Image Reconstruction Tutorial

Part 2: Computed Tomography (CT) reconstruction

Aleksandra Pižurica[1] and Bart Goossens[2]

[1]Group for Artificial Intelligence and Sparse Modelling (GAIM)

[2]Image Processing and Interpretation (IPI) group

TELIN, Ghent University - imec

Yearly workshop FWO–WOG

Turning images into value through statistical parameter estimation

Ghent, Belgium, 20 September 2019

Computed Tomography

- The goal of **tomography** (from the Greek *tomos* for *section*) is to recover the interior structure of a body using external measurements.

- Various probes, including X-rays, gamma rays, visible light, electrons, protons, neutrons, sound waves, and nuclear magnetic resonance signals can be used to study a large variety of objects ranging from complex molecules through astronomical objects.

- The most popular application of tomography is Computed Tomography (CT) for medical imaging, widely used for medical diagnostic.

$\rightarrow$ Over 80 million CT scans were performed in the USA in 2015.

Computed Tomography

- CT involves the exposure of the patient to x-ray radiation.

This is associated with health risks (radiation-induced carcinogenesis) essentially **proportional to the levels of radiation exposure**.
$\Rightarrow$ 2% of cancers in the United States attributed to CT radiation.

- Radiation exposure can be directly reduced, this often leads to a lower SNR and/or lower image resolution $\Rightarrow$ trade-off diagnostic quality vs. radiation dose.

- Another technique consists of sparse sampling (e.g., sparse-angle CT reconstruction)

- In some cases, only a “small” region-of-interest (ROI) needs to be reconstructed.

Computed Tomography acquisition

Tomography: A series of planar images is acquired from different angles around the patent.

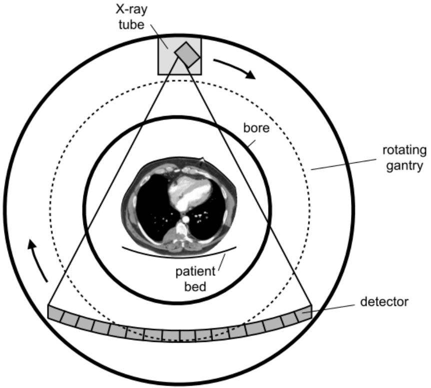

Picture taken from [Vandeghinste, 2014]

The Radon Transform: Simple Backprojection

In 2D, the measurements can be mathematically represented by the **Radon transform** R, which maps a density function f into linear projections.

A line ℓ can be parametrized with respect to $e_\theta = (\cos\theta, \sin\theta) \in S^1$ and $t \in \mathbb{R}$:

$$\ell(\theta, t) = \{y = (u, v) \in \mathbb{R}^2 : e_\theta \cdot y = t\}.$$

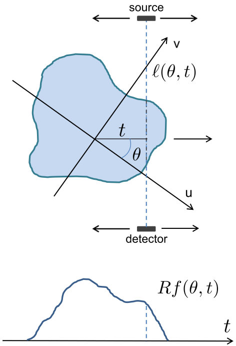

In 1917, Johann Radon proved that an object can be reconstructed exactly from an infinite number of projections, when taken over 360° around the object.

The Radon Transform of the Shepp Logan phantom

Shepp Logan image

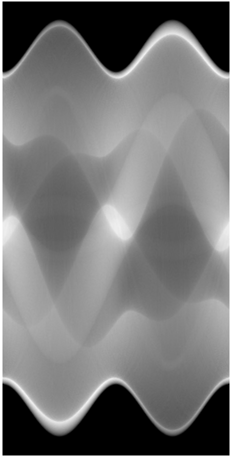

Radon transform (sinogram)

The Radon Transform: Simple Backprojection

Radon transform: For each $\theta \in S^1$ and $t \in \mathbb{R}$

$$p(\theta, t) = Rf(\theta, t) \quad = \quad \int_{\ell(\theta,t)} f(\mathbf{y})\, d\mathbf{y}$$

Backprojection (mathematically incorrect):

$$f(\mathbf{y}) = R^*\{p\} = \int_0^{\pi} p(\theta, u\cos\theta + v\sin\theta) d\theta$$

Why incorrect?

- To explain: Fourier Slice Theorem needed
- Consequence: image reconstruction techniques required!

Fourier Slice Theorem: the 1D Fourier transform of a parallel projection of an object $f(\mathbf{y})$ obtained at an angle θ equals one line in the 2D Fourier transform of $f(\mathbf{y})$ at the same angle θ.

