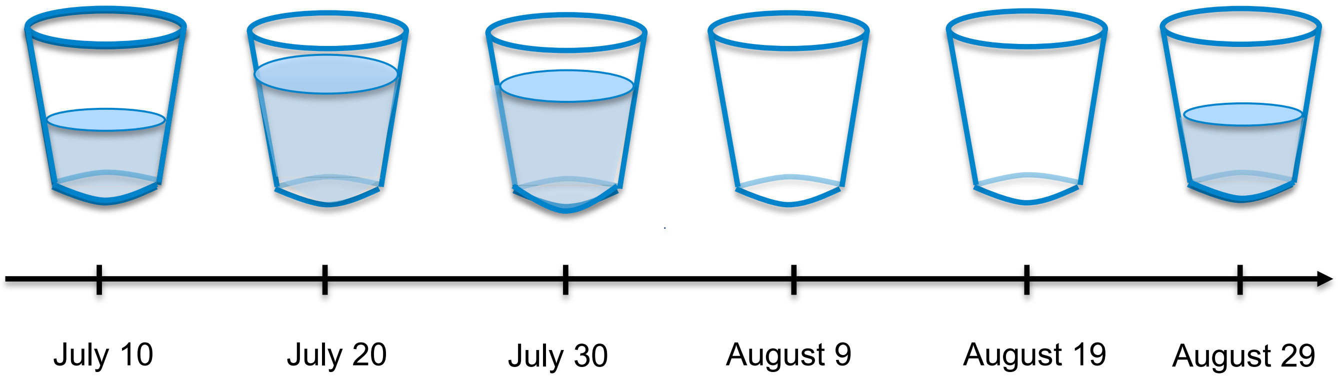

Imperial College London

Bio-Inspired Energy Efficient Sensing

Approach 2

- Only record the day when the bucket is full and then empty it

Bio-Inspired Energy Efficient Sensing

Approach 2 maps analogue information into a time sequence and is used by nature (e.g., integrate-and-fire neurons)

Time encoding appears in nature, as a mechanism used by neurons to represent sensory information as a sequence of action potentials, allowing them to process information **very efficiently**.

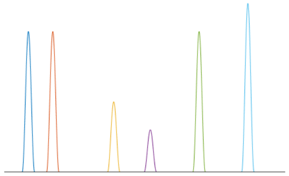

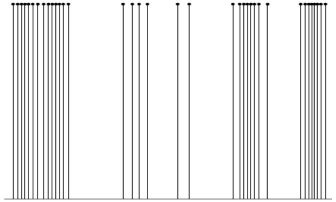

Imperial College London

Time-Encoding Machines

Integrate-and-fire System

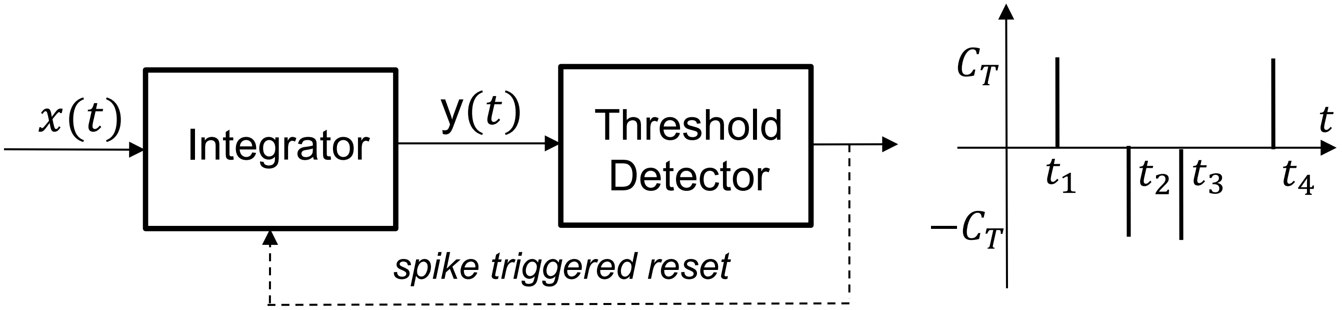

A. A. Lazar and L. T. Toth, “Perfect recovery and sensitivity analysis of time encoded bandlimited signals,” IEEE Trans. Circuits Syst. I, Oct. 2004.

Imperial College London

Reconstruction from time-encoded information

- Reconstruction achieved by imposing iteratively:
  - Consistency constraint
  - Signal prior (e.g., bandlimited function) constraint

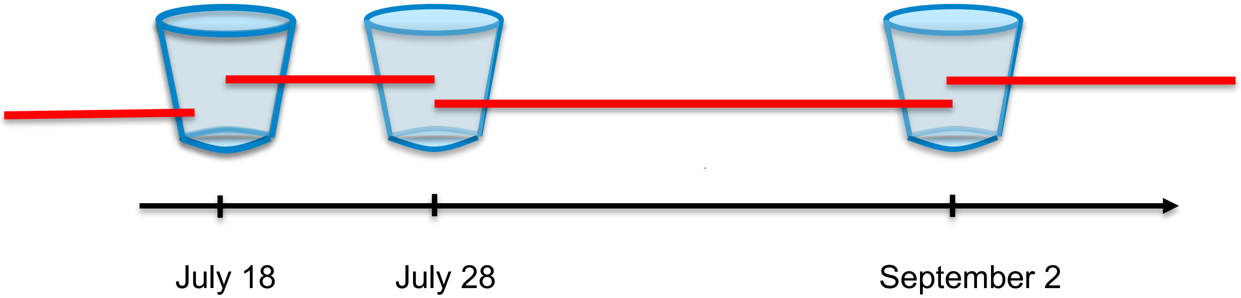

Imperial College London

Reconstruction from time-encoded information

- Reconstruction achieved by imposing iteratively:
  - Consistency constraint
  - Signal prior (e.g., bandlimited function) constraint

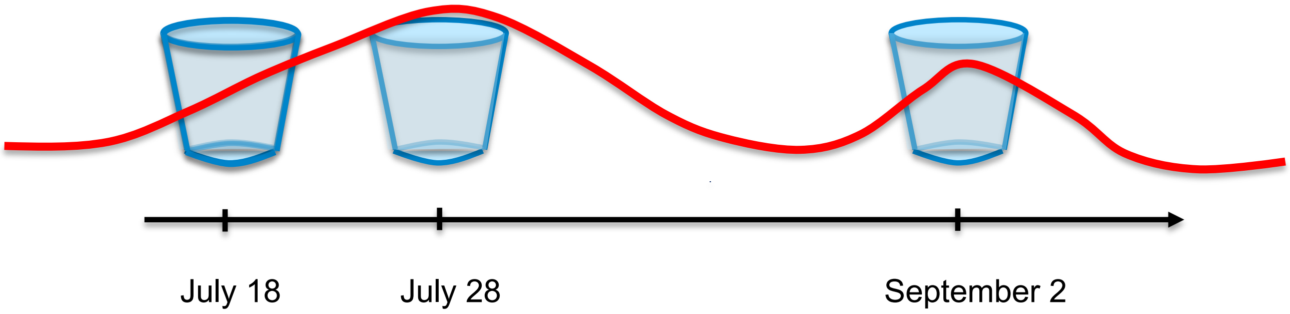

Reconstruction from time-encoded information

- **Key result:** if the density of samples $D \geq 1$ then perfect reconstruction can be achieved (Aldroubi and Grochenig[1])

- **Key Issue 1**: In the case of uniform sampling the density is $D = 1$. This means that current TEMs are **less** energy efficient than uniform sampling!

- **Key Issue 2:** Cannot sample sparse (non-bandlimited) signals with the current methods.

[1] A. Aldroubi and K, Grochenig, "Non-Uniform Sampling and Reconstruction in shift-invariant spaces" SIAM Review 2001

Imperial College London

Reconstruction from time-encoded information

- For integrate-and-fire machines exact reconstruction proved here: A. A. Lazar and L. T. Toth, “Time encoding and perfect recovery of bandlimited signals”, ICASSP 2003

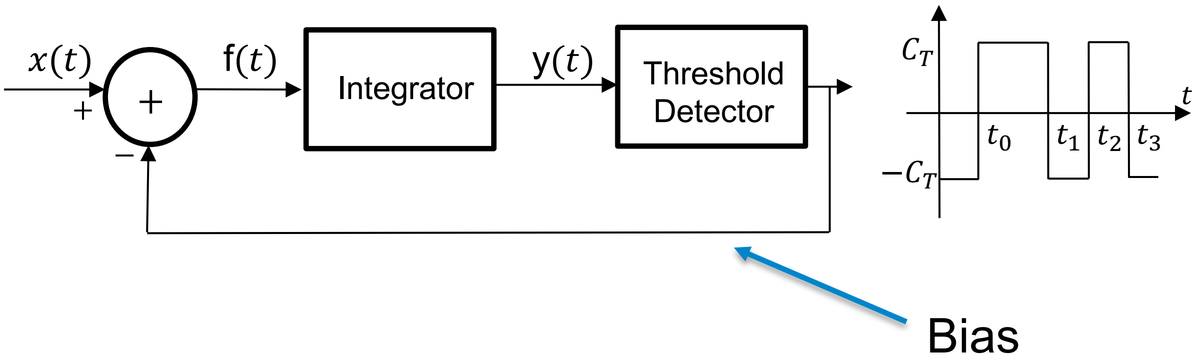

See also: Gauntier-Vetterli-2014, Adam et al 2019,

Imperial College London

Time-based Sampling of Sparse Signals

Signals:

- We consider sparse continuous-time signals like stream of pulses, piecewise constant or regular signals

Sensing Systems:

- We filter before using a TEM

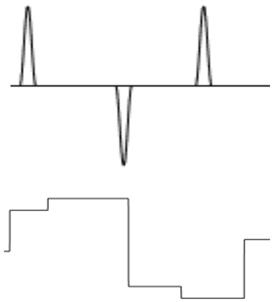

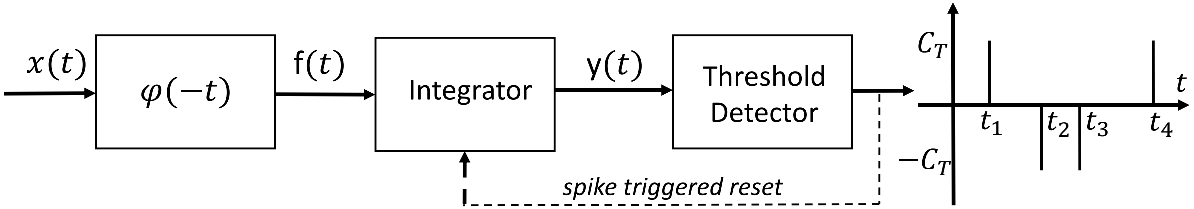

Imperial College London

Integrate and Fire – Reconstruction of Pulses

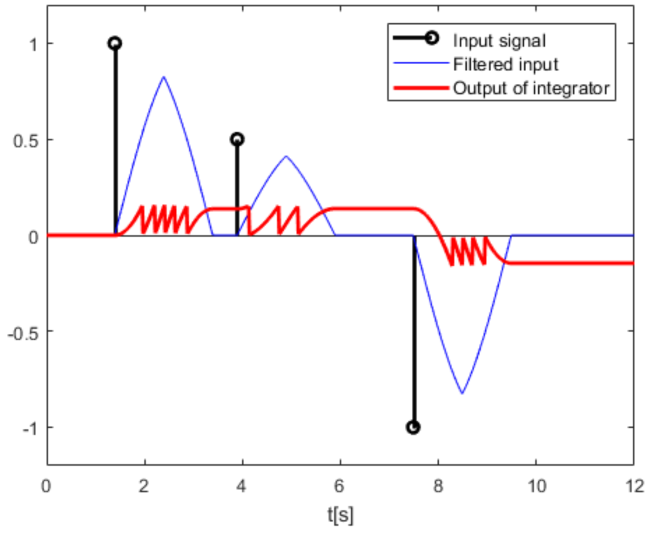

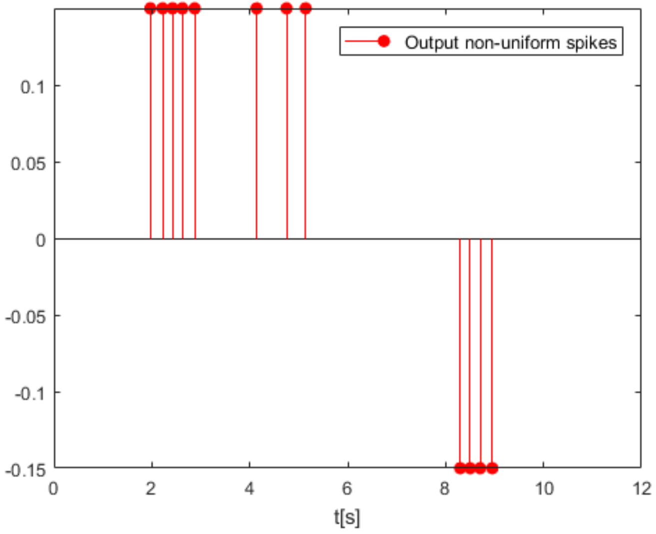

Imperial College London

Our approach for time decoding of signals

- Reconstruction of $x(t)$ depends on the
  - sampling kernel $\varphi(t)$
  - the density of time instants $\{t_n\}$

- We achieve a sufficient density of output samples by imposing conditions on:
  - The trigger mark of the integrator (**integrate-and-fire TEM**).

