Memory-efficient reconstruction for 3D sparse-view X-ray CT

Romain Vo

POPILSS 2025

X-Ray Tomography

X-Ray Computed Tomography

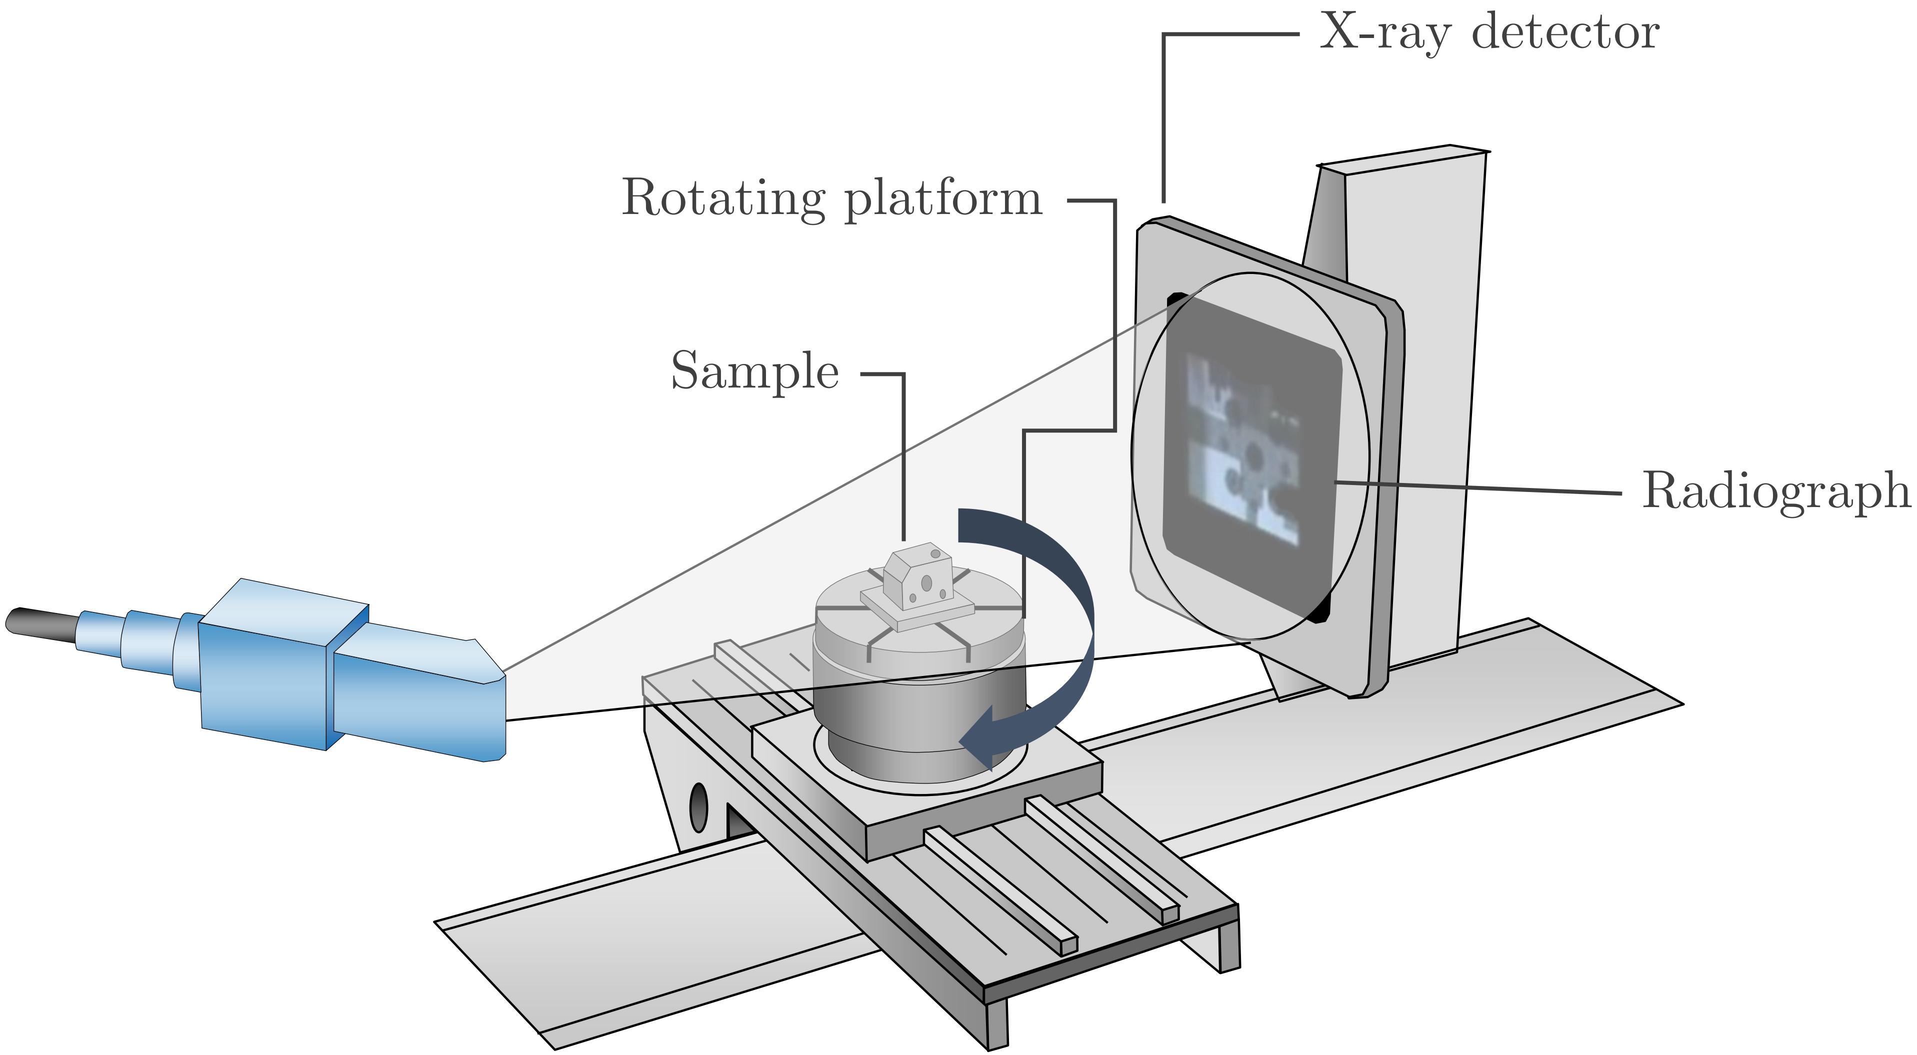

[Liu 2023]

X-Ray Computed Tomography

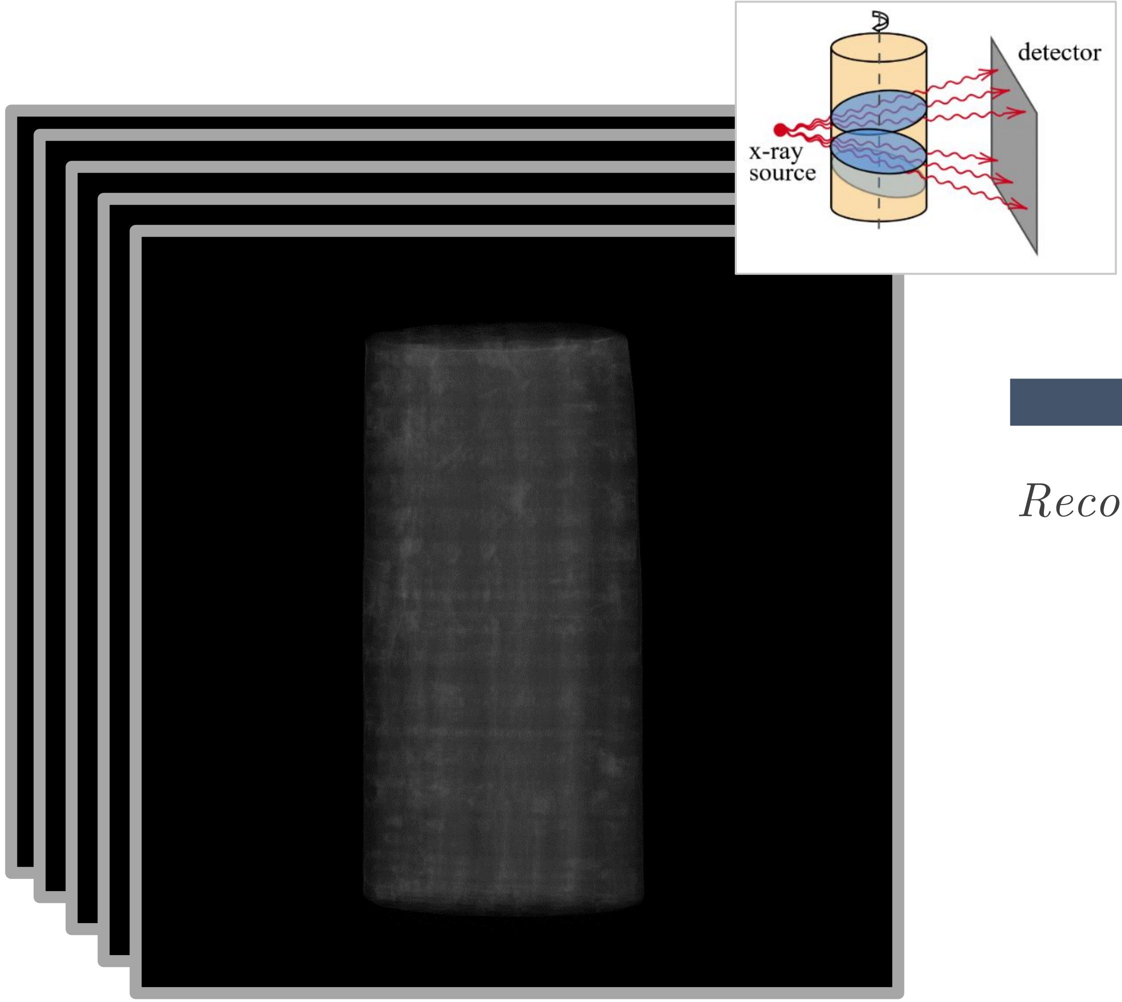

X-ray Radiographs (2D)