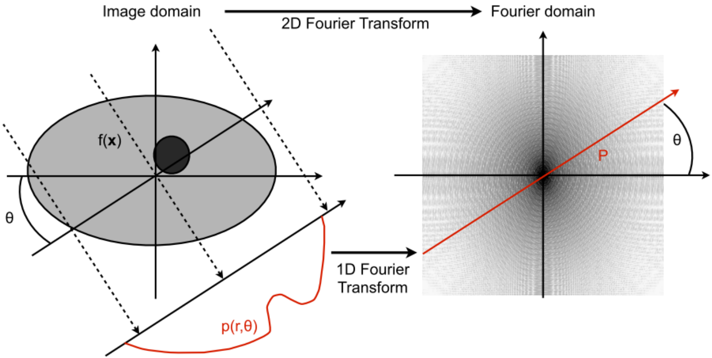

Picture taken from [Vandeghinste, 2014]

CT reconstruction by Filtered Backprojection

Backprojection blur is caused by a polar sampling pattern in Fourier space.

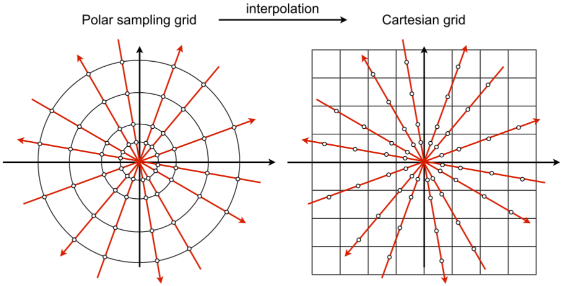

Picture taken from [Vandeghinste, 2014]

The density of samples near the center is a factor $1/r$ higher than at the outer regions, with r the radial distance to the center.

CT reconstruction by Filtered Backprojection

Solution: uniform sampling density requires the Fourier transform of each projection to be multiplied with a ramp filter proportional with this $1/r$ factor:

$$f(\mathbf{y}) = R^*(q \star R\{f(\mathbf{y})\})$$

where the filter q has Fourier transform:

$$\mathcal{F}\{q\}(\omega) = \left|\frac{\omega}{2\pi}\right| G(\omega)$$

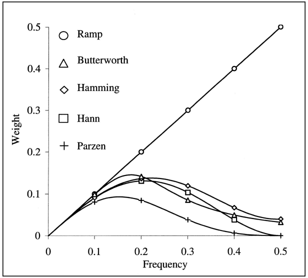

with $G(\omega)$ a smoothing filter (e.g., sinc filter, cosine filter, Parzen filter, a Hamming window, Hann window, ...). The smoothing filter directly influences the quality of the reconstructed image in terms of noise, resolution, contrast and other measures.

Simple versus Filtered Backprojection

Illustration of the difference between simple backprojection and filtered Backprojection

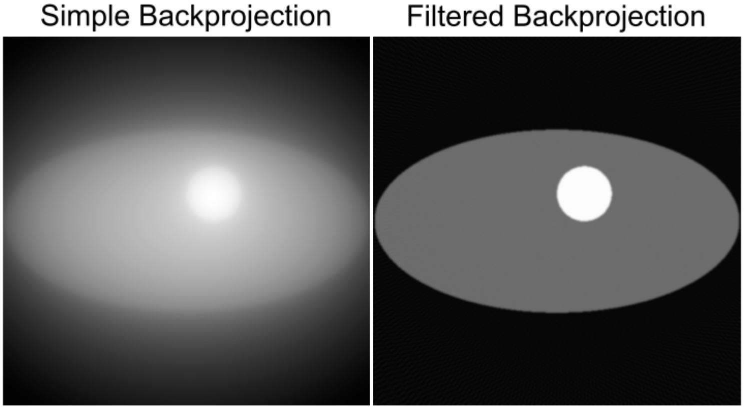

Picture taken from [Vandeghinste, 2014]

Filtered Backprojection (FBP) characteristics

- The most commonly used image reconstruction method in CT due to
  1. being very fast
  2. having low memory requirements
  3. yielding good results on many data
- Originally defined for parallel-beam geometry; extensions exist for current systems (e.g., fan-beam, cone-beam, helical cone-beam).
- Exact solution in absense of *noise*, *complete data* and for *uniform spatial resolution*
- In practice, these conditions usually do not apply $\Rightarrow$ iterative reconstruction

Iterative Reconstruction Methods

Solve linear systems numerically

$$\mathbf{y} = \mathbf{W}\mathbf{f} \Rightarrow \mathbf{f} = ?$$

- $\mathbf{f}$: the input image, arranged as a vector (e.g., using column stacking)
- $\mathbf{y}$: the output sinogram, arranged as a vector
- $\mathbf{W}$: system matrix of elements w_{ij}, which relates the contribution of every pixel (voxel) j in $\mathbf{f}$ to every detector element i.