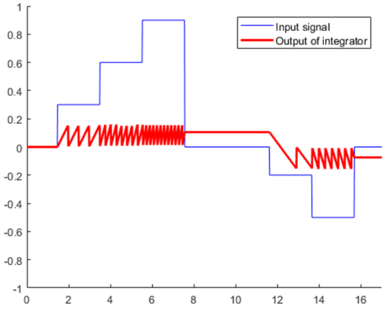

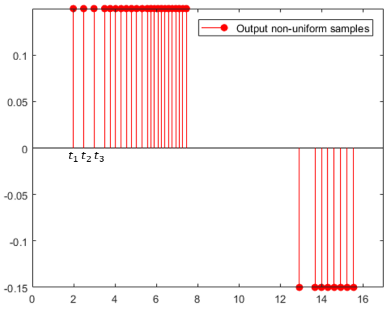

Imperial College London

Integrate and Fire TEM

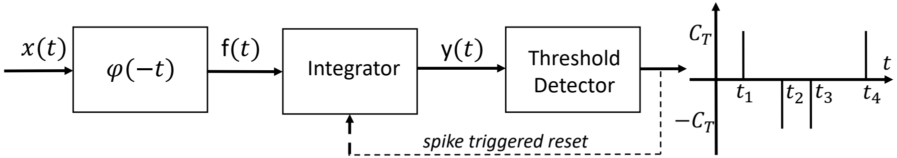

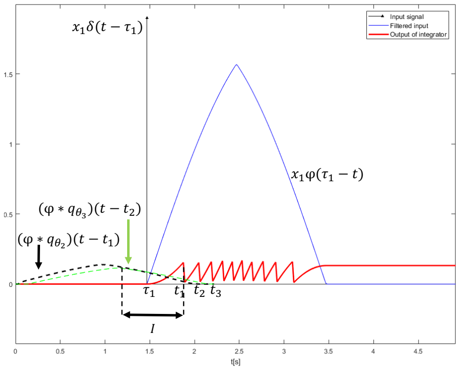

- Given the times $t_1, t_2, \dots, t_n$, the amplitude values are

$$y_n = y(t_n) = \pm C_T = \int_{t_{n-1}}^{t_n} f(\tau)d\tau = \int_{t_{n-1}}^{t_n} \int x(\alpha)\varphi(\alpha - t)d\alpha d\tau.$$

Integrate and Fire TEM

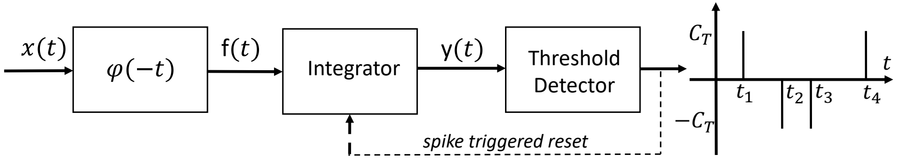

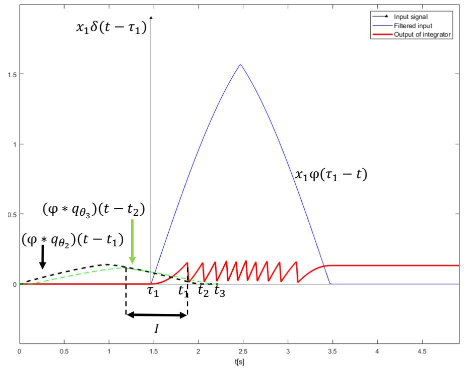

- Equivalently the output samples can be expressed as:

$$y(t_n) = \langle x(t), (\varphi * q_{\theta_n})(t - t_{n-1}) \rangle,$$

where $\theta_n = t_n - t_{n-1}$ and $q_{\theta_n}(t)$ is defined as:

$$q_{\theta_n}(t) = \begin{cases} 1, & 0 \leq t \leq \theta_n, \\ 0, & otherwise. \end{cases}$$

Imperial College London

Integrate and Fire TEM

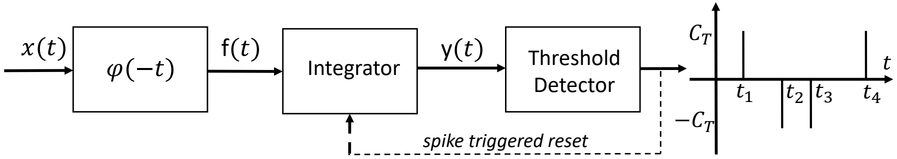

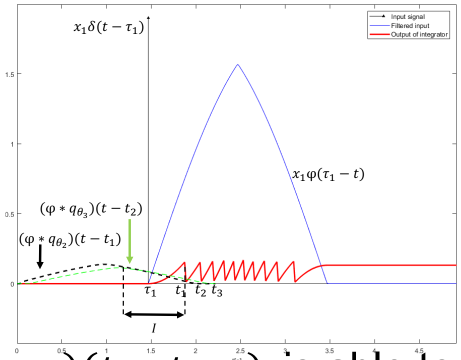

- When $\varphi(t)$ is e.g., an E-spline, the equivalent kernel $(\varphi * q_{\theta_n})(t - t_{n-1})$ is able to reproduce exponentials
- So trigger mark must guarantee enough samples in a short interval

- *Proposition:* when $C_T < \frac{A_{min}}{4\omega_0^2}\left(1 - \cos\left(\frac{\omega_0 L}{2}\right)\right)$ then $t_1, t_2, t_3 \in \left[\tau_1, \tau_1 + \frac{L}{2}\right]$ and perfect reconstruction is possible

Imperial College London

Reproduction of Exponentials

$$\sum_n c_{m,n}\varphi(n-t) \approx e^{j\omega_m t}$$

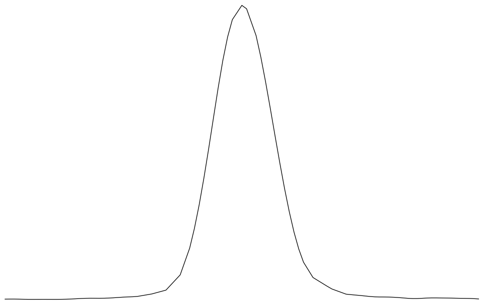

Pulse shape

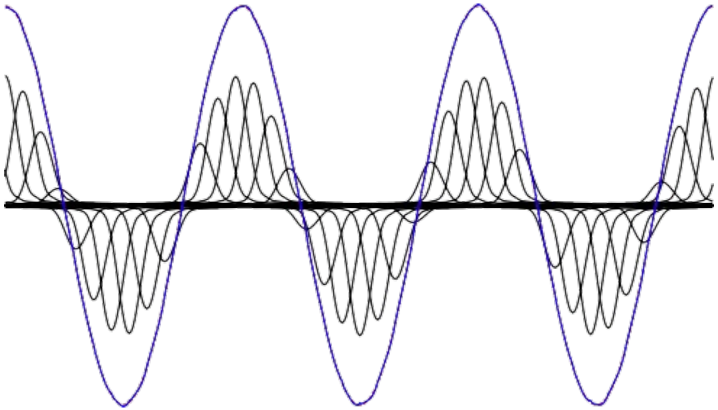

Reproduction of exponentials

Reconstruction of an input Dirac from time-encoded information

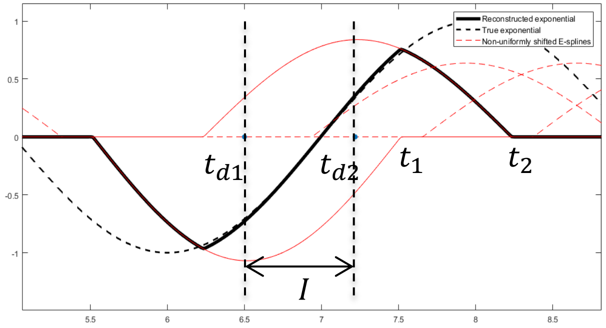

- The output samples are: $y(t_n) = \langle x(t), (\varphi * q_n)(t)\rangle = x_1\varphi_n(\tau_1)$
- Since $\varphi_n(t) = a_{0,n}e^{\alpha_0 t} + a_{1,n}e^{\alpha_1 t}$, we find $\boldsymbol{c_1}$, $\boldsymbol{c_2}$, $\boldsymbol{d_1}$, $\boldsymbol{d_2}$ such that in $I_1 = [t_2 - T, t_1]$:

$$\boldsymbol{c_1}\ \varphi_1(t) + \boldsymbol{c_2}\ \varphi_2(t) = e^{\alpha_0 t}$$

$$\boldsymbol{d_1}\ \varphi_1(t) + \boldsymbol{d_2}\ \varphi_2(t) = e^{\alpha_1 t}$$

- We then use these coefficients to define the signal moments, in $I_1 = [t_2 - T, t_1]$:

$$s_0 = c_1 y(t_1) + c_2 y(t_2) = x_1[c_1\varphi_1(\tau_1) + c_2\varphi_2(\tau_1)] = x_1 e^{\alpha_0\tau_1}$$

$$s_1 = d_1 y(t_1) + d_2 y(t_2) = x_1[d_1\varphi_1(\tau_1) + d_2\varphi_2(\tau_1)] = x_1 e^{\alpha_1\tau_1}$$

Imperial College London

Integrate and Fire – Reconstruction of Pulses

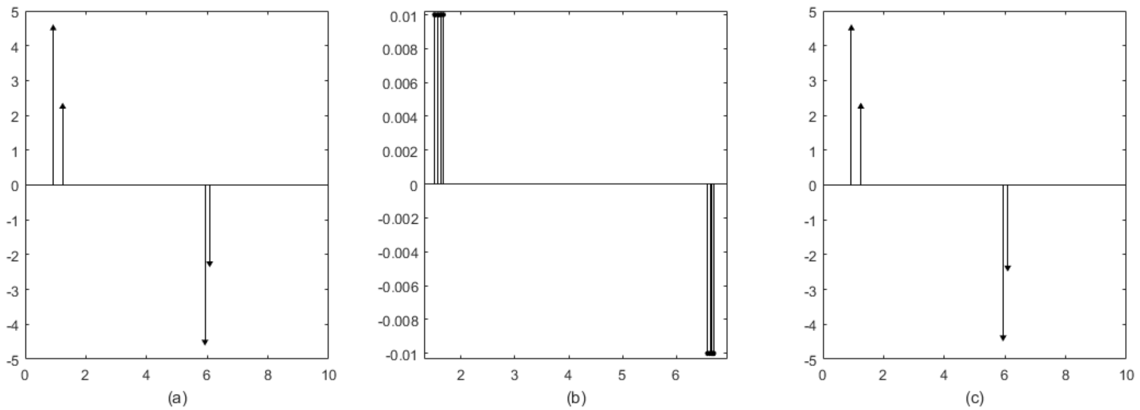

Energy Efficient Sampling -Results

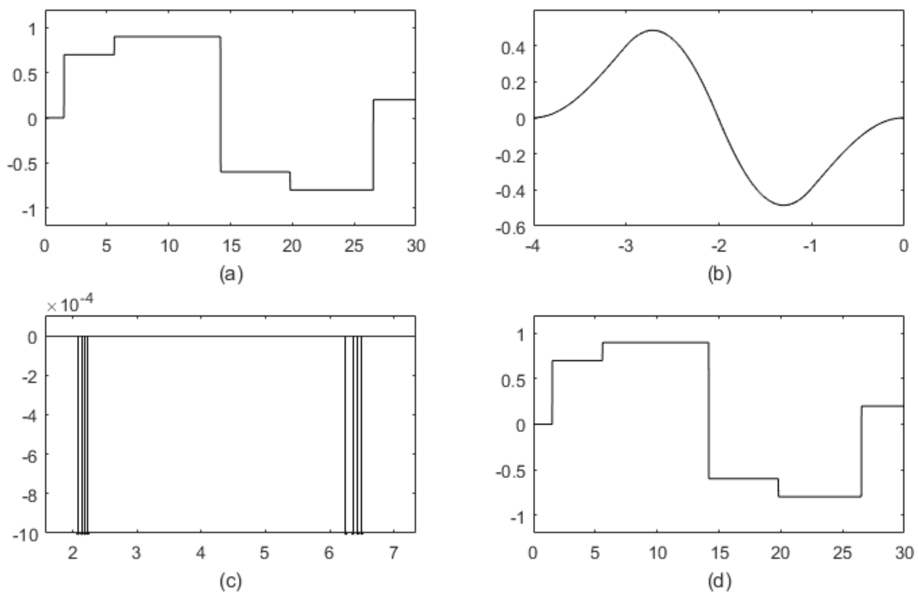

If the distance S between discontinuities is on average $S > (L-1)T$ with T being the sampling period in uniform sparse sampling then the new time encoding framework[2] is **more efficient** than sparse sampling (lower sampling density

[2]R. Alexandru and P.L. Dragotti, Reconstructing Classes of Non-bandlimited Signals from Time Encoded Information, IEEE Trans. on Signal Processing, vol.68, 2020.