Reconstruction

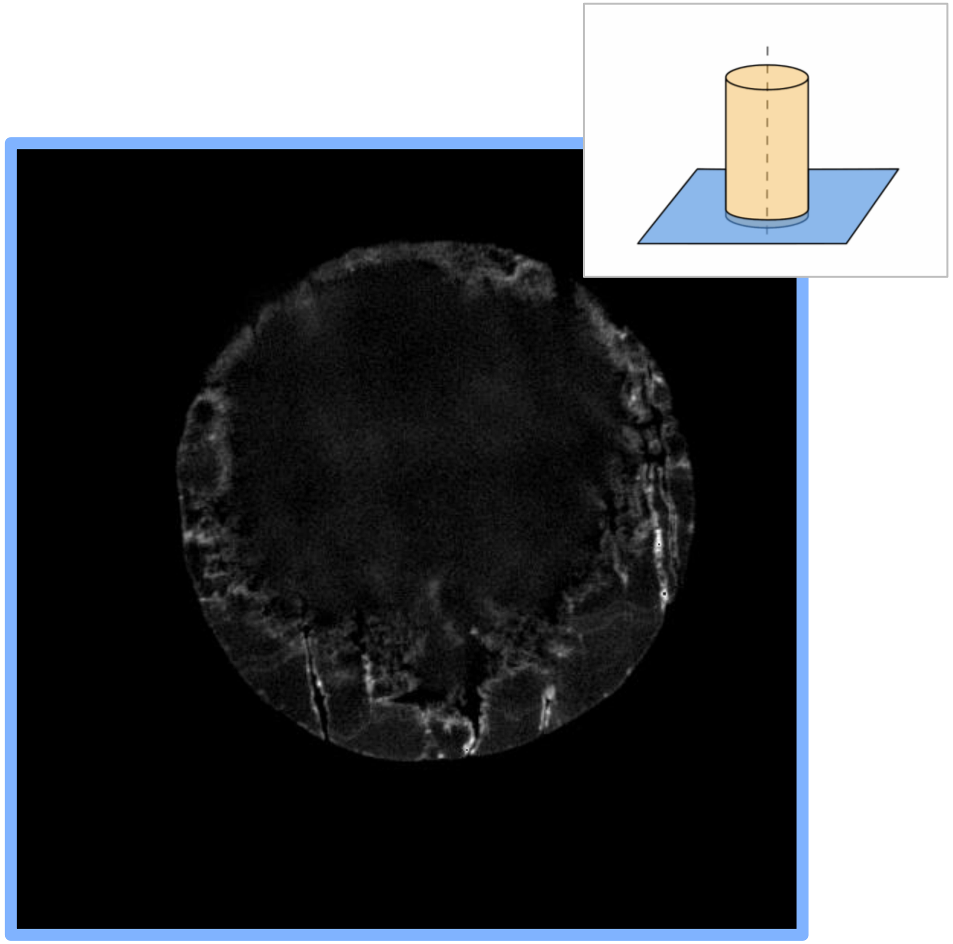

Slices of the 3D attenuation map

CT inspection – online monitoring

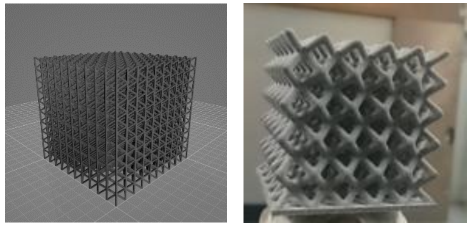

Design / Process → Online production → Deployment

Online production ⇄ CT Inspection

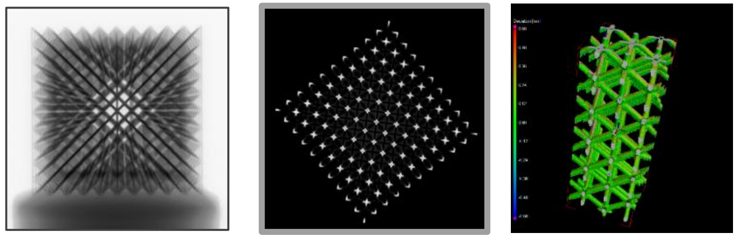

Very slow ≈ hour(s)

Speeding up CT

Application

→ sparse-view CT

Decrease the number of angular **measurements** to reduce the acquisition time

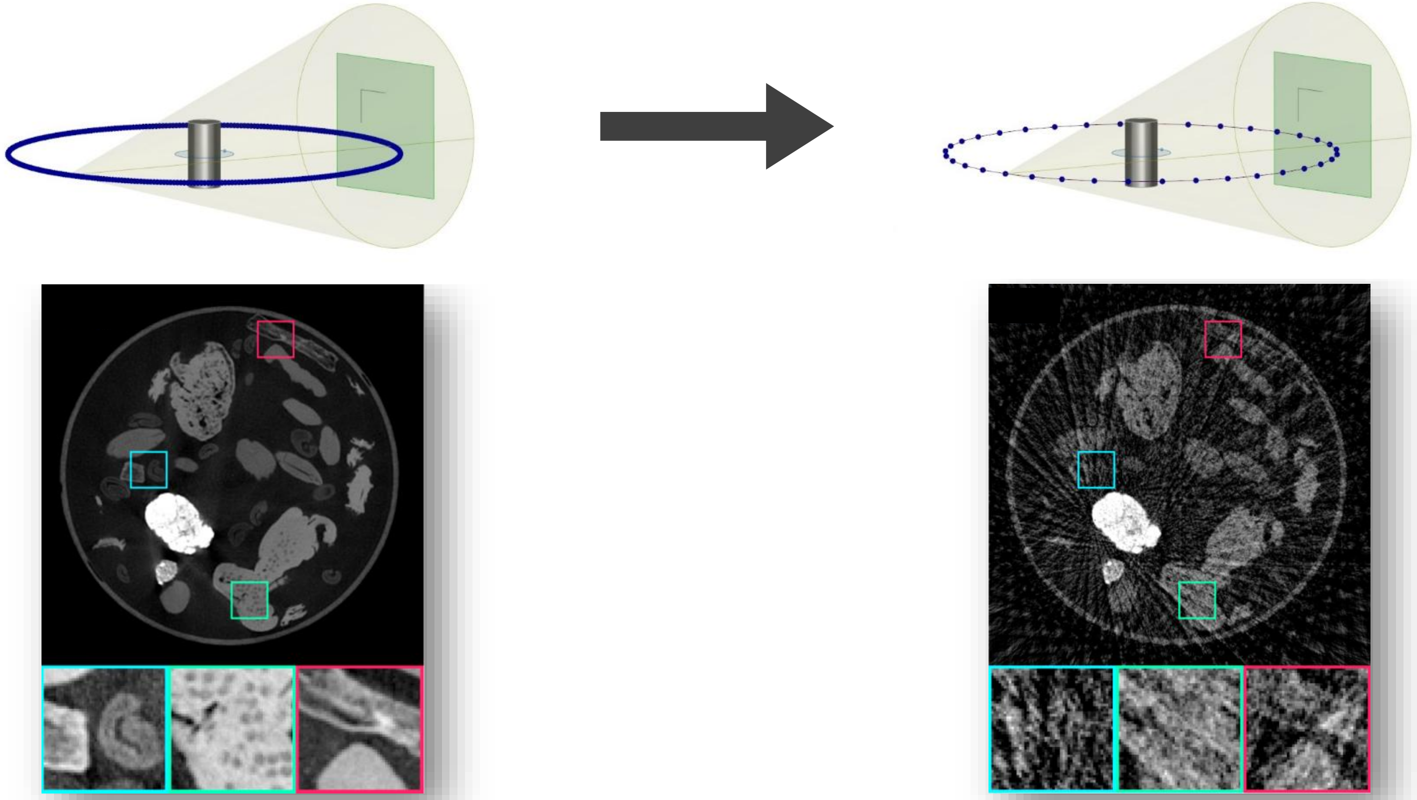

2DeteCT - [Kiss 2024]

Fine-grained inspection

Application

→ sparse-view CT

Decrease the number of angular **measurements** to reduce the acquisition time

> *Reducing the number of measurements leads to **artifacts** in the reconstruction*

→ 3D high-resolution inspection

We typically deal with **high-resoluted 3D object**, to be able to observe and control fine-grained details of the part we inspect

- Volume $x \in \mathbb{R}^n$ $n = 1024^3$ voxels (= 4 GigaBytes in float32)
- Radiographs $b \in \mathbb{R}^m = \mathrm{p} \times 1024^2$ (< 1 GB depending on p the number of *angular positions*)

2DeteCT **[2D]**
512^2

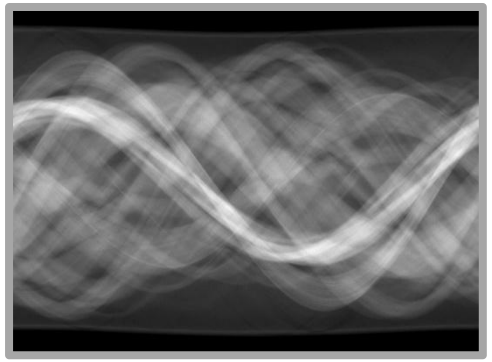
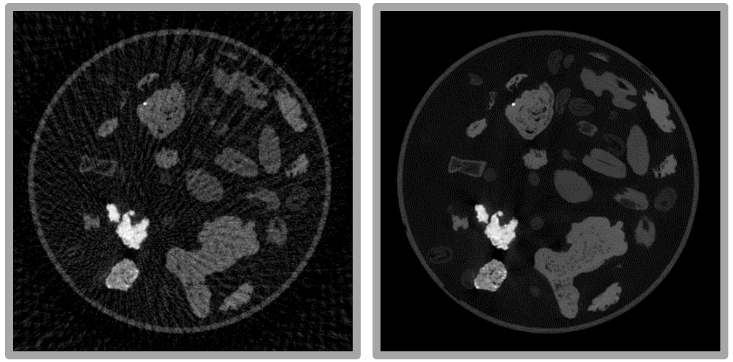
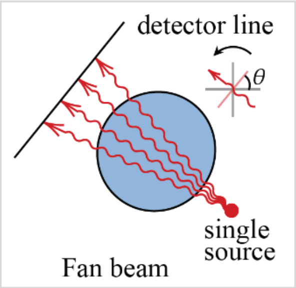

Cork-CBCT **[3D]**
1024^3

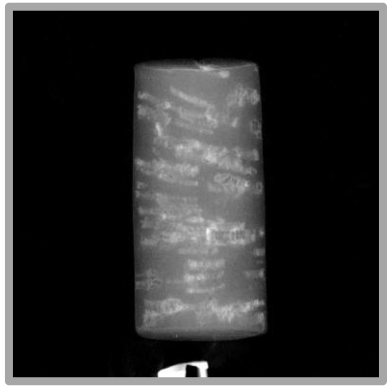
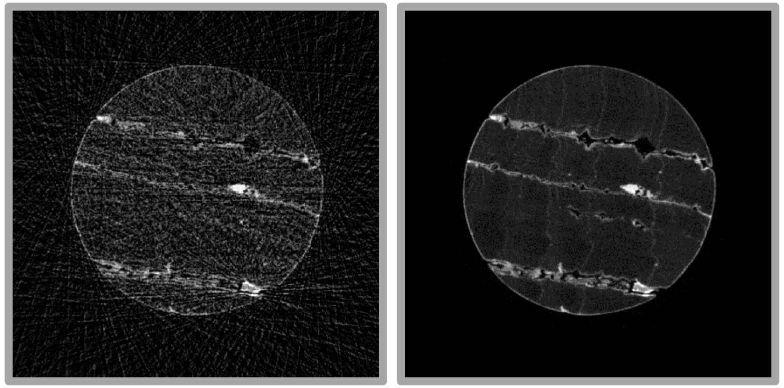
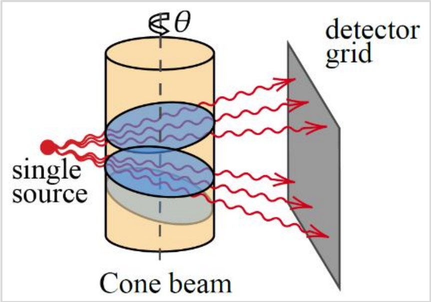

Walnut-CBCT **[3D]**
501^3

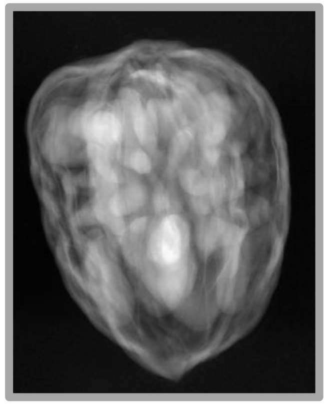
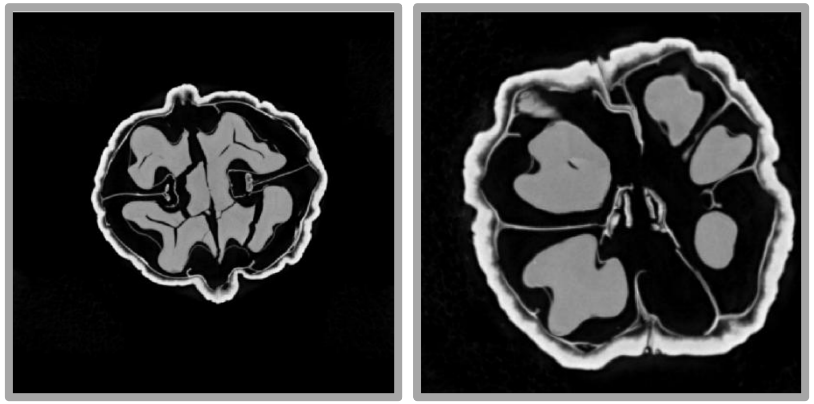
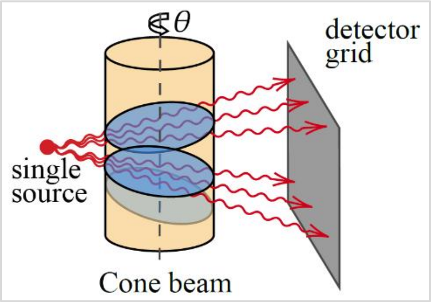