Linear system too large to solve directly $\Rightarrow$ instead, use iterative solvers.

Algebraic Iterative Reconstruction (ART)

Algebraic Iterative Reconstruction (ART), by Gordon, Bender and Herman in 1970:

$$\mathbf{f}^{(k+1)} = \mathbf{f}^{(k)} + \lambda_k \frac{y_i - \mathbf{w}_i^T \mathbf{f}^{(k)}}{\mathbf{w}_i^T \mathbf{w}_i} \mathbf{w}_i$$

with $\mathbf{w_i} = (w_{i1}, w_{i2}, ..., w_{iJ})$ the i-th row of the system matrix $\mathbf{W}$.

Intuitively, the current image estimate $\mathbf{f}^{(k)}$ is forward projected and compared to the measured data. The error due to mis-estimation is redistributed to the current estimate, bringing it closer to the final solution.

Filtered Backprojection vs ART

Example with 3% noise and projection angles $15^\circ, 30^\circ, 45^\circ, \ldots, 180^\circ$
Incorporate constraints (e.g., non-negativity)

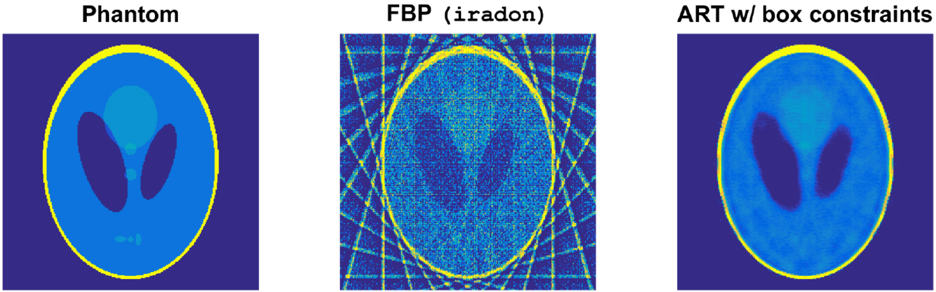

Iterative reconstruction techniques: the good...

Compared to Filtered Backprojection, iterative reconstruction offers:

- Improved image quality (in particular in presence of noise and limited data), at a higher computational cost (compute on GPU).
- More flexibility to adapt the reconstruction to *incomplete data*, *noise characteristics* and *image prior knowledge*.
- Several improvements of ART have been proposed, including Simultaneous Iterative Reconstruction Technique (SIRT) [Herman and Lent, 1976], Image Space Reconstruction Algorithm (ISRA), Maximum Likelihood for Transmission Tomography (MLTR) [Yu et al., 2000], ...
- In 2015, Siemens integrated their own Sinogram Affirmed Iterative Reconstruction (SAFIRE) algorithm in their CT scanners.

Iterative reconstruction techniques: the bad...

Iterative reconstruction techniques faces several challenges, especially in presence of noise / undersampling:

- Data fidelity ($\hat{\mathbf{y}} \approx \mathbf{y}$), even with regularization is not enough to guarantee a good image!
  $\Rightarrow$ Problem is not always uniquely solvable
  $\Rightarrow$ Given a projection error $||\hat{\mathbf{y}} - \mathbf{y}||^2$ we want to control the image reconstruction error $||\hat{\mathbf{f}} - \hat{\mathbf{f}}||^2$
  $\Rightarrow$ Challenging problem, due to the null-space of $\mathbf{W}$
- Sometimes, iterative reconstruction algorithms are stopped after a fixed number of iterations (best image quality(?)), rather than at convergence.
  $\Rightarrow$ Study of the relation between reconstruction parameters, noise and image quality is very important!
  $\Rightarrow$ Research domain: medical image quality assessment and optimization.

Sparsity-Inducing Reconstruction Algorithm (SIRA)

Joint work with Demetrio Labate and Bernhard Bodmann from Univ. of Houston.

ROI Computed Tomography

ROI Computed Tomography is concerned with reconstructing an ROI within the field of view using ROI-focused scanning only.