Imperial College London

Integrate and Fire and Neuromorphic Cameras

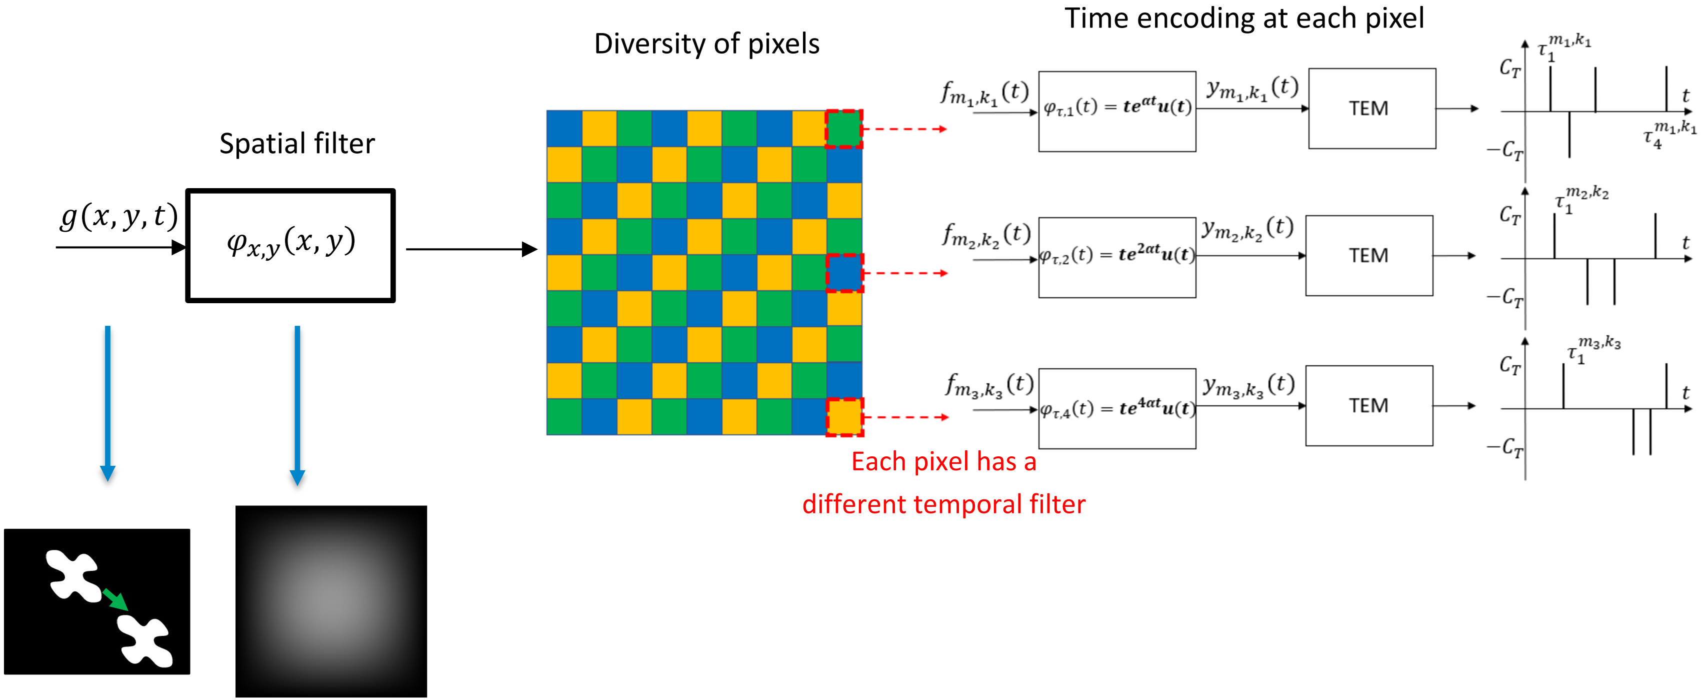

Integrate and Fire and Neuromorphic Cameras

Sensing Diversity for Neuromorphic Cameras

- Key insight: design an end-to-end neural network where the acquisition
- process is part of the learned architecture
- Key approach: each pixel behaves differently
- The network architecture for reconstruction is model-based

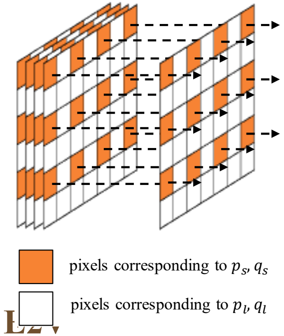

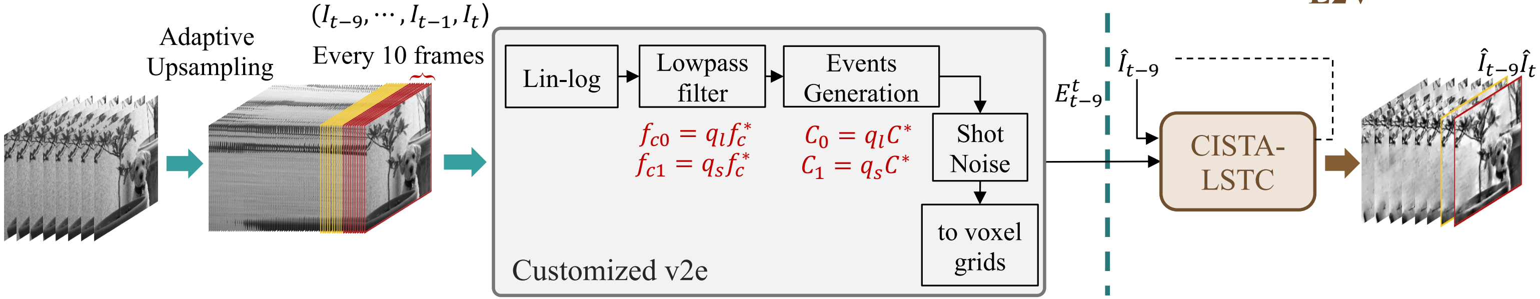

Imperial College London

Sensing Diversity for Neuromorphic Cameras

- The network architecture for reconstruction is model-based: intensity and event frames share the same sparse representation

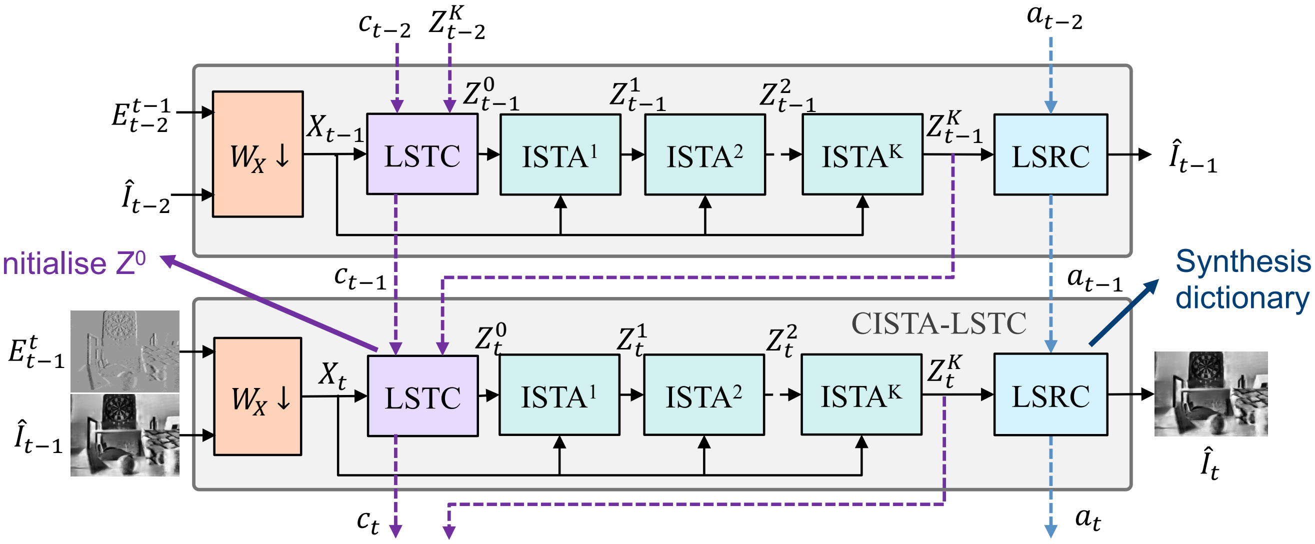

Imperial College London

Deep Unfolding Strategy

Explicit embedding of priors and constraints in deep networks

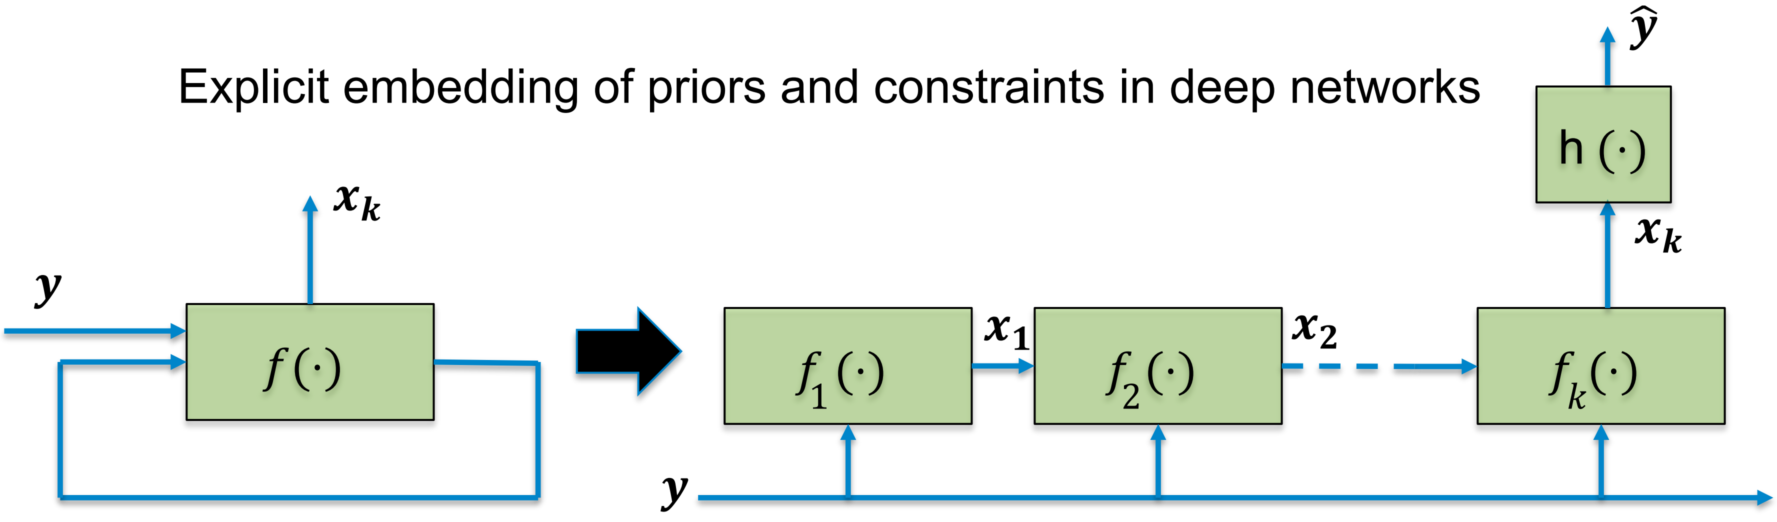

Iterative algorithm with $\boldsymbol{y}$ as input and $\boldsymbol{x}$ as output

Unfolded version of the iterative algorithm with learnable parameters

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

Imperial College London

Dictionary Learning

- The network architecture for reconstruction is model-based: intensity and event frames share the same sparse representation
- The dictionary is usually learned