ENS DE LYON

Linear Inverse Problem

Inverse problem

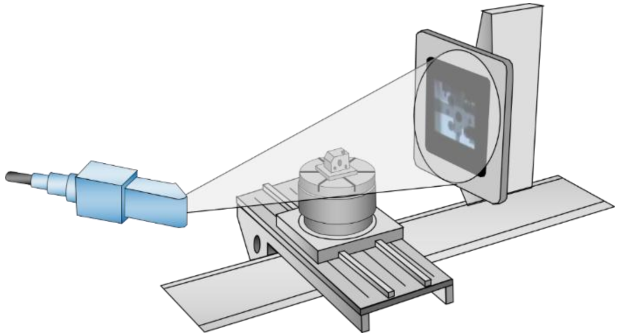

Forward model

The CT reconstruction problem can be formulated as solving a system of linear equations, **we want to find** x^* such that :

$$Ax^* + \varepsilon = b, \quad A \in \mathbb{R}^{m\times n}, x^* \in \mathbb{R}^n, b \in \mathbb{R}^m, \varepsilon \in \mathbb{R}^m$$

Forward operator

Signal (Object)

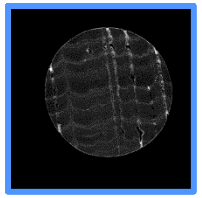

Noise

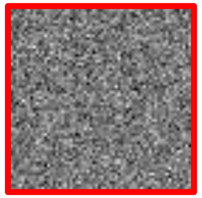

Measurements (Radiographs)

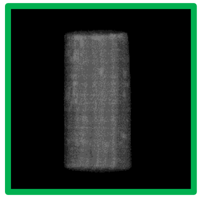

Ill-posed inverse problem

Forward model

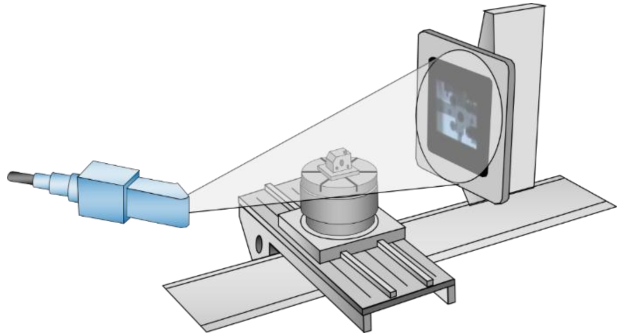

The CT reconstruction problem can be formulated as solving a system of linear equations, **we want to find** x^* such that :

$$Ax^* + \varepsilon = b, \quad A \in \mathbb{R}^{m \times n}, x^* \in \mathbb{R}^n, b \in \mathbb{R}^m, \varepsilon \in \mathbb{R}^m$$

Sparse-view CT → *The problem is* ***ill-posed :***

the number of measurements $m \ll$ *size of the volume* n

Standard methods

Direct/Analytical Inversion

$$\tilde{x} = A^{+} b$$

→ Fast estimation[1] using

- *Filtered Back Projection* (FBP)
- *Feldkamp,Davis,Kress* (FDK) [Feldkamp1984]

[1]Anastasio et al. 2001. 'Comments on the Filtered Backprojection Algorithm, Range Conditions, and the Pseudoinverse Solution'. *IEEE Transactions on Medical Imaging*

Iterative Reconstruction

$$\underset{x \in \mathbb{R}^n}{\arg\min} \; \underbrace{\frac{1}{2} \|Ax - b\|_2^2}_{\text{Data-fidelity}} + \underbrace{\lambda \mathcal{R}(x)}_{\text{Regularization}}$$

Data-fidelity

Regularization

Proximal Gradient Descent

$$\begin{cases} z_{k+1} = x_k - \eta A^\top (Ax_k - b) \\ x_{k+1} = \mathrm{prox}_{\eta\lambda\mathcal{R}}(z_{k+1}) \end{cases} \quad \text{(PGD)}$$

[Combettes & Pesquet 2010]

Standard methods

Direct/Analytical Inversion

$$\tilde{x} = A^{+}b$$

→ Fast estimation using

- *Filtered Back Projection* (FBP)
- *Feldkamp,Davis,Kress* (FDK) [Feldkamp1984]

Iterative Reconstruction

$$\underset{x \in \mathbb{R}^n}{\arg\min} \; \underbrace{\mathcal{D}(x)}_{\text{Data-fidelity}} + \lambda \underbrace{\mathcal{R}(x)}_{\text{Regularization}}$$

Proximal Gradient Descent

$$\begin{cases} z_{k+1} = x_k - \eta A^\top (Ax_k - b) \\ x_{k+1} = \mathrm{prox}_{\eta\lambda\mathcal{R}}(z_{k+1}) \end{cases} \quad \text{(PGD)}$$

[Combettes & Pesquet 2010]

Standard methods

Direct/Analytical Inversion

$$\tilde{x} = A^{+}b$$

→ Fast estimation using

- *Filtered Back Projection* (FBP)
- *Feldkamp,Davis,Kress* (FDK) [Feldkamp1984]

Iterative Reconstruction

$$\arg\min_{x \in \mathbb{R}^n} \; \mathcal{D}(x) + \lambda \mathcal{R}(x)$$

➡ *Total Variation:* $\mathcal{R}(x) = \|\nabla x\|_1$

Standard methods

Direct/Analytical Inversion

$$\tilde{x} = A^{+} b$$

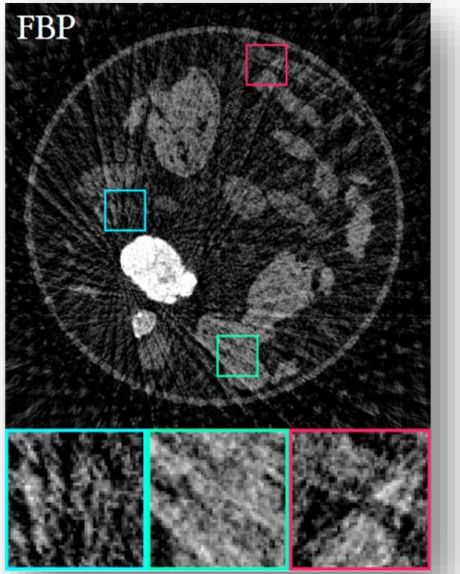

Iterative Reconstruction

$$\arg\min_{x \in \mathbb{R}^n} \; \mathcal{D}(x) + \lambda \mathcal{R}(x)$$

→ *Total Variation:* $\mathcal{R}(x) = \|\nabla x\|_1$

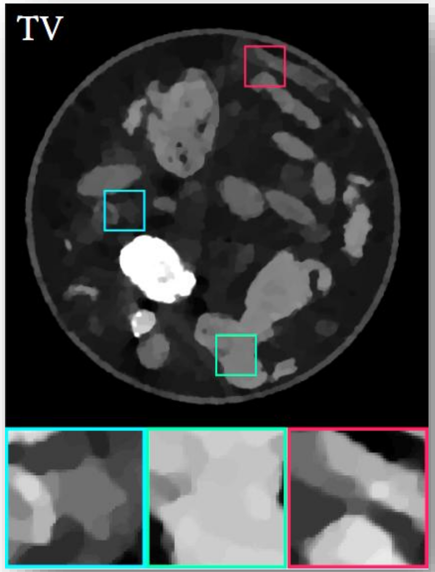

[Sidky & Pan 2008]

Standard methods - Limitations

Direct/Analytical Inversion

$$\tilde{x} = A^{+}b$$

- *Fast but artifacts in sparse setup*

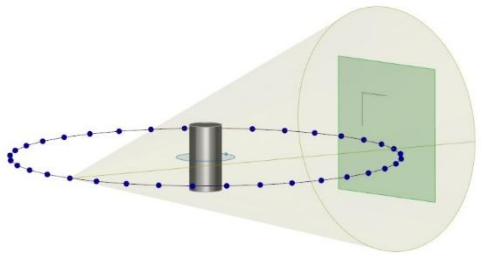

Iterative Reconstruction

$$\arg\min_{x \in \mathbb{R}^n} \; \mathcal{D}(x) + \lambda \mathcal{R}(x)$$

- *Reduced artifacts but slow*
- *Explicit formulation*
- *Hand-crafted regularization*

Deep Learned Reconstruction

Model-agnostic formulation

Find the function $f_\phi : \mathbb{R}^m \to \mathbb{R}^n$ *such that :*

$$x = f_\phi(b)$$

Neural Network

Deep Learned Reconstruction

Find the function $f_\phi : \mathbb{R}^n \to \mathbb{R}^n$ *such that :*

$$x = f_\phi(\tilde{x}), \quad \tilde{x} = A^+ b$$

image-translation mapping

Deep Learned Reconstruction

Given a collection of **degraded** and **reference** reconstruction pairs $\{(\tilde{x}_i, x_i^*)\}_{i=1}^N$

Training : $$\boxed{\min_{\phi} \; \mathbb{E}_{(\tilde{\mathrm{x}},\mathrm{x}^*)} \| f_\phi(\tilde{\mathrm{x}}) - \mathrm{x}^* \|_2^2}$$

Parametrization ?

Post-processing

Post-processing

D_ϕ : ***single-pass*** convolutional network (U-Net)

[Ronneberger 2015]

$$\min_\phi \ \mathbb{E}_{(\tilde{\mathrm{x}},\mathrm{x}^*)} \| D_\phi(\tilde{\mathrm{x}}) - \mathrm{x}^* \|_2^2$$

[Han 2016]

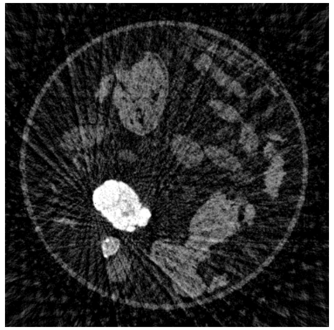

D_ϕ

remove artifacts

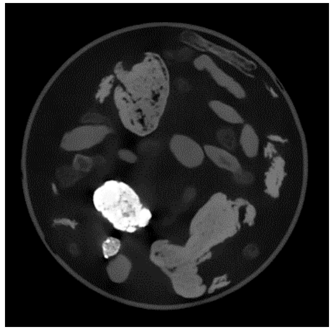

$\tilde{x} = A^+ b$

$\hat{x} = D_\phi(\tilde{x})$

Unrolled optimization

Post-processing

Post-processing

D_ϕ : ***single-pass*** convolutional network (U-Net)

[Ronneberger 2015]

$$\min_{\phi} \ \mathbb{E}_{(\tilde{\mathrm{x}}, \mathrm{x}^*)} \|D_\phi(\tilde{\mathrm{x}}) - \mathrm{x}^*\|_2^2$$

[Han 2016]

→ What about minimizing $\|Ax - b\|_2^2$?

Unrolled optimization

Unrolled network

Post-processing

D_ϕ : ***single-pass*** convolutional network (U-Net)

[Ronneberger 2015]

$$\min_{\phi} \; \mathbb{E}_{(\tilde{\mathrm{x}},\mathrm{x}^*)} \|D_\phi(\tilde{\mathrm{x}}) - \mathrm{x}^*\|_2^2$$

[Han 2016]

→ What about minimizing $\|Ax - b\|_2^2$?

Unrolled optimization

$$\begin{cases} z_{k+1} = x_k - \eta \nabla \mathcal{D}(x) \\ x_{k+1} = \mathrm{prox}_{\eta\lambda\mathcal{R}}(z_{k+1}) \end{cases}$$

hand-crafted

Unrolled network

Post-processing

D_ϕ : ***single-pass*** convolutional network (U-Net)

[Ronneberger 2015]

$$\min_\phi \ \mathbb{E}_{(\tilde{\mathrm{x}},\mathrm{x}^*)} \|D_\phi(\tilde{\mathrm{x}}) - \mathrm{x}^*\|_2^2$$

[Han 2016]

→ What about minimizing $\|Ax - b\|_2^2$?