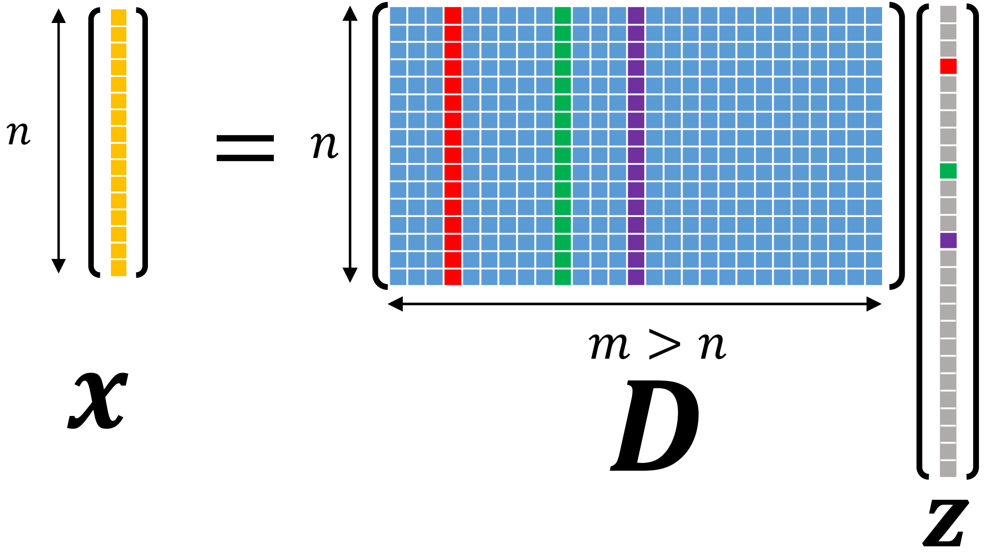

Model of dependency between intensity and events

Assumption: intensity and event frames share the same sparse representation

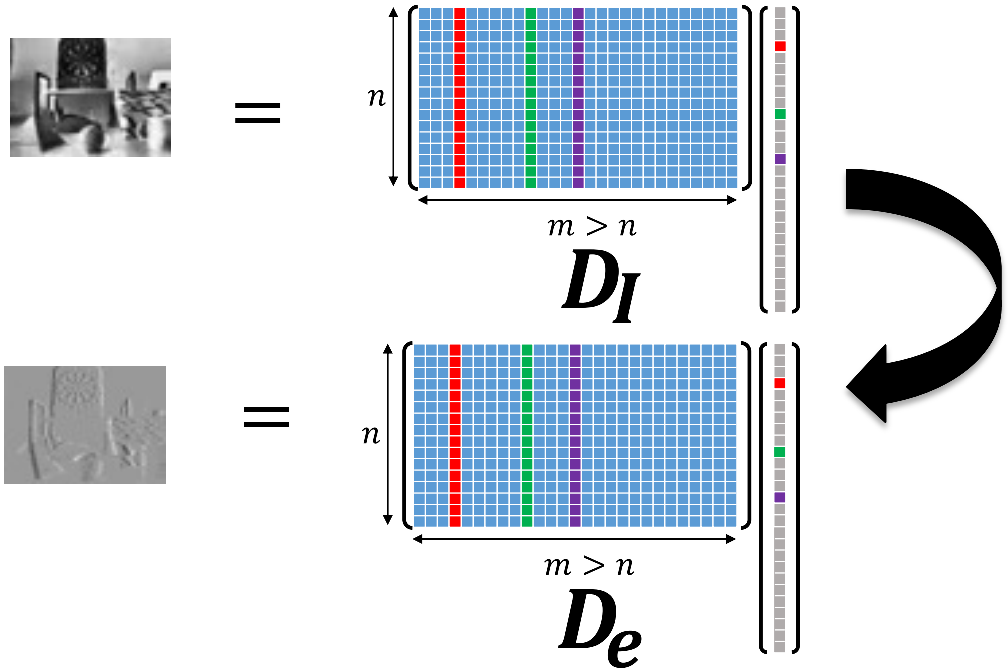

Deep Unfolding Strategy

- The network architecture for reconstruction is model-based: intensity and event frames share the same sparse representation Z_t
- The sparse vector can be found using ISTA: $Z_t^k = h_\theta(Z_t^{k-1} + D_k^T(X_t - D_k Z_t^{K-1}))$

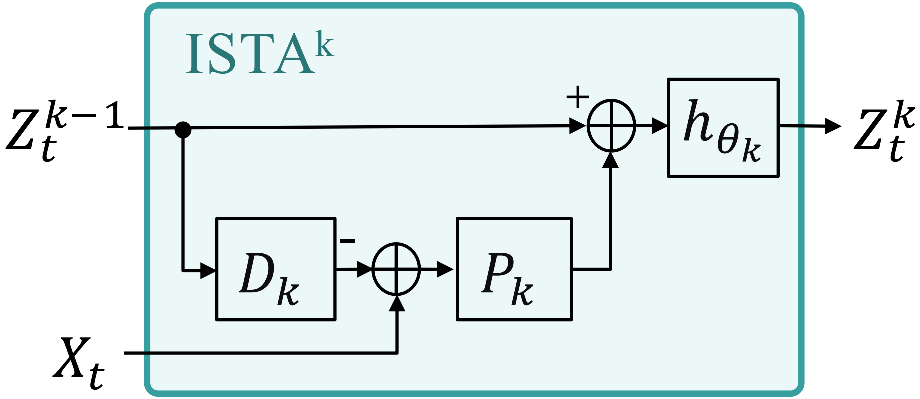

Imperial College London

Sensing Diversity for Neuromorphic Cameras

- The network architecture for reconstruction is model-based: intensity and event frames share the same sparse representation

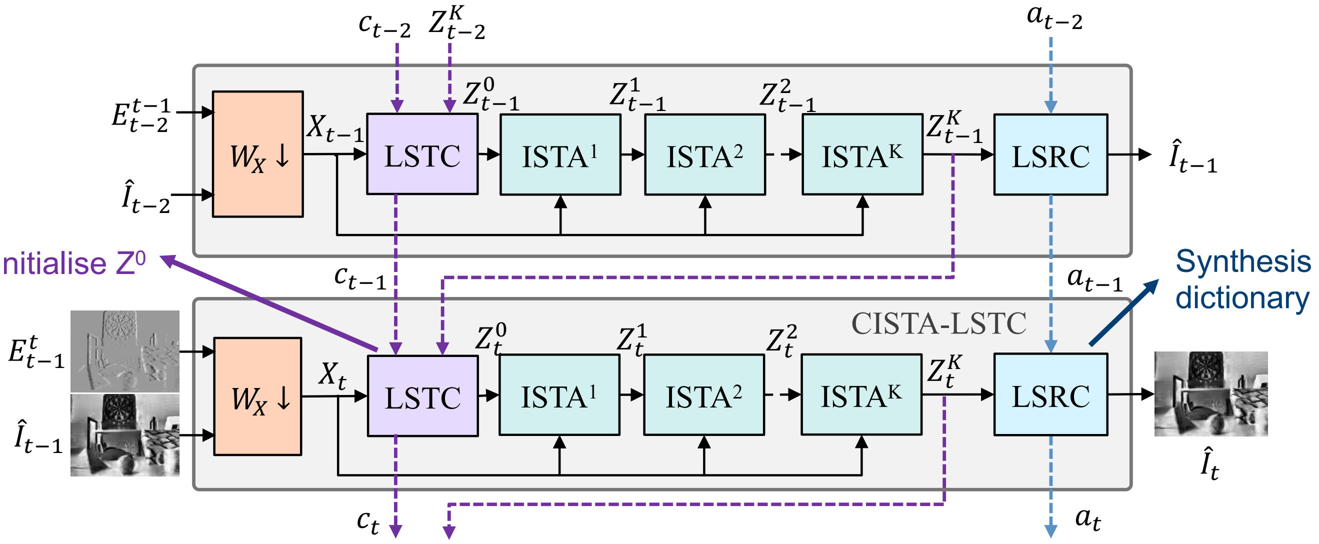

Imperial College London

Reconstruction using DU

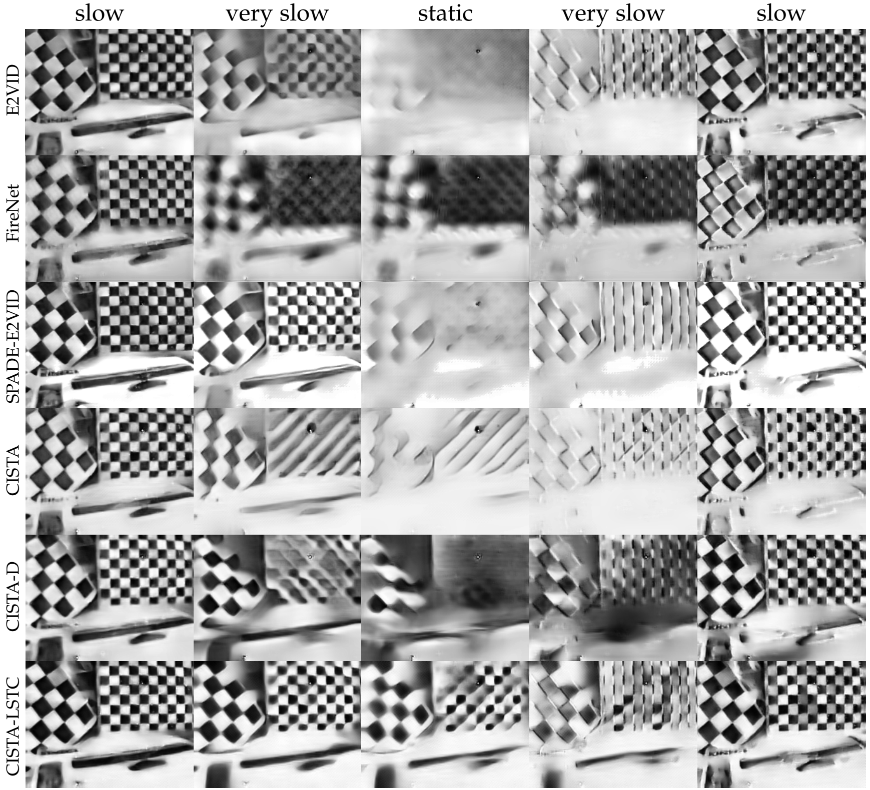

Model-based reconstruction from events

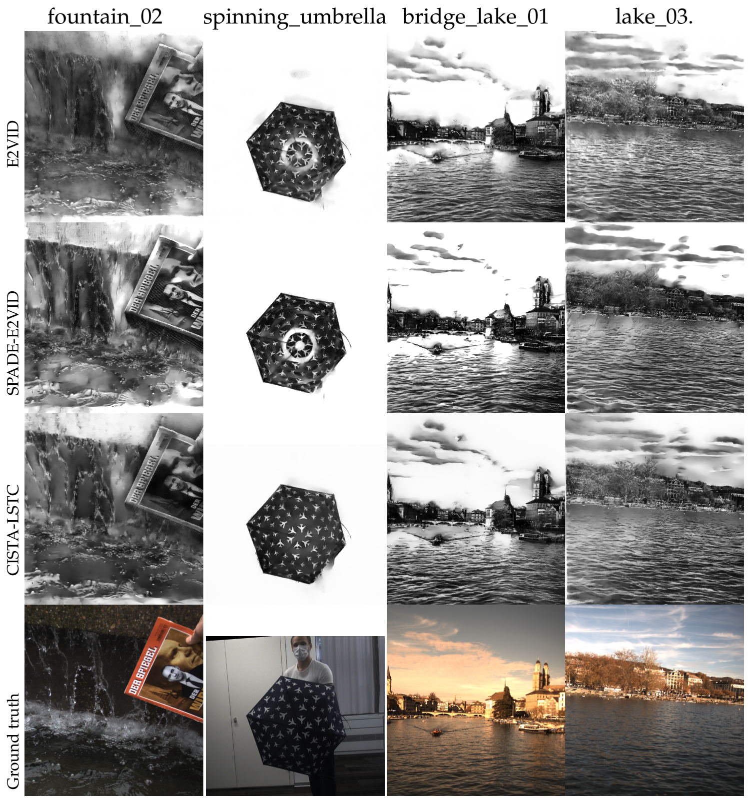

Model-based reconstruction from events

gun_bullet_mug

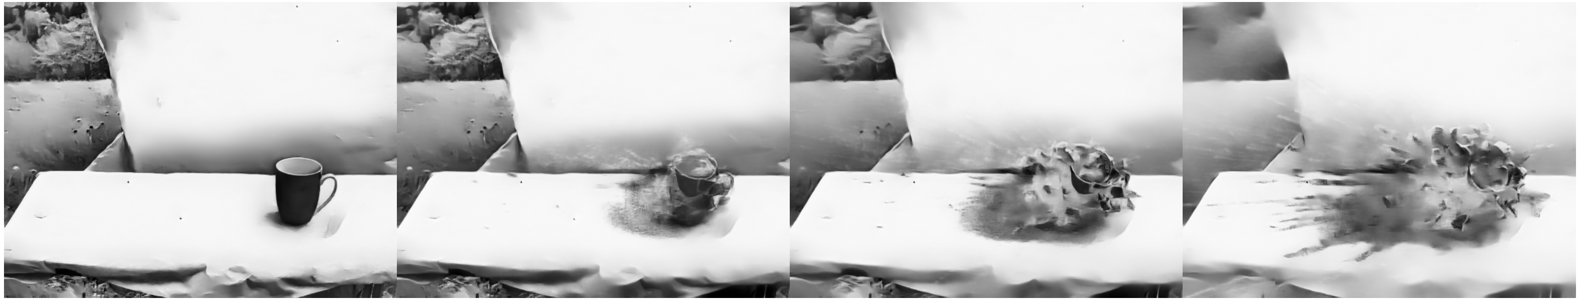

Imperial College London

Reconstruction using DU

Imperial College London

Sensing Diversity for Neuromorphic Cameras

- Key insight: design an end-to-end neural network where the acquisition process is part of the learned architecture
- Key approach: each pixel behaves differently
- The network architecture for reconstruction is model-based

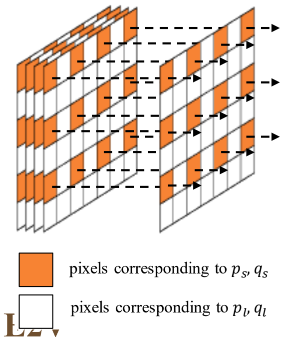

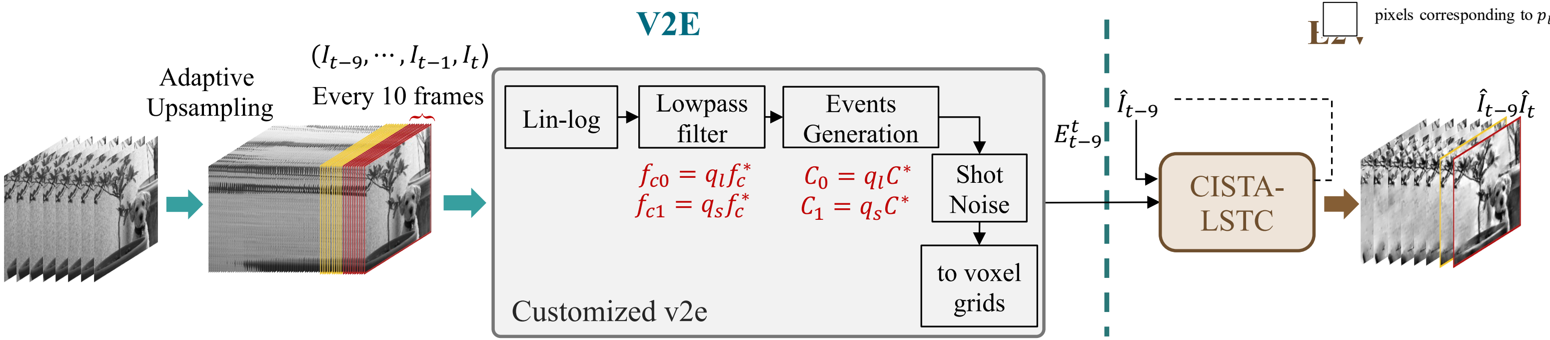

Imperial College London

End-to-end learning - Results

without sensing diversity

with sensing diversity

Event-driven systems – first set of conclusions

- Neuromorphic sensing systems inspire a new paradigm for sampling
- Sampling provides insights into the design of event-driven systems (end-to-end learning)
- Model-based deep learning leads to lighter and more universal architectures

Imperial College London

Two-Photon Microscopy for Neuroscience

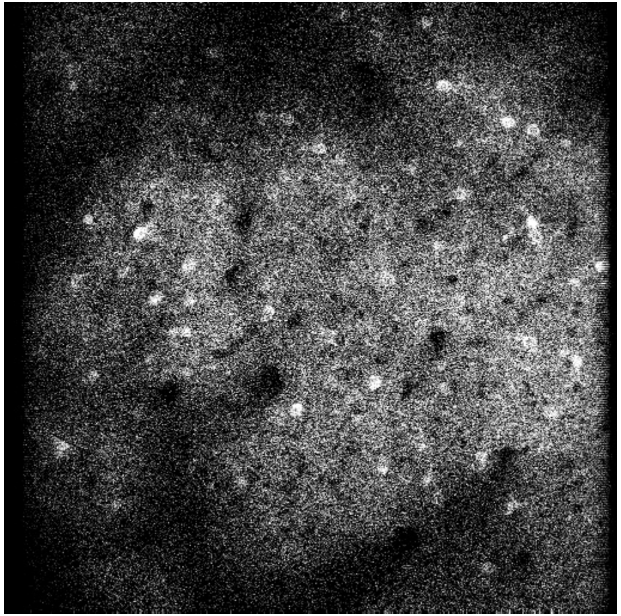

- Goal of Neuroscience: to study how information is processed in the brain
- Neurons communicate through pulses called Action Potentials (AP)
- Need to measure in-vivo the activity of large populations of neurons at cellular level resolution
- Two-photon microscopy combined with right indicators is the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