Unrolled optimization

$$\begin{cases} z_{k+1} = x_k - \eta \nabla \mathcal{D}(x) \\ x_{k+1} = D_\phi(z_{k+1}) \end{cases}$$

learned

Unrolled network w. parameters ϕ

$$G_\phi(\tilde{x}, K) = [D_\phi \circ g \circ \cdots \circ D_\phi \circ g] (\tilde{x})$$

K steps

with $g(x) = x - \eta \nabla \mathcal{D}(x)$

Unrolled network

Post-processing

D_ϕ : ***single-pass*** convolutional network (U-Net)

[Ronneberger 2015]

$$\min_{\phi} \; \mathbb{E}_{(\tilde{\mathrm{x}},\mathrm{x}^*)} \| D_\phi(\tilde{\mathrm{x}}) - \mathrm{x}^* \|_2^2$$

[Han 2016]

Unrolled optimization

Unrolled network w. parameters $\boldsymbol{\phi}$

$$G_\phi(\tilde{x}, K) = [D_\phi \circ g \circ \cdots \circ D_\phi \circ g]\,(\tilde{x})$$

$$\min_{\phi} \; \mathbb{E}_{(\tilde{\mathrm{x}},\mathrm{x}^*)} \| G_\phi(\tilde{\mathrm{x}}, K) - \mathrm{x}^* \|_2^2$$

[Adler 2017]

« End-to-end iterative training »

Deep Reconstruction - Qualitative results

Post-processing

D_ϕ : ***single-pass*** convolutional network (U-Net)

[Ronneberger 2015]

Unrolled optimization

Unrolled network w. parameters ϕ

$$G_\phi(\tilde{x}, K) = [D_\phi \circ g \circ \cdots \circ D_\phi \circ g](\tilde{x})$$

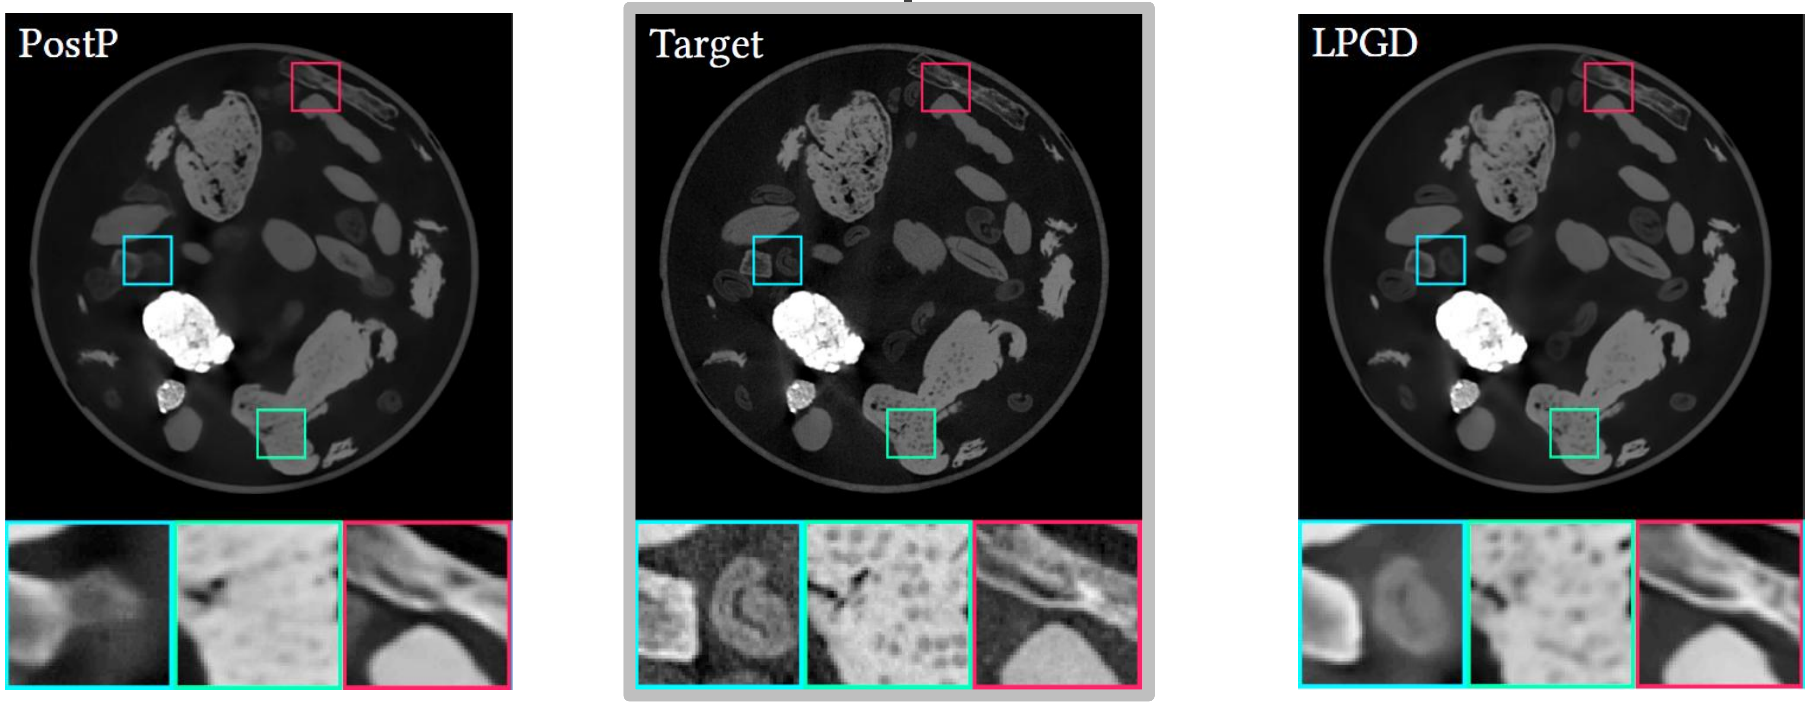

Deep Reconstruction - Recap

Post-processing

D_ϕ : ***single-pass*** convolutional network (U-Net)

- Fast but does not ensure $AD_\phi(\tilde{x}) = b$
- Reduced performances in out-of-distribution settings [Vo 2025]

Unrolled optimization

Unrolled network w. parameters $\boldsymbol{\phi}$

$$G_\phi(\tilde{x}, K) = [D_\phi \circ g \circ \cdots \circ D_\phi \circ g](\tilde{x})$$

- State-of-the-art since 2017 on every benchmarks
- « Interpretable »

Limitations

- Limited scalability for high-resolution 2D and 3D
  - 2DeteCT (512×512) – Training = 57GB VRAM & 30mn/epoch

Alternative ?

$$G_\phi(\tilde{x}, K) = [D_\phi \circ g \circ \cdots \circ D_\phi \circ g](\tilde{x})$$

K steps

1. *Plug & Play* (**PnP**) [Venkatakrishnan 2013]

pre-trained ***denoiser***

$$\min_\phi \ \mathbb{E}_{\mathrm{x}^*} \mathbb{E}_{\xi_\sigma \sim \mathcal{N}(0,\sigma^2 I)} \| D_\phi(\mathrm{x}^* + \xi_\sigma) - \mathrm{x}^* \|_2^2$$

[Zhang 2017]

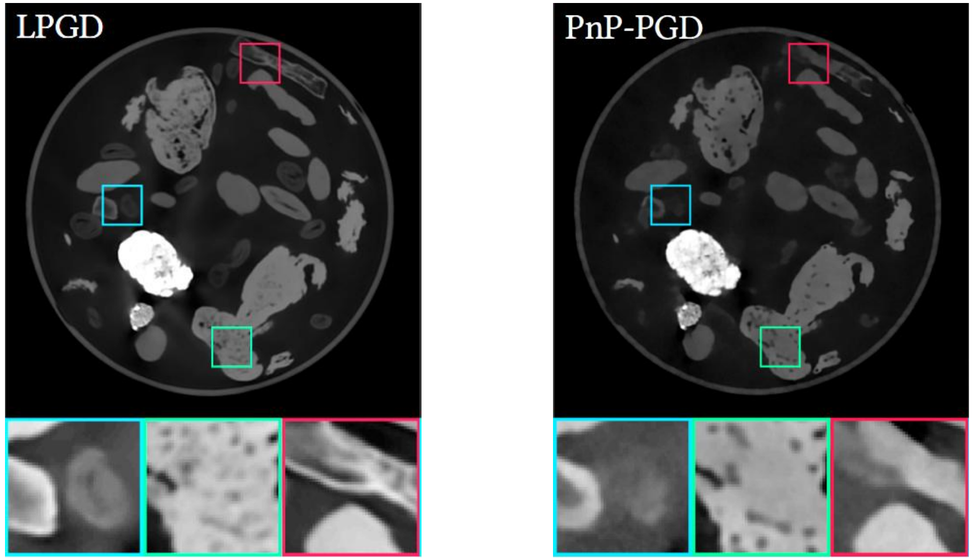

Alternative ?

$$G_\phi(\tilde{x}, K) = \underbrace{[D_\phi \circ g \circ \cdots \circ D_\phi \circ g]}_{K \text{ steps}}(\tilde{x})$$

1. *Plug & Play* (**PnP**) [Venkatakrishnan 2013]

*pre-trained **denoiser***

2. *Deep Equilibrium Networks* [Bai 2019]

constant memory

training

`stop-gradient`

Jacobian Free Backpropagation [Fung 2021]

$$\min_\phi \ \mathbb{E}_{(\tilde{\mathrm{x}},\mathrm{x}^*)} \| G_\phi(\tilde{\mathrm{x}}, K) - \mathrm{x}^* \|_2^2$$

- From 30mn/epoch, → ~10mn/epoch
- 57 GB → 9 GB

Limitations

- You still need to compute the forward loop at train-time
- For non-separable $\boldsymbol{A}$, patch-training is not possible

Alternative ?

$$G_\phi(\tilde{x}, K) = [D_\phi \circ g \circ \cdots \circ D_\phi \circ g](\tilde{x})$$

K steps

1. *Plug & Play* (**PnP**) [Venkatakrishnan 2013]

pre-trained ***denoiser***

2. *Deep Equilibrium Networks* [Bai 2019]

constant memory

training

3. *Neural Fields*

compressed representation

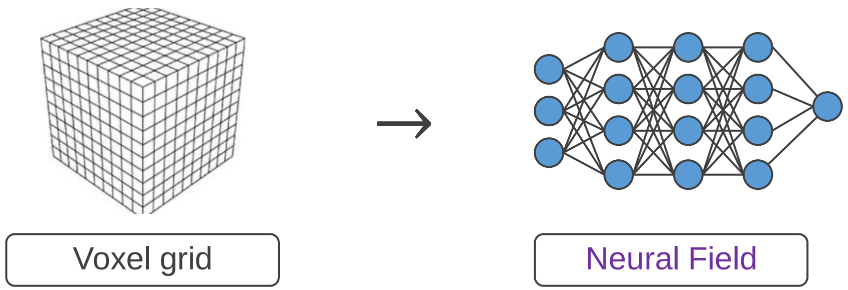

- Lightweight representation
  - lighter deep network ?
  - faster forward operator evaluation ?
- Inherent regularization
  - easy modulation of the information associated with a given frequency

Neural Fields

Efficient representation using Deep learning

« Compress » the parametrization of $\boldsymbol{x}$

At the root, x is just a scalar field $\mathbb{R}^3 \to \mathbb{R}$

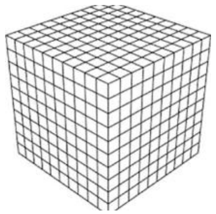

Voxel grid

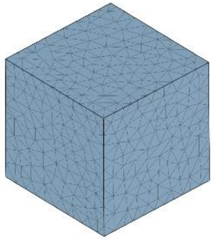

Mesh

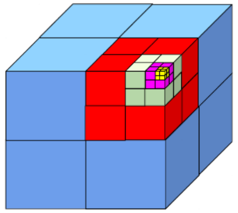

Octree

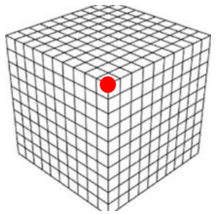

$x \in \mathbb{R}^n$

$x_i = F(x, u_i) \quad u_i \in \mathbb{R}^3$

$\rightarrow$

u_i

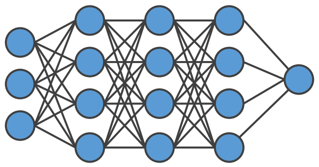

Neural Field

$x_i = F(\Theta, u_i)$

$\Theta \in \mathbb{R}^d$ and $u_i \in \mathbb{R}^3$

with F an interpolation function

u_i a coordinate in space

x_i the attenuation value at position u_i

with F a **neural network**

and d the number of parameters

$\Theta \in \mathbb{R}^d$ "stores" the information about the object

Efficient representation using Deep learning

Neural Fields, INR, etc..