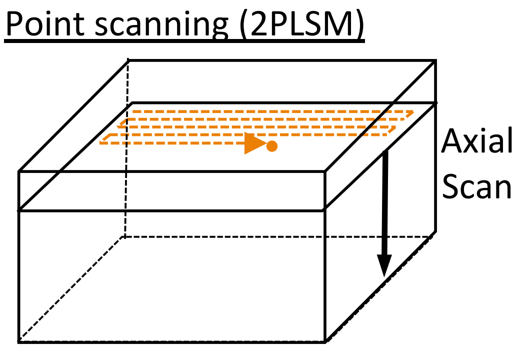

Imperial College London

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

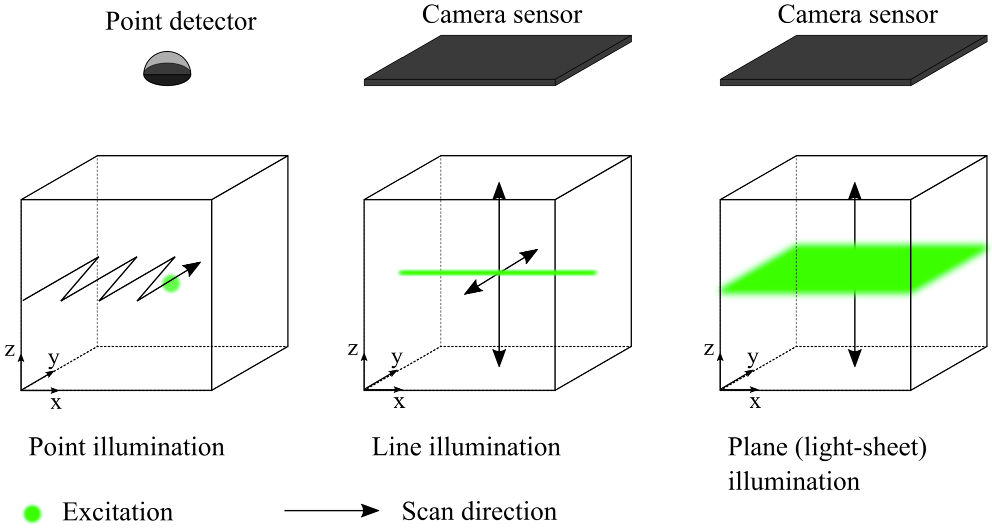

Imperial College London

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

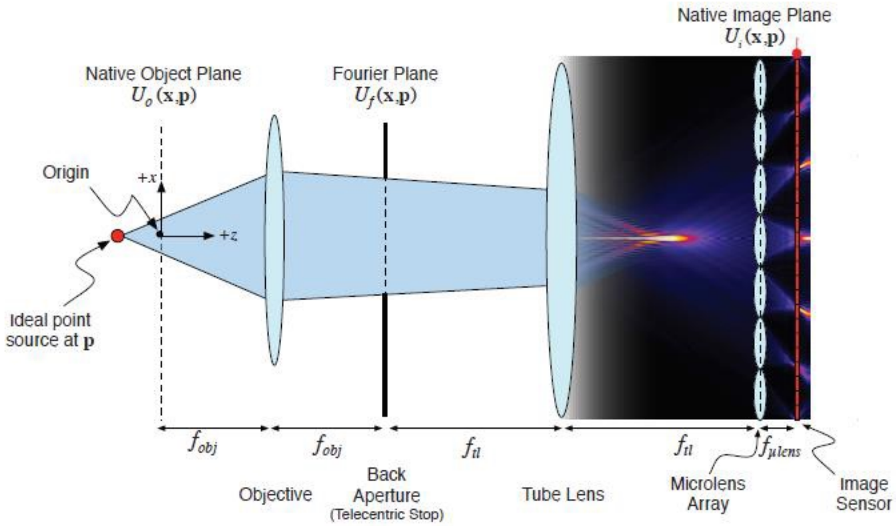

Imperial College London

Light-field Microscopy and EPI

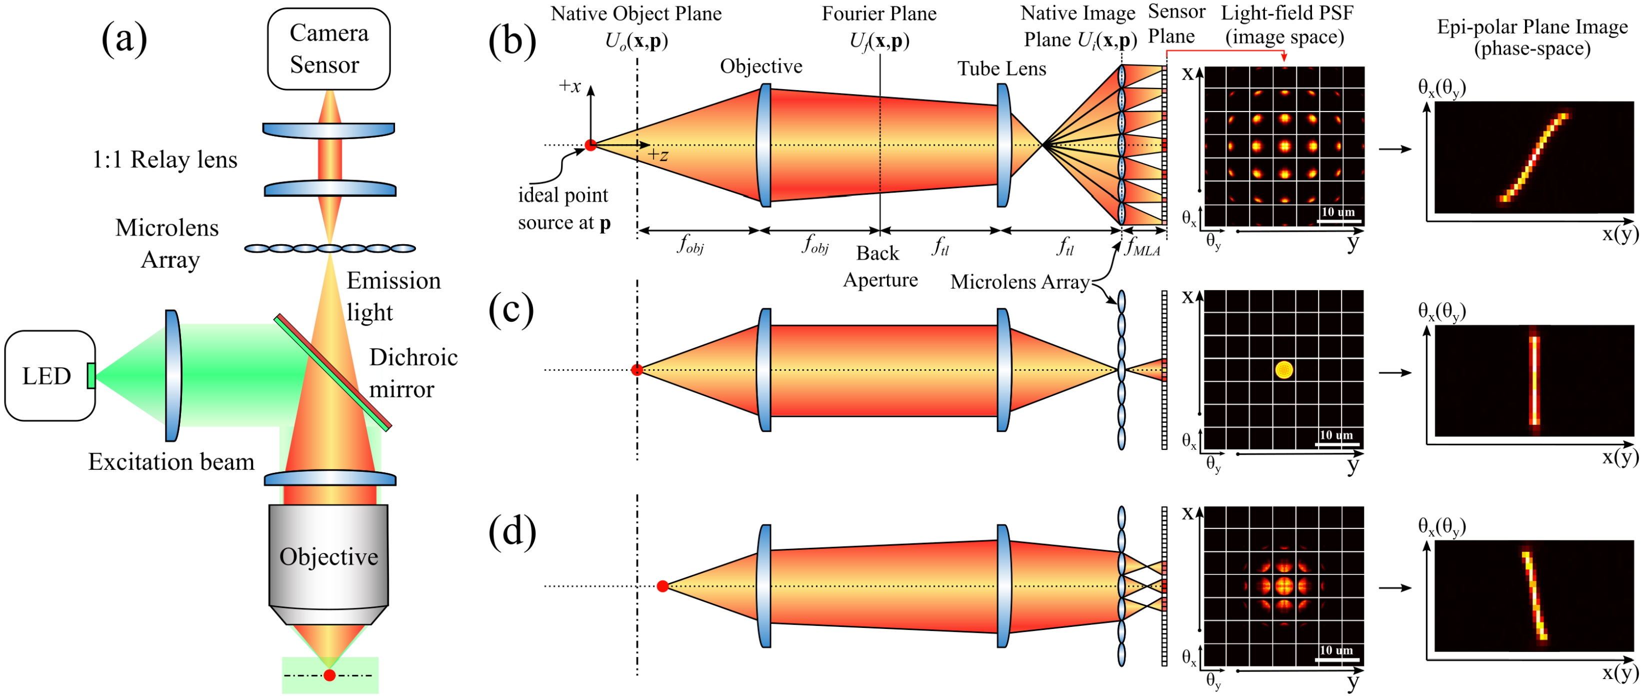

Imperial College London

Light-field Microscopy and Illumination Strategies

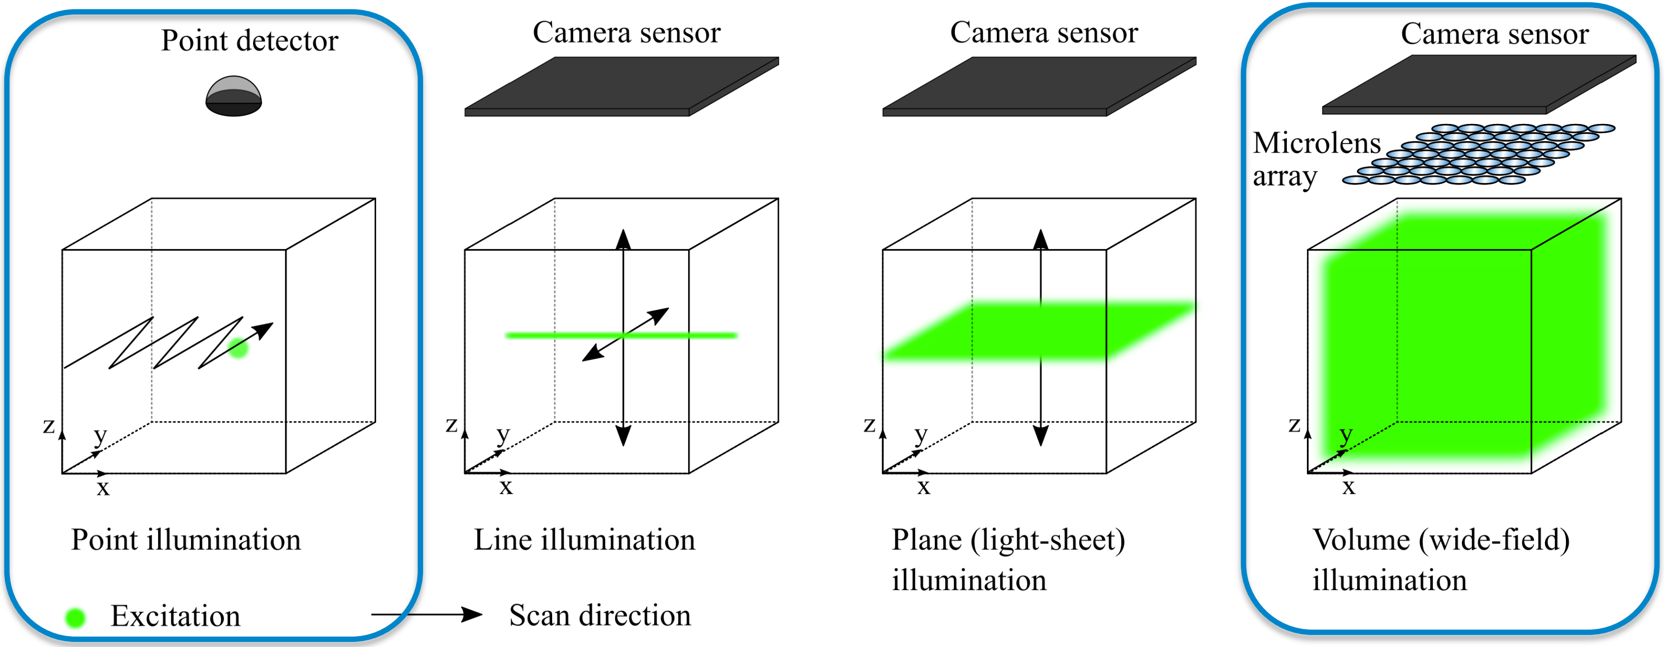

Key insight: use the 2P microscope for high-resolution structural information and the LFM for monitoring the activity of neurons.

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)

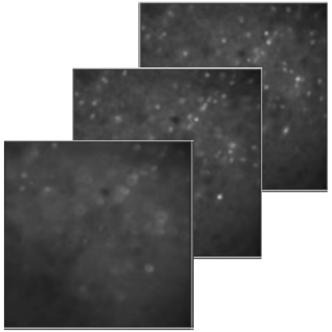

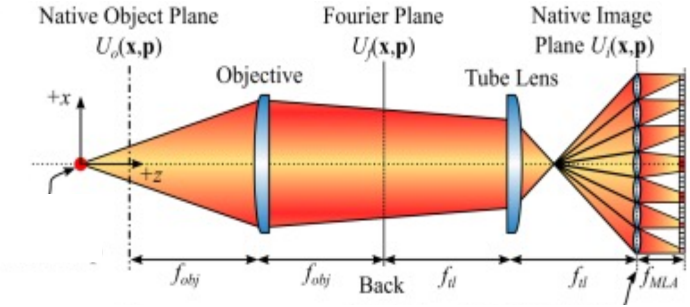

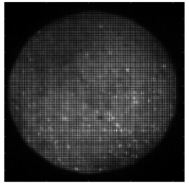

2D Measured LF image

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

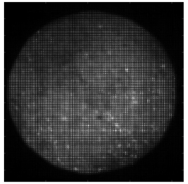

2-D LF

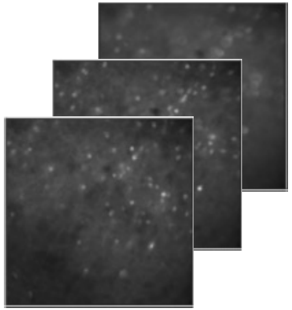

Volume

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)

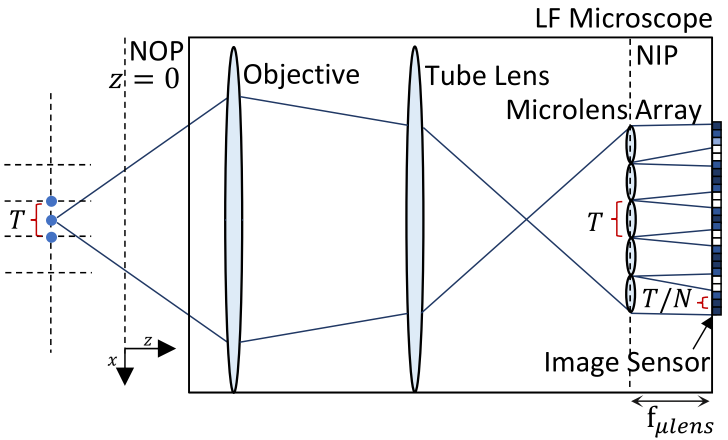

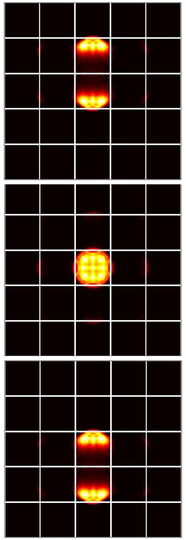

Imperial College London

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

Neural network for volume reconstruction

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

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

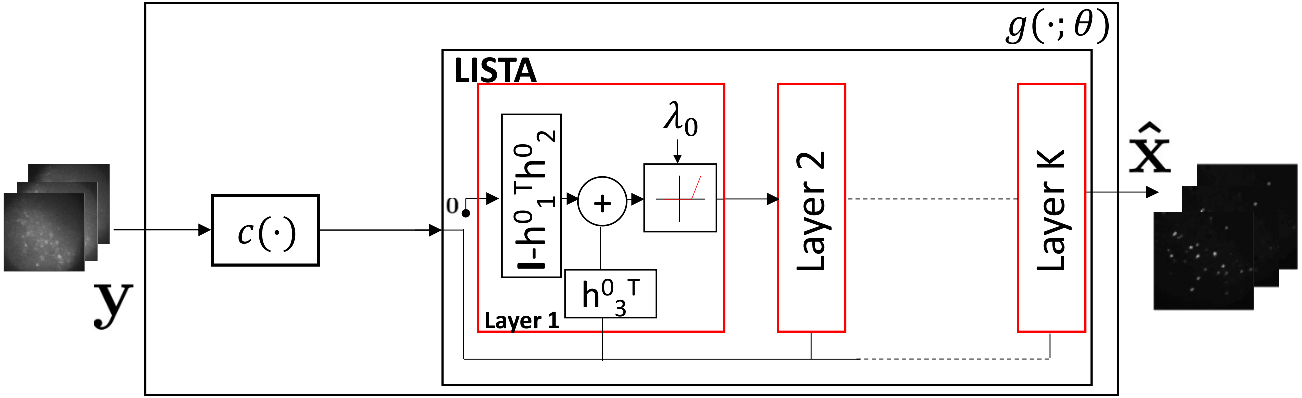

Training of the neural network

- 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

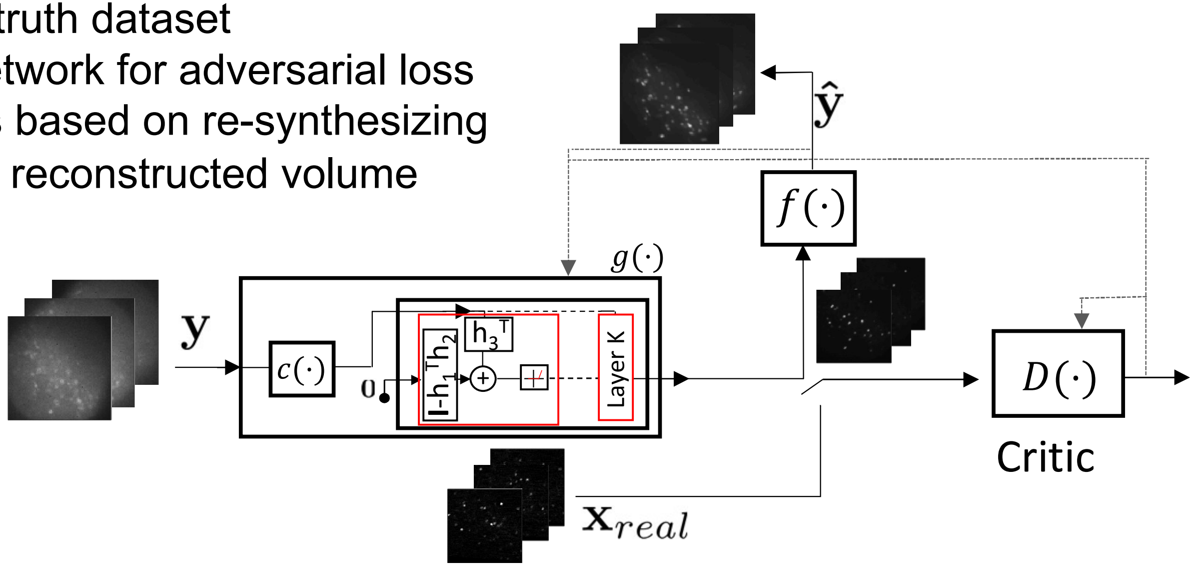

Training of the neural network

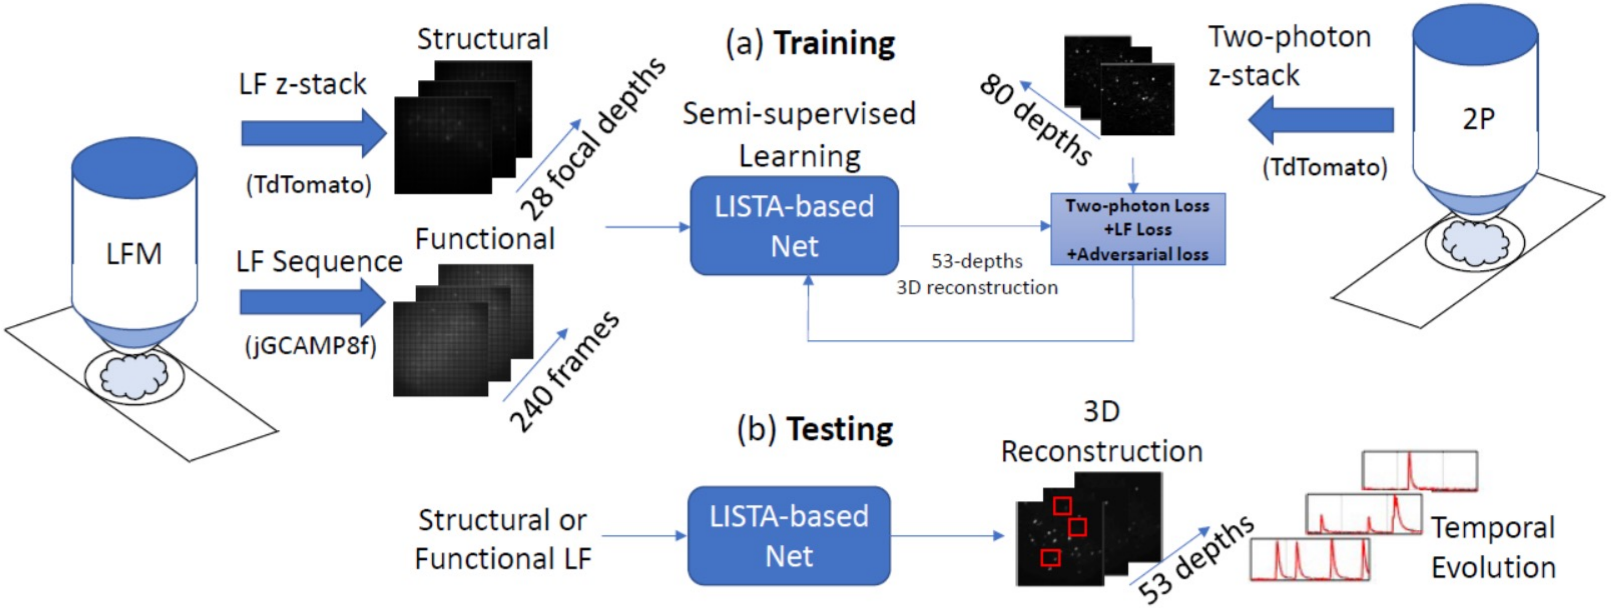

Results – Structural Data

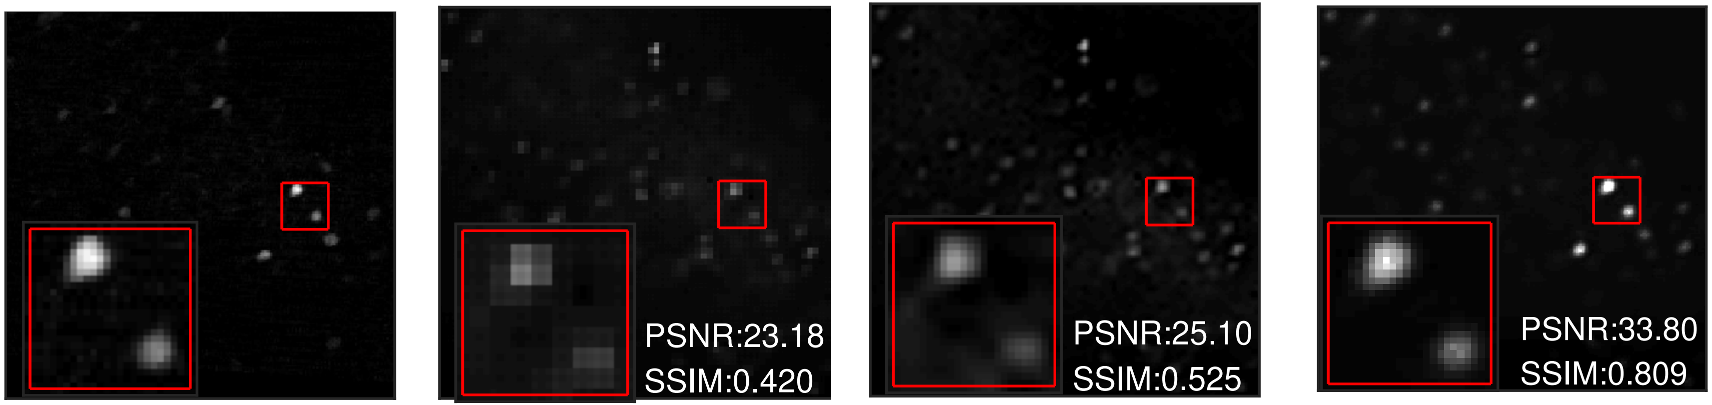

Ground-truth ISRA ADMM New method (0.3s to reconstruct one volume)

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

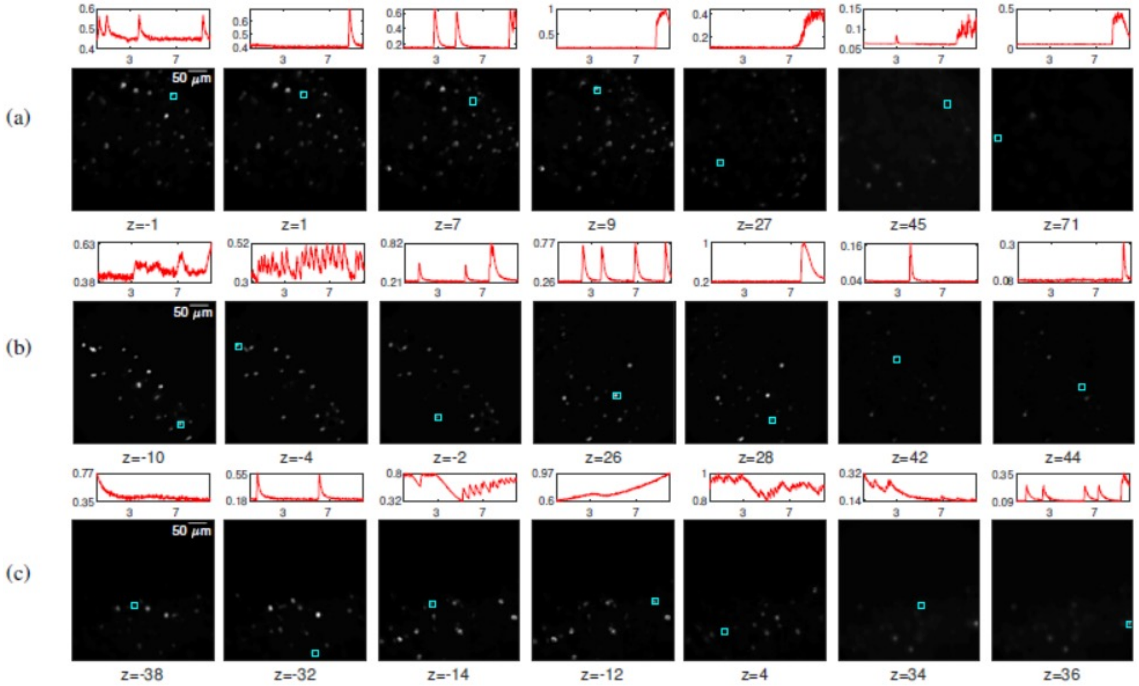

Three brain samples are shown in parts (a), (b), and (c)

Imperial College London

Technical Examination of Paintings

Imperial College London

Structure of a painting

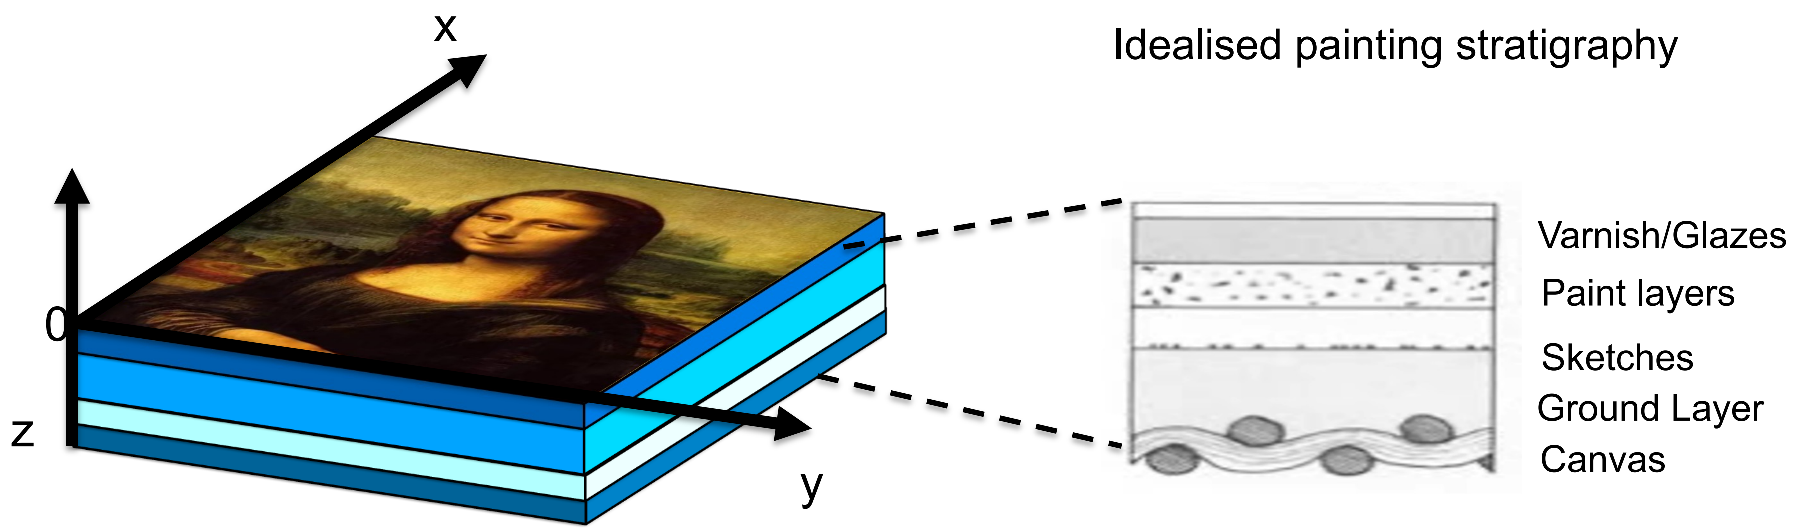

Imperial College London

Art-Investigation

Visible

X-ray

Imperial College London

Demixing using a connected auto-encoder

Machine Learning to extract painting underneath[1]

(a)

(b)

Francisco de Goya, Dona Isabel de Porcel (NG1473), before 1805. Oil on canvas. (a). RGB image. (b). X-ray image.

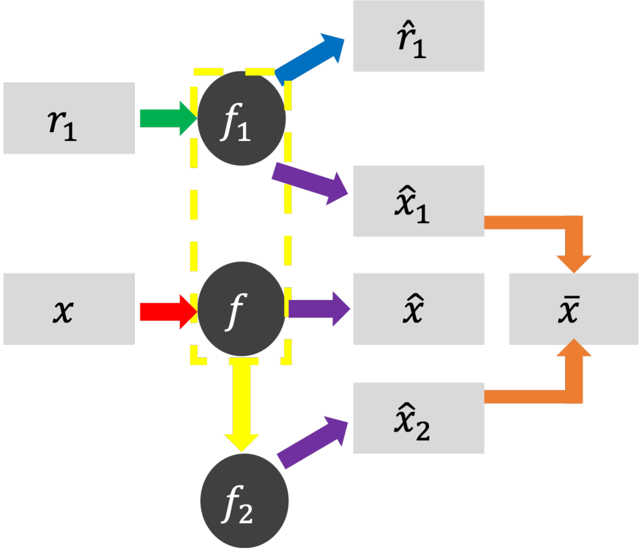

[1] W. Pu, J. Huang, B. Sober, N. Daly, C. Higgitt, I. Daubechies, P.L. Dragotti and M. Rodrigues, "Mixed X-Ray Image Separation for Artworks With Concealed Designs", IEEE Trans. on Image Processing, 2022.