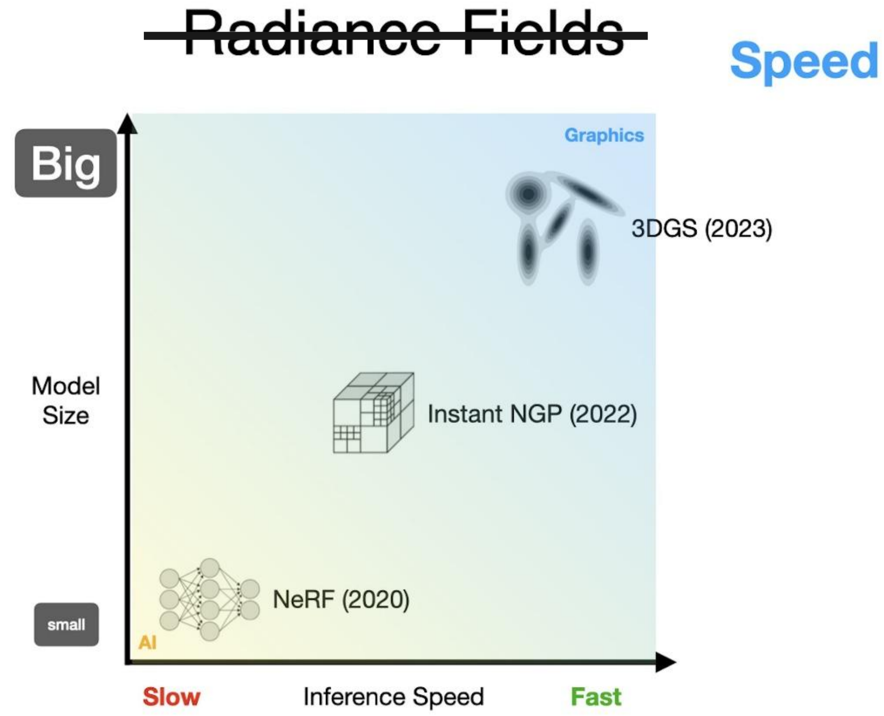

MLP-based

$\gamma(u_i)$ 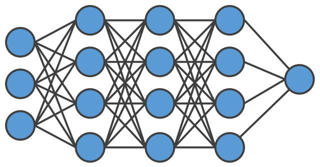 $x_i = [F_\Theta \circ \gamma](u_i)$

F_Θ

$$\gamma(u) = [a_1 \cos(2\pi b_1^\top u), a_1 \sin(2\pi b_1^\top u), \ldots, a_m \cos(2\pi b_m^\top u), a_m \sin(2\pi b_m^\top u)]$$

Grid-based (+ shallow MLP)

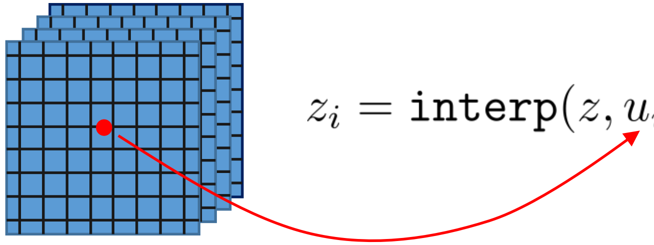

$z_i = \texttt{interp}(z, u_i)$

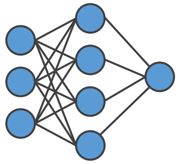

$x_i = F_\Theta(z_i)$

latent grid $\boldsymbol{z}$

F_Θ

lightweight representation of $\boldsymbol{x}$ → here we have $\boldsymbol{d} \ll \boldsymbol{n}$ **(4 times lighter)**

[Müller 2022]

Optimizing a neural field

$$\arg\min_{x\in\mathbb{R}^n} \frac{1}{2}\|Ax - b\|_2^2 + \lambda\mathcal{R}(x) \quad \rightarrow \quad \arg\min_{\Theta\in\mathbb{R}^d} \frac{1}{2}\|AF_\Theta(\mathcal{X}) - b\|_2^2 + \lambda\mathcal{R}(F_\Theta(\mathcal{X}))$$

coordinates of the grid: $\mathcal{X} \in \mathbb{R}^{n\times 3}$

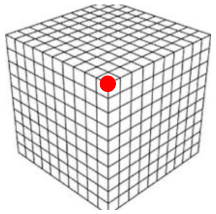

Voxel grid

$$x_i = F(x, u_i)$$

$$x \in \mathbb{R}^n \quad \text{and} \quad u_i \in \mathbb{R}^3$$

$\rightarrow$

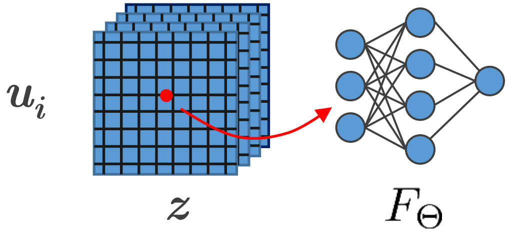

$x_i = F_\Theta(z_i)$

Neural Field

Optimizing a neural field

$$\underset{x\in\mathbb{R}^n}{\arg\min}\ \frac{1}{2}\|Ax-b\|_2^2+\lambda\mathcal{R}(x)\quad\rightarrow\quad\underset{\Theta\in\mathbb{R}^d}{\arg\min}\ \frac{1}{2}\|AF_\Theta(\mathcal{X})-b\|_2^2+\lambda\mathcal{R}(F_\Theta(\mathcal{X}))$$

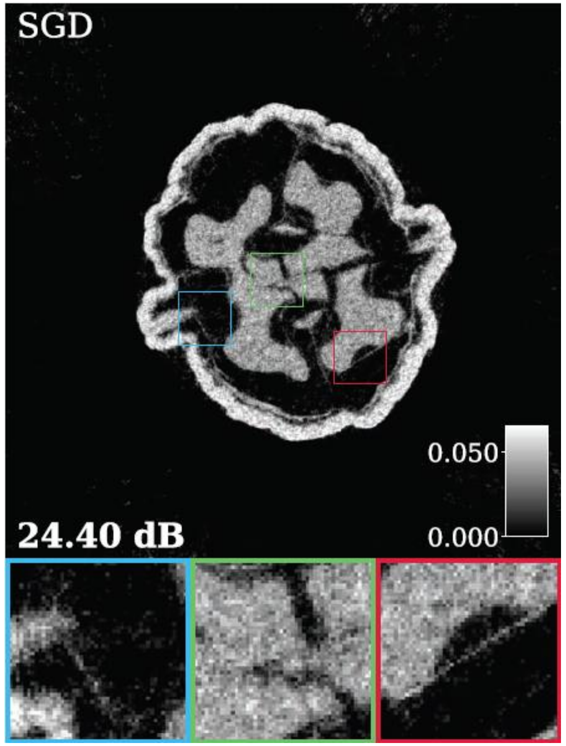

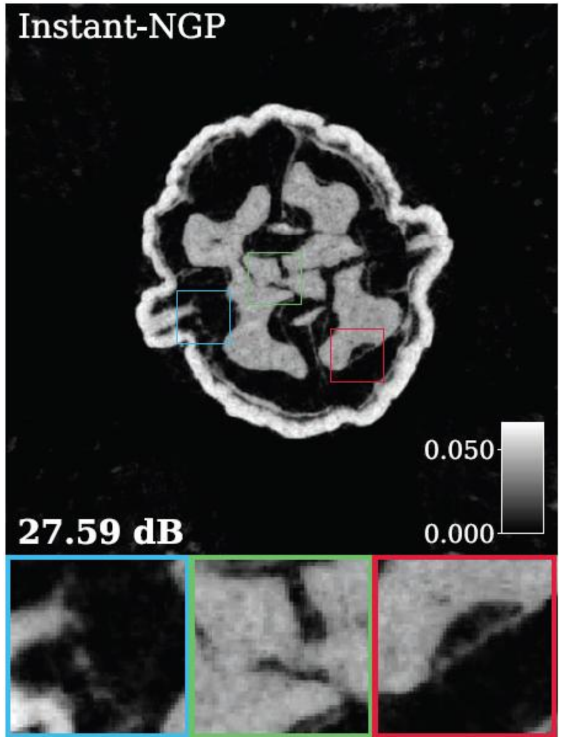

Regularized Neural Field

$$\arg\min_{\Theta \in \mathbb{R}^d} \frac{1}{2}\|AF_\Theta(\mathcal{X}) - b\|_2^2 + \lambda \mathcal{R}(F_\Theta(\mathcal{X}))$$

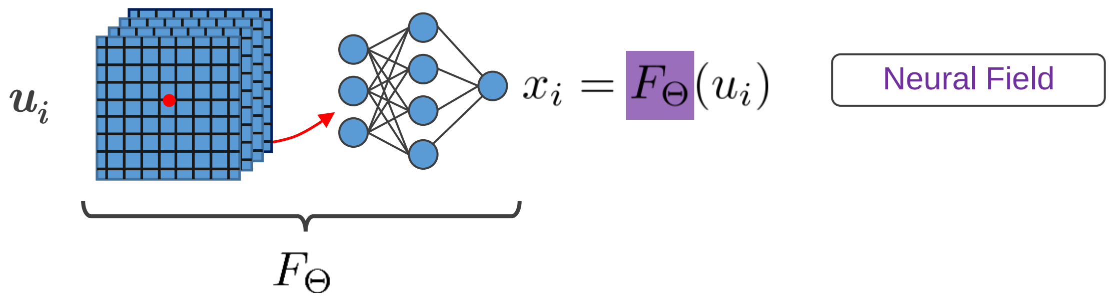

Proximal Gradient Descent

$$\begin{cases} z_{k+1} = x_k - \eta \nabla \mathcal{D}(x) \\ x_{k+1} = \text{prox}_{\eta\lambda\mathcal{R}}(z_{k+1}) \end{cases} \rightarrow \begin{cases} \Theta_{k+\frac{1}{2}} = \Theta_k - \eta \nabla \mathcal{D}(\Theta_k) \\ \Theta_{k+1} = \text{prox}_{\eta\lambda\mathcal{R}}(\Theta_{k+\frac{1}{2}}) \end{cases} \quad ?$$

Using a **gradient descent** instead:

$$x^{(k+1)} = x^{(k)} - \eta \nabla_x \mathcal{D}(x^{(k)}) - \eta\lambda \nabla \mathcal{R}(x^{(k)}) \rightarrow \Theta^{(k+1)} = \Theta^{(k)} - \eta \frac{\partial \mathcal{D}}{\partial x}\frac{\partial F}{\partial \Theta} - \eta\lambda \frac{\partial \mathcal{R}}{\partial x}\frac{\partial F}{\partial \Theta}$$

ENS DE LYON

Regularized Neural Field

$$\arg\min_{\Theta \in \mathbb{R}^d} \quad \frac{1}{2}\|AF_\Theta(\mathcal{X}) - b\|_2^2 + \lambda \mathcal{R}(F_\Theta(\mathcal{X}))$$

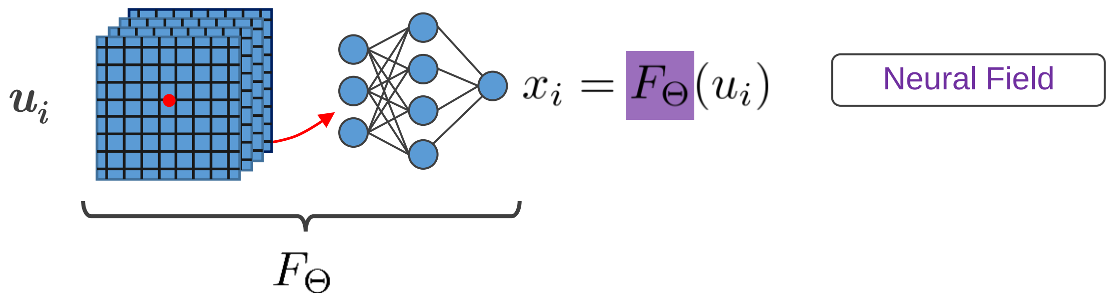

$$\Theta^{(k+1)} = \Theta^{(k)} - \eta \frac{\partial \mathcal{D}}{\partial x}\frac{\partial F}{\partial \Theta} - \eta\lambda \frac{\partial \mathcal{R}}{\partial x}\frac{\partial F}{\partial \Theta}$$

We use *Regularization by Denoising* (**RED**)

[Romano 2017]

$$\nabla \mathcal{R}(x) \approx h(x) = x - D_\phi(x)$$

Results – empirical convergence

$$\hat{\Theta} \in \underset{\Theta \in \mathbb{R}^d}{\arg\min} \; \mathcal{D}(F_\Theta) + \lambda \mathcal{R}(F_\Theta)$$

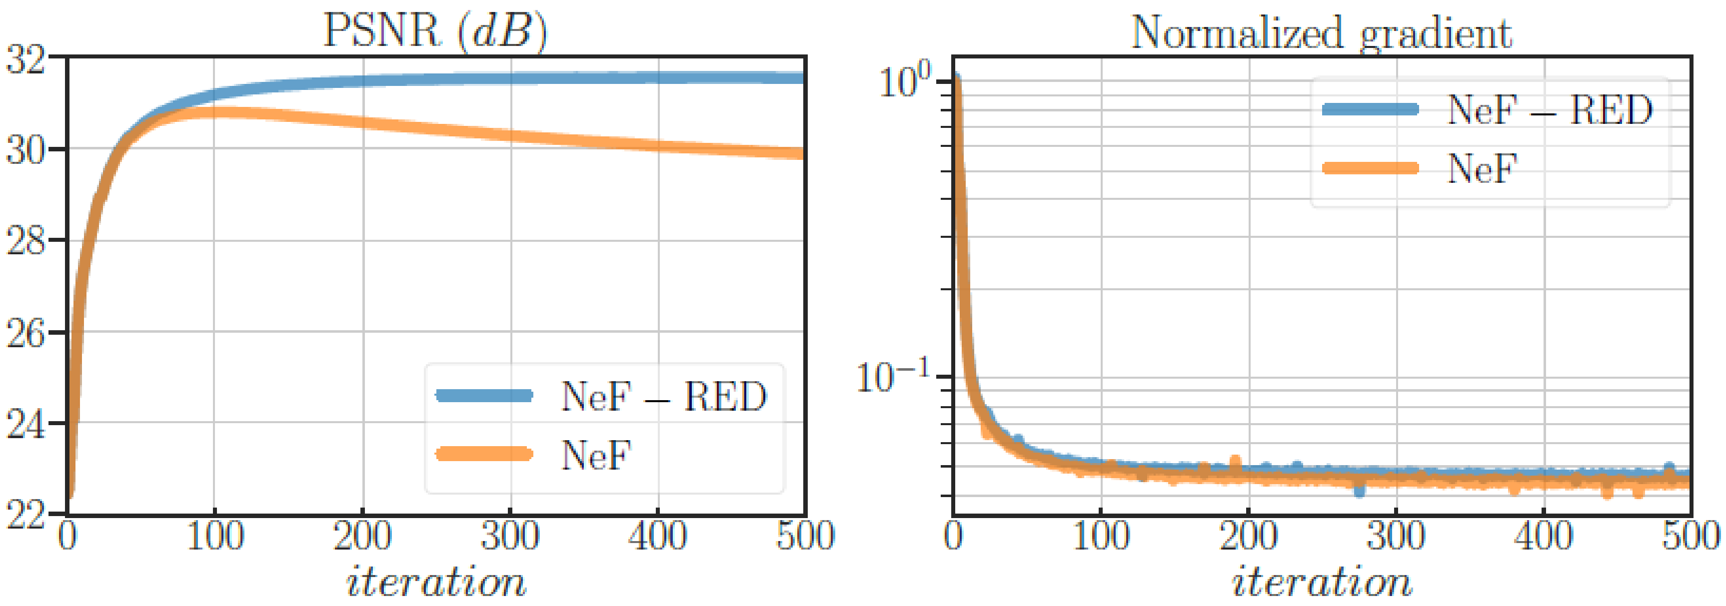

Results – performances

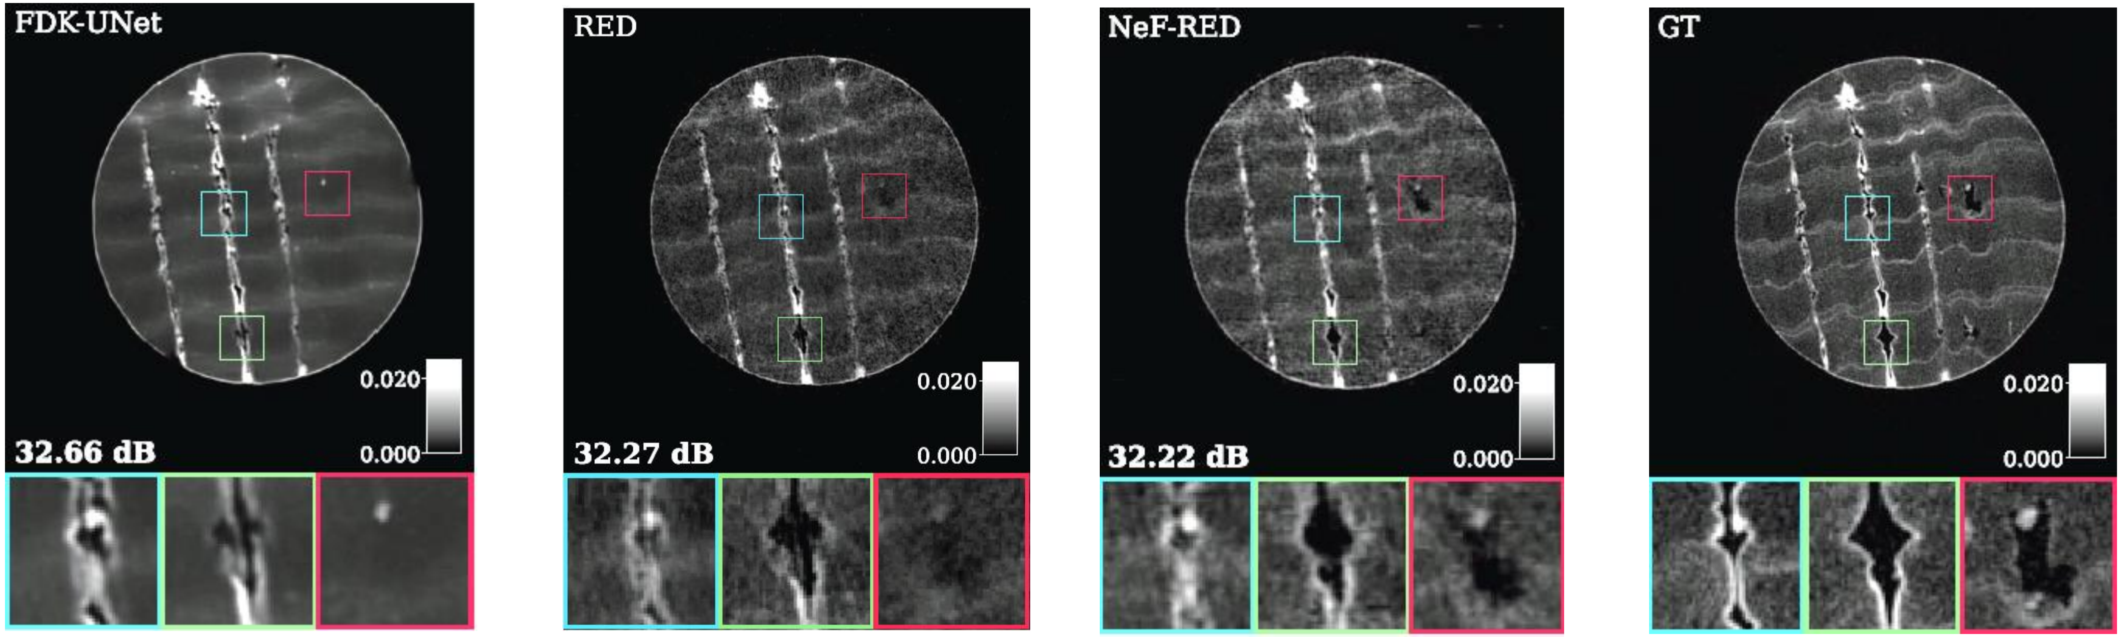

Cork-CBCT - **3D**	sampling	optimization	**SSIM** ↑	**PSNR** ↑	**VRAM** ↓
FDK	-	-	0.258	24.11	4.9
SGD	-	*per-scene*	0.652	30.52	4.7
FDK-UNet	-	***generalized***	0.827	32.96	4.9
Instant-NGP [Mül+22; ZZL22]	-	*per-scene*	0.705	30.84	7.2
RED	***ours***	*per-scene*	0.790	33.24	33.2
Ours: NeF-RED	***ours***	*per-scene*	0.797	33.27	10.0

Regularization tradeoff

$$\Theta^{(k+1)} = \Theta^{(k)} - \eta \frac{\partial \mathcal{D}}{\partial x} \frac{\partial F}{\partial \Theta} - \eta \lambda h(F_{\Theta^{(k)}}) \frac{\partial F}{\partial \Theta}$$