Imperial College London

Separation Results

Machine Learning to extract painting underneath

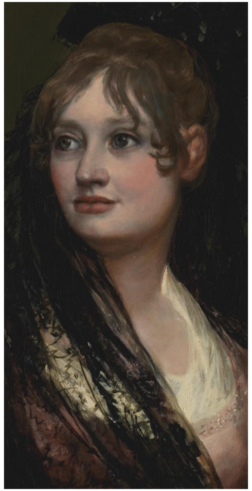

Separation Results

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

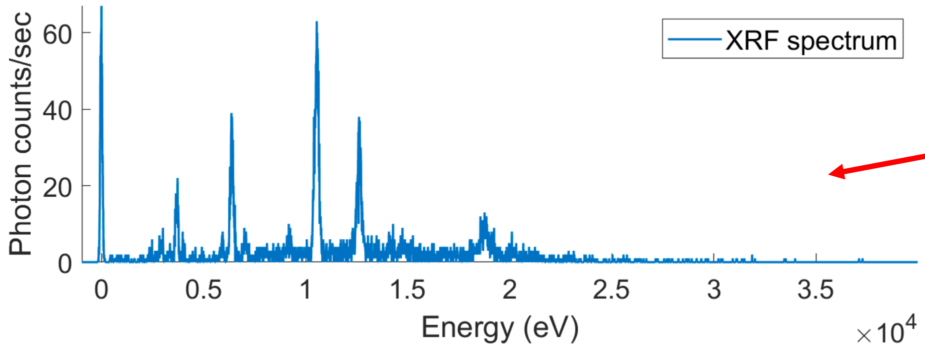

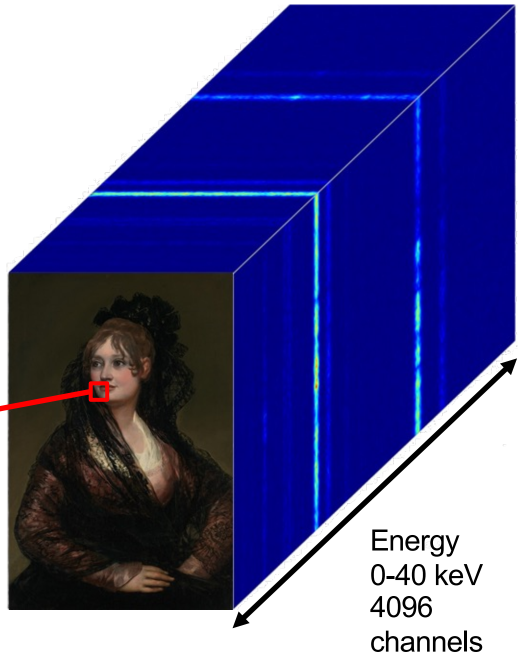

Imperial College London

Extraction of Elemental Maps

Our XRF

Deconvolution

Algorithm

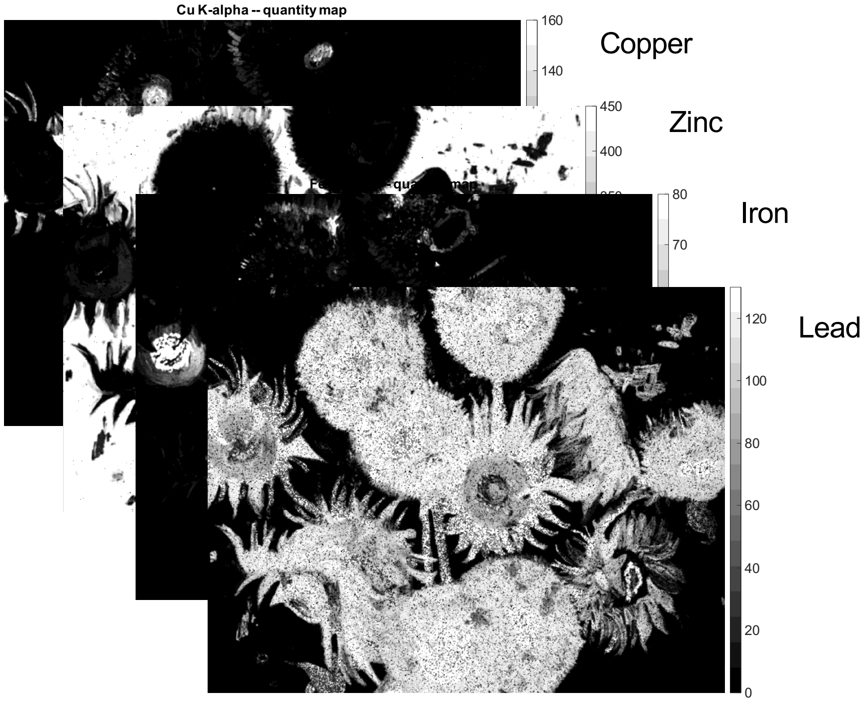

Vincent van Gogh, “Sunflowers (NG3863)”, © The National Gallery, London.

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.

Results

Copper (Cu) distribution maps

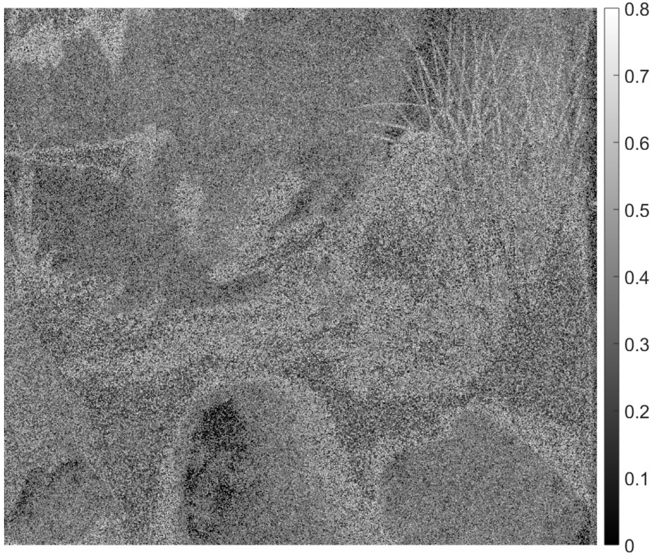

Cu confidence map

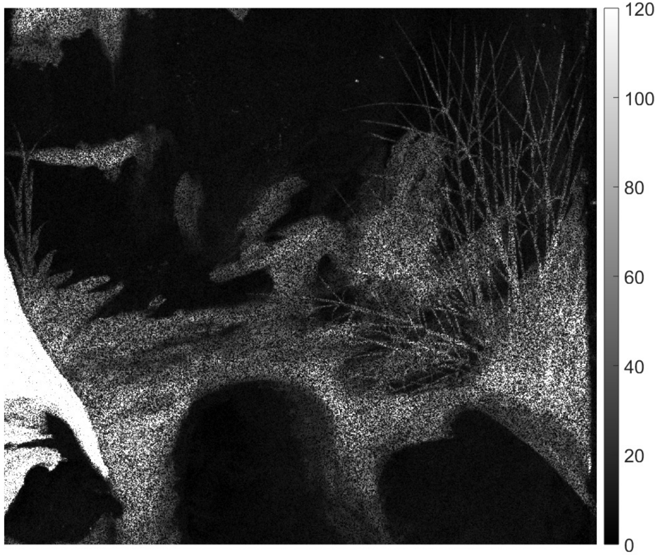

Cu quantity map

Imperial College London

Results

Zinc (Zn) distribution maps

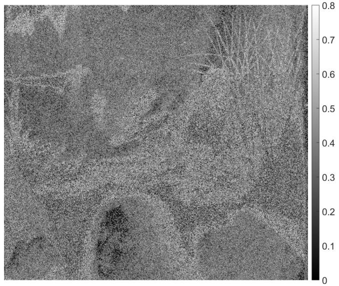

Cu confidence map

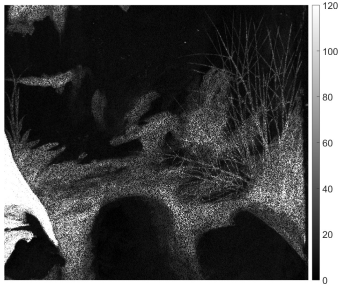

Cu quantity map

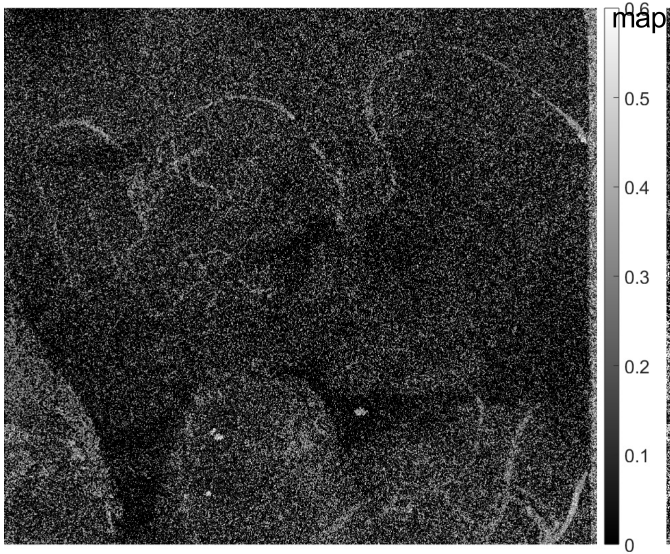

Zn confidence map

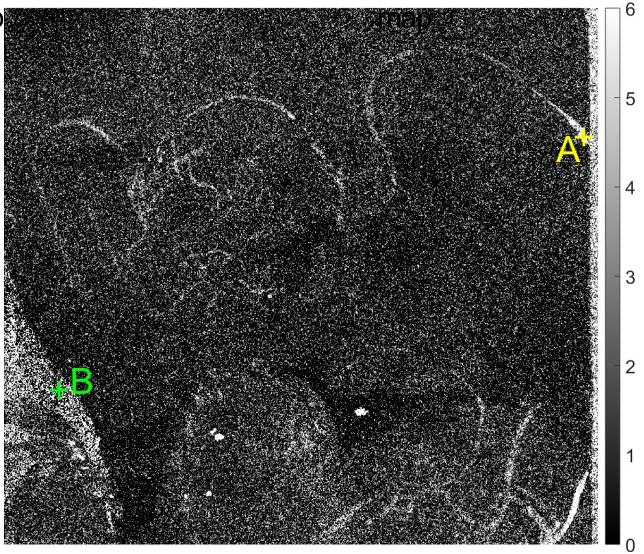

Zn quantity map

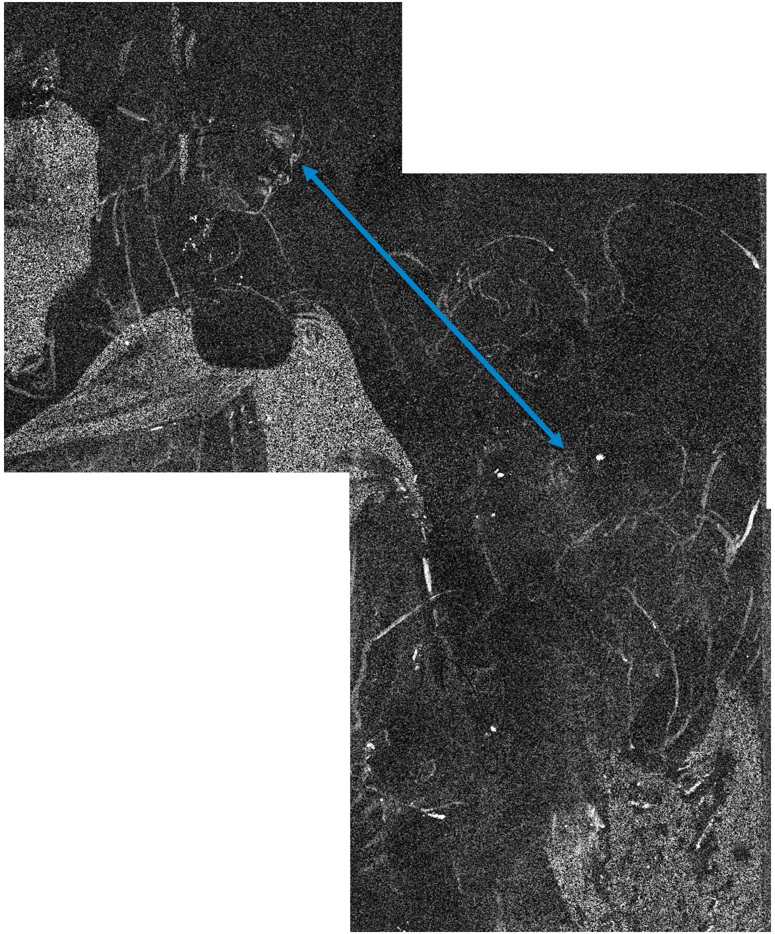

Imperial College London

Conclusions

- Why Computational Imaging?
  - It is fun 🙂
  - It is inter-disciplinary
  - It is the right way to handle ‘big data’: joint sensing, representation, analysis and inference
  - Models and priors can reduce complexity and lead to better results
  - Some computational approaches are transfearable

Thank you!

Imperial College London

References

On Sparse Sampling

- J. Uriguen, T. Blu, and P.L. Dragotti ‘FRI Sampling with Arbitrary Kernels’, IEEE Trans. on Signal Processing, November 2013

On time-based Sampling and event-driven cameras

- R. Alexandru and P. L. Dragotti, "Reconstructing classes of non-bandlimited signals from time encoded information", IEEE Transactions on Signal Processing, Vol.68, pp. 747-763, Year 2020
- S. Liu and P.L. Dragotti, Sensing Diversity and Sparsity Models for Event Generation and Video Reconstruction from Events, IEEE Trans. on Pattern Recognition and Machine Intelligence, 2023

Imperial College London

References

On Art Investigation

- 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.
- W. Pu, J. Huang, B. Sober, N. Daly, C. Higgitt, I. Daubechies, P.L. Dragotti and M. Rodrigues, “Mixed X-Ray Image Separation for Artworks With Concealed Designs”, IEEE Trans. on Image Processing, 2022.
- Y. Su et al. "A Fast Automatic Method for Deconvoluting Macro X-ray Fluorescence Data Collected from Easel Painting", IEEE Transactions on Computational Imaging, 2023.

On Light-Field Microscopy and Neuroscience

- P. Song P, H. Verinaz Jadan, C. Howe, A. Foust and P.L. Dragotti, Light-Field Microscopy for optical imaging of neuronal activity: when model-based methods meet data-driven approaches, IEEE Signal Processing Magazine, March 2022.
- 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.