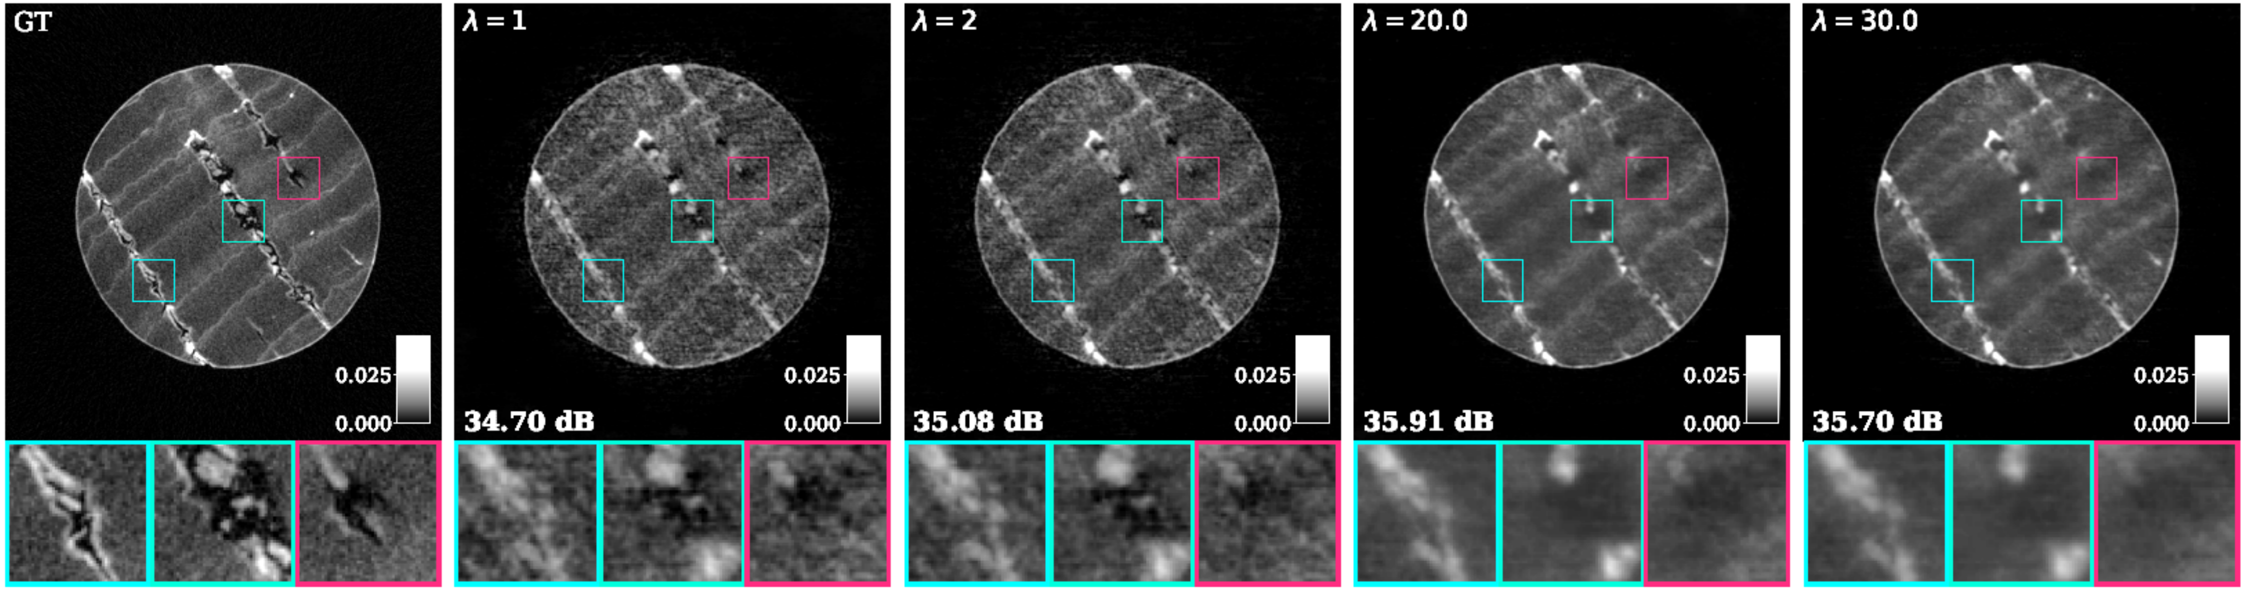

Figure 3.9: Effect of the regularization weight λ on the perceptual quality (qualitative quality) and distortion (quantitative quality / PSNR). The higher we set lambda, the smoother the reconstruction. However, the effect on the perception is not as clear (best-viewed zoomed-in).

Conclusions & Perspectives

Limitations & Perspectives

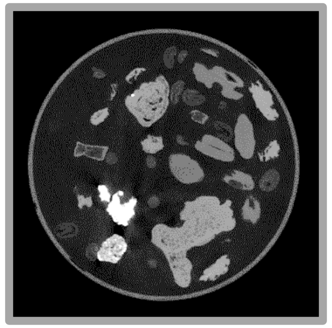

$F_{\Theta} : \mathbb{R}^2 \rightarrow \mathbb{R}$

$\Theta \in \mathbb{R}^d$

low-dimensional

$d \ll n$

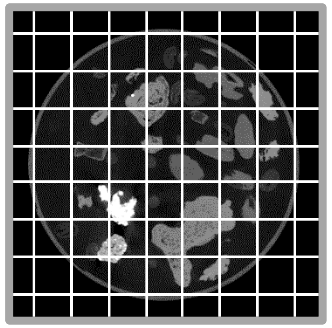

$\boldsymbol{x} \in \mathbb{R}^n$

costly

$\boldsymbol{A} \in \mathbb{R}^{m \times n}$

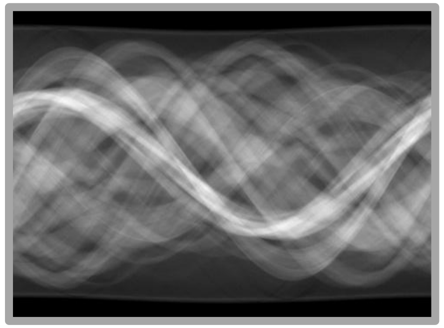

$\boldsymbol{b} \in \mathbb{R}^m$

$$\nabla_{\boldsymbol{x}} \mathcal{R}(\boldsymbol{x}) = \boldsymbol{x} - D(\boldsymbol{x})$$

$$\text{with} \quad \boldsymbol{x} = F_{\Theta}(\texttt{grid})$$

Limitations & Perspectives

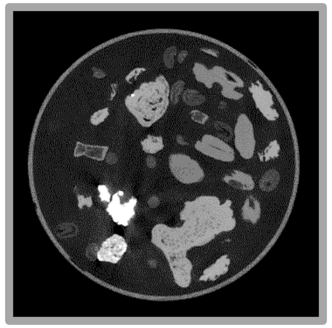

$F_{\Theta} : \mathbb{R}^2 \to \mathbb{R}$

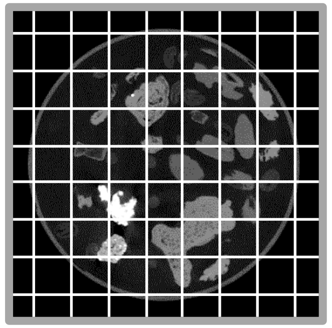

$\boldsymbol{x} \in \mathbb{R}^n$

costly

$\boldsymbol{A} \in \mathbb{R}^{m \times n}$

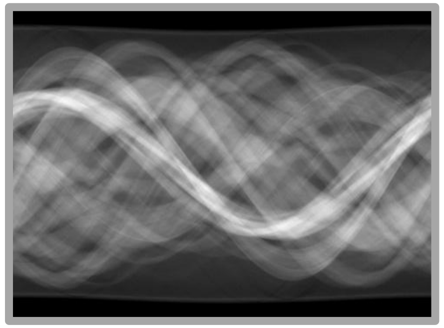

$\boldsymbol{b} \in \mathbb{R}^m$

$\boldsymbol{\Theta} \in \mathbb{R}^d$

low-dimensional

$d \ll n$

Latent denoising instead of grid denoising

$$\nabla_x \mathcal{R}(x) = x - D(x) \quad \to \quad \nabla_\Theta \mathcal{R}(\Theta) = \Theta - D(\Theta)$$

Limitations & Perspectives

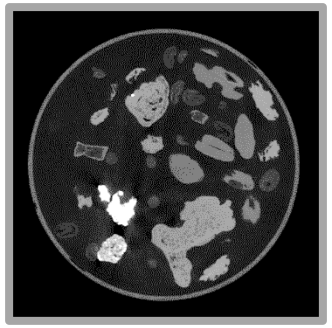

$F_{\Theta} : \mathbb{R}^2 \to \mathbb{R}$

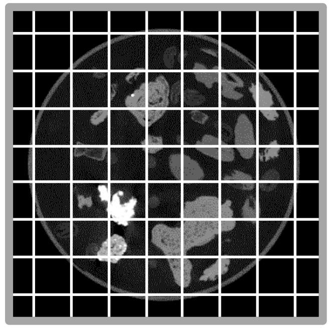

$\boldsymbol{x} \in \mathbb{R}^n$

costly

$\boldsymbol{A} \in \mathbb{R}^{m \times n}$

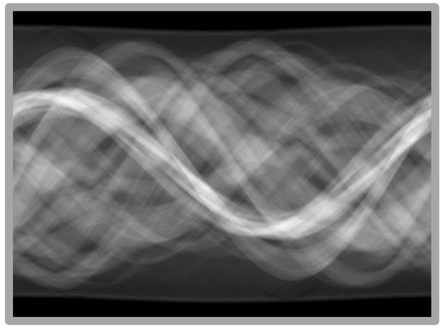

$\boldsymbol{b} \in \mathbb{R}^m$

$\boldsymbol{\Theta} \in \mathbb{R}^d$

low-dimensional

$d \ll n$

Latent denoising instead of grid denoising

- ❑ **build** operator acting on the latent
- ❑ **learning** this operator

$$\nabla_x \mathcal{R}(x) = x - D(x) \quad \to \quad \nabla_\Theta \mathcal{R}(\Theta) = \Theta - D(\Theta)$$

Limitations & Perspectives

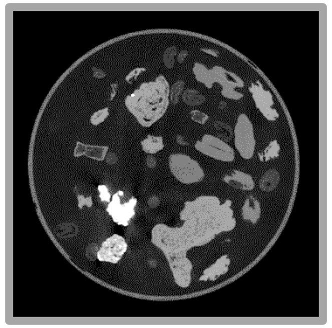

$F_{\Theta} : \mathbb{R}^2 \to \mathbb{R}$

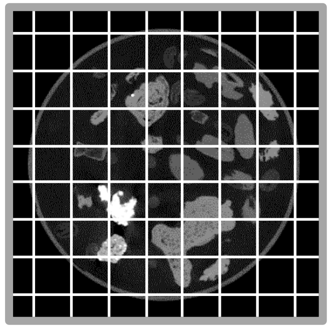

$\boldsymbol{A} \in \mathbb{R}^{m \times n}$

costly

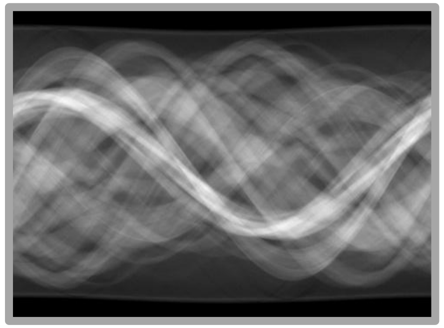

$\boldsymbol{x} \in \mathbb{R}^n$

$\boldsymbol{b} \in \mathbb{R}^m$

costly

$\Theta \in \mathbb{R}^d$

low-dimensional

$d \ll n$

Data-fidelity

- **Variational approach** $\to \boldsymbol{A}, \boldsymbol{A}^\top$ [Venkatakrishnan 2013, Adler 2017]
- **One-step** conditioning [Terris 2025]

$$\hat{\boldsymbol{x}} = D_\phi(\tilde{\boldsymbol{x}}, \boldsymbol{A}, \boldsymbol{b})$$

cnrs ENS DE LYON

Limitations & Perspectives

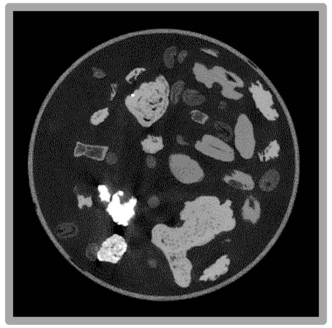

$F_{\Theta} : \mathbb{R}^2 \to \mathbb{R}$

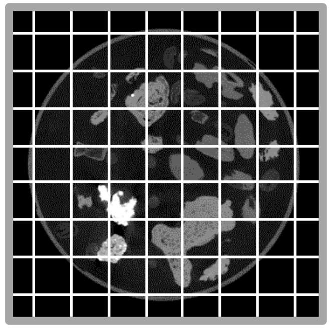

$\boldsymbol{A} \in \mathbb{R}^{m \times n}$

$\boldsymbol{x} \in \mathbb{R}^n$

costly

costly

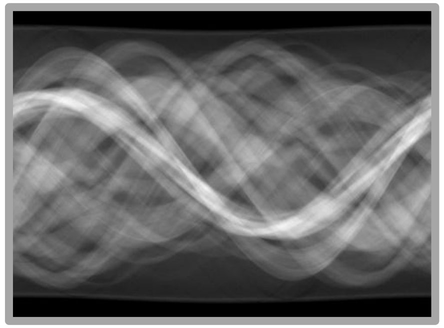

$\boldsymbol{b} \in \mathbb{R}^m$

$\Theta \in \mathbb{R}^d$

low-dimensional

$d \ll n$

Data-fidelity

- **Variational approach** → $\boldsymbol{A}, \boldsymbol{A}^\top$ [Venkatakrishnan 2013, Adler 2017]
- **One-step** conditioning [Terris 2025]
- **Null-space** denoising [Schwab 2019]

$$\hat{\boldsymbol{x}} = A^+ b + (I - A^+ A) D_\phi(\tilde{\boldsymbol{x}})$$

Limitations & Perspectives

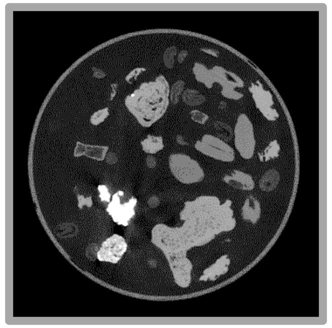

$F_{\Theta} : \mathbb{R}^2 \to \mathbb{R}$

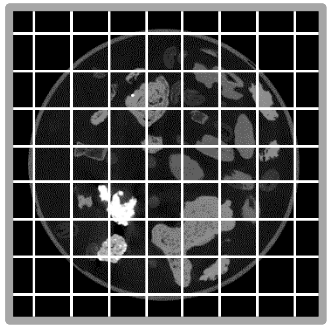

$\boldsymbol{x} \in \mathbb{R}^n$

costly

$\boldsymbol{A} \in \mathbb{R}^{m \times n}$

costly

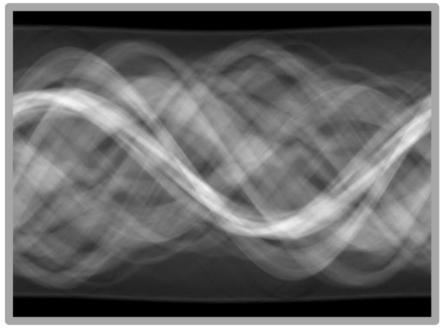

$\boldsymbol{b} \in \mathbb{R}^m$

$\Theta \in \mathbb{R}^d$

low-dimensional

$d \ll n$

Data-fidelity $\mathcal{D}(\Theta) = \frac{1}{2}\|AF_{\Theta}(\mathcal{X}) - b\|_2^2$

❑ Approximation ? $f_{\phi}(\Theta) \approx \nabla\mathcal{D}(\Theta)$

[Lindell 2021]
[Raphaeli 2025]

Call for Data !

Develop solutions for **large-scale** imaging inverse problem in 3D

- ❑ 3D CT (ray-based partitioning)
- ❑ 3D MRI
- ❑ **others ?**

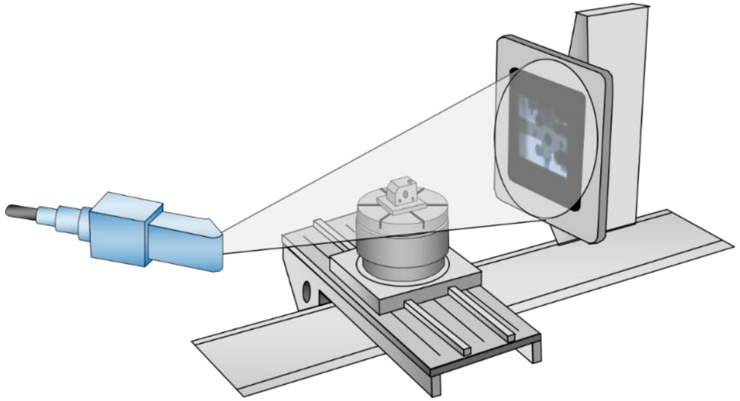

?

Thank you for your attention

References - I

- Qing Liu, Mengdi Sun, Xianda Sun, Bo Liu, Mehdi Ostadhassan, Wanxia Huang, Xiaoxia Chen, Zhejun Pan, Pore network characterization of shale reservoirs through state-of-the-art X-ray computed tomography: A review, Gas Science and Engineering, Volume 113, 2023.
- Kiss, Maximilian B., Sophia B. Coban, K. Joost Batenburg, Tristan van Leeuwen, and Felix Lucka. 2023. '2DeteCT - A Large 2D Expandable, Trainable, Experimental Computed Tomography Dataset for Machine Learning'. Scientific Data 10 (1): 576. https://doi.org/10.1038/s41597-023-02484-6.
- Der Sarkissian, Henri, Felix Lucka, Maureen van Eijnatten, Giulia Colacicco, Sophia Bethany Coban, and Kees Joost Batenburg. 2019. 'A Cone-Beam X-Ray Computed Tomography Data Collection Designed for Machine Learning'. Scientific Data 6 (1): 215. https://doi.org/10.1038/s41597-019-0235-y.
- Anastasio, M.A., X. Pan, and E. Clarkson. 2001. 'Comments on the Filtered Backprojection Algorithm, Range Conditions, and the Pseudoinverse Solution'. IEEE Transactions on Medical Imaging 20 (6): 539–42. https://doi.org/10.1109/42.929620.
- Feldkamp, L. A., L. C. Davis, and J. W. Kress. 1984. 'Practical Cone-Beam Algorithm'. Journal of the Optical Society of America A 1 (6): 612. https://doi.org/10.1364/JOSAA.1.000612.
- Combettes, Patrick L., and Jean-Christophe Pesquet. 2010. 'Proximal Splitting Methods in Signal Processing'. arXiv:0912.3522 [Math], May. http://arxiv.org/abs/0912.3522.
- Sidky, Emil Y., Jakob H. Jørgensen, and Xiaochuan Pan. 2012. 'Convex Optimization Problem Prototyping for Image Reconstruction in Computed Tomography with the Chambolle–Pock Algorithm'. Physics in Medicine and Biology 57 (10): 3065–91. https://doi.org/10.1088/0031-9155/57/10/3065.
- Han, Yo Seob, Jaejun Yoo, and Jong Chul Ye. 2016. 'Deep Residual Learning for Compressed Sensing CT Reconstruction via Persistent Homology Analysis'. arXiv:1611.06391 [Cs], November. http://arxiv.org/abs/1611.06391.
- Ronneberger, Olaf, Philipp Fischer, and Thomas Brox. 2015. 'U-Net: Convolutional Networks for Biomedical Image Segmentation'. In Medical Image Computing and Computer-Assisted Intervention – MICCAI 2015, edited by Nassir Navab, Joachim Hornegger, William M. Wells, and Alejandro F. Frangi, 234–41. Lecture Notes in Computer Science. Cham: Springer International Publishing. https://doi.org/10.1007/978-3-319-24574-4_28.
- Adler, Jonas, and Ozan Öktem. 2018. 'Learned Primal-Dual Reconstruction'. IEEE Transactions on Medical Imaging 37 (6): 1322–32. https://doi.org/10.1109/TMI.2018.2799231.
- Venkatakrishnan, Singanallur V., Charles A. Bouman, and Brendt Wohlberg. 2013. 'Plug-and-Play Priors for Model Based Reconstruction'. In 2013 IEEE Global Conference on Signal and Information Processing, 945–48. Austin, TX, USA: IEEE. https://doi.org/10.1109/GlobalSIP.2013.6737048.
- Zhang, Kai, Wangmeng Zuo, Shuhang Gu, and Lei Zhang. 2017. 'Learning Deep CNN Denoiser Prior for Image Restoration'. arXiv. https://doi.org/10.48550/arXiv.1704.03264.
- Cohen, Regev, Michael Elad, and Peyman Milanfar. 2021. 'Regularization by Denoising via Fixed-Point Projection (RED-PRO)'. SIAM Journal on Imaging Sciences 14 (3): 1374–1406. https://doi.org/10.1137/20M1337168.
- Vo, Romain, Julie Escoda, Caroline Vienne, and Étienne Decencière. 2025. 'Plug-and-Play Learned Proximal Trajectory for 3D Sparse-View X-Ray Computed Tomography'. In Computer Vision – ECCV 2024, edited by Aleš Leonardis, Elisa Ricci, Stefan Roth, Olga Russakovsky, Torsten Sattler, and Gül Varol, 15093:221–38. Lecture Notes in Computer Science. Cham: Springer Nature Switzerland. https://doi.org/10.1007/978-3-031-72761-0_13.

ENS DE LYON

References - II

- Mescheder, Lars, Michael Oechsle, Michael Niemeyer, Sebastian Nowozin, and Andreas Geiger. 2019. 'Occupancy Networks: Learning 3D Reconstruction in Function Space'. In , 4460–70. https://openaccess.thecvf.com/content_CVPR_2019/html/Mescheder_Occupancy_Networks_Learning_3D_Reconstruction_in_Function_Space_CVPR_2019_paper.html.
- Mildenhall, Ben, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. 2020. 'NeRF: Representing Scenes as Neural Radiance Fields for View Synthesis'. In Computer Vision – ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part I, 405–21. Berlin, Heidelberg: Springer-Verlag. https://doi.org/10.1007/978-3-030-58452-8_24.
- Müller, Thomas, Alex Evans, Christoph Schied, and Alexander Keller. 2022. 'Instant Neural Graphics Primitives with a Multiresolution Hash Encoding'. ACM Transactions on Graphics 41 (4): 1–15. https://doi.org/10.1145/3528223.3530127.
- Romano, Yaniv, Michael Elad, and Peyman Milanfar. 2017. 'The Little Engine That Could: Regularization by Denoising (RED)'. SIAM Journal on Imaging Sciences 10 (4): 1804–44. https://doi.org/10.1137/16M1102884.
- Fercoq, Olivier. 2019. 'Coordinate Descent Methods for Non-Differentiable Convex Optimisation Problems'. Thesis, Sorbonne Université. https://hal.science/tel-02944077.
- Sun, Yu, Jiaming Liu, and Ulugbek Kamilov. 2019. 'Block Coordinate Regularization by Denoising'. In Advances in Neural Information Processing Systems. Vol. 32. Curran Associates, Inc. https://papers.nips.cc/paper/2019/hash/9872ed9fc22fc182d371c3e9ed316094-Abstract.html.
- Whalen, A. 1984. 'Statistical Theory of Signal Detection and Parameter Estimation'. IEEE Communications Magazine 22 (6): 37–44. https://doi.org/10.1109/MCOM.1984.1091969.
- Barrett, H H, J Yao, J P Rolland, and K J Myers. 1993. 'Model Observers for Assessment of Image Quality.' Proceedings of the National Academy of Sciences 90 (21): 9758–65. https://doi.org/10.1073/pnas.90.21.9758.
- Zhou, Weimin, Hua Li, and Mark A. Anastasio. 2019. 'Approximating the Ideal Observer and Hotelling Observer for Binary Signal Detection Tasks by Use of Supervised Learning Methods'. IEEE Transactions on Medical Imaging 38 (10): 2456–68. https://doi.org/10.1109/TMI.2019.2911211.
- Adler, Jonas, Sebastian Lunz, Olivier Verdier, Carola-Bibiane Schönlieb, and Ozan Öktem. 2018. 'Task Adapted Reconstruction for Inverse Problems'. arXiv:1809.00948 [Cs, Math], August. http://arxiv.org/abs/1809.00948.
- Zhang, Jiajin, Hanqing Chao, Xuanang Xu, Chuang Niu, Ge Wang, and Pingkun Yan. 2021. 'Task-Oriented Low-Dose CT Image Denoising'. In Medical Image Computing and Computer Assisted Intervention – MICCAI 2021, edited by Marleen de Bruijne, Philippe C. Cattin, Stéphane Cotin, Nicolas Padoy, Stefanie Speidel, Yefeng Zheng, and Caroline Essert, 12906:441–50. Lecture Notes in Computer Science. Cham: Springer International Publishing. https://doi.org/10.1007/978-3-030-87231-1_43.
- Selvaraju, Ramprasaath R., Michael Cogswell, Abhishek Das, Ramakrishna Vedantam, Devi Parikh, and Dhruv Batra. 2020. 'Grad-CAM: Visual Explanations from Deep Networks via Gradient-Based Localization'. International Journal of Computer Vision 128 (2): 336–59. https://doi.org/10.1007/s11263-019-01228-7.

ENS DE LYON