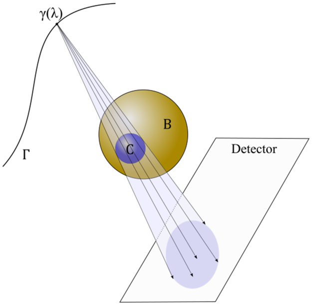

Challenge: since projections are truncated, the reconstruction problem may become severely ill-posed.

$\rightarrow$ *Interior problem* (projections are known only for rays intersecting an ROI strictly inside the field of view) is in general not uniquely solvable.

Existing methods for local ROI CT reconstruction require **restrictions on the geometry and location** of the ROI or some **prior knowledge** about the density function.

- **Differentiated Back-Projection** [Clackdoyle et al., 2004]

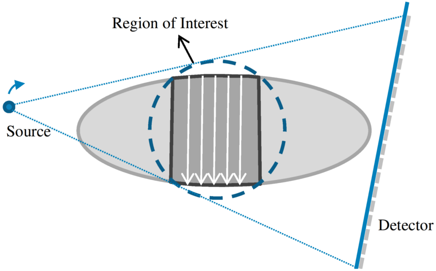

There exists a projection angle such that for angles in its vicinity, non-truncated projection data is available

ROI Computed Tomography

- **Known subregion** [Kudo et al., 2008]

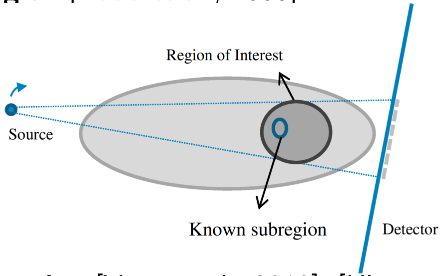

Density function to be recovered is known on a subregion inside the ROI

- **Special assumption** [Yang et al., 2010], [Klann et al., 2015]

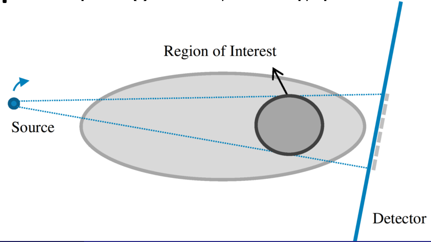

Density function satisfies special assumption, e.g., piecewise constant in ROI

ROI Computed Tomography

Remark: even when local reconstruction is theoretically guaranteed, the practical solution might be numerically unstable.
$\Rightarrow$ There is no theoretical guarantee in the presence of **noise**.

Our ROI CT reconstruction method includes performance guarantees in the setting of **noisy projection data**.

Novelty:
- we treat image and projection data **jointly** in the recovery
- a robust width **prior assumption** that relies on **sparsity norms** and measurement models supported by the theory of compressed sensing.

ROI Reconstruction problem

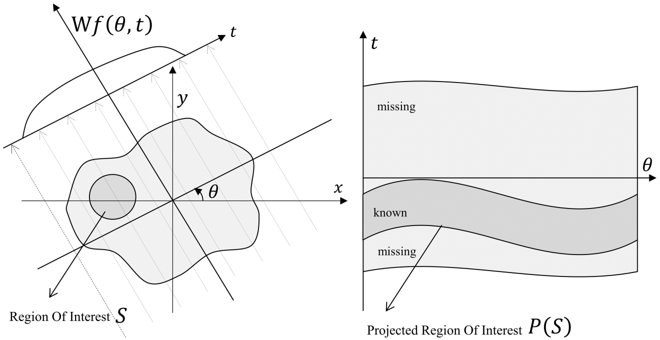

$W \rightarrow$ **projection operator** (e.g., Radon transform, fan-beam transform)
it maps a density function f into linear projections defined in the tangent space of the circle $\mathcal{T} = \{(\theta, t) : \theta \in [0, 2\pi), t \in \mathbb{R}\}$

$S \subset \mathbb{R}^2 \rightarrow$ **ROI** (image space)

$P(S) = \{(\theta, t) \in \mathcal{T} : \ell(\theta, t) \cap S \neq 0\} \rightarrow$ **ROI** (projection space)

$M = \chi_{P(S)} \rightarrow$ **ROI mask** (projection space)

ROI Reconstruction problem

ROI reconstruction problem

$$\text{Find } f \text{ on } S \text{ given } y_0(\theta, t) = M(\theta, t)\, Wf(\theta, t)$$

or

ROI reconstruction problem (with noise)

$$\text{Find } f \text{ on } S \text{ given } y_0(\theta, t) = M(\theta, t)\, (Wf(\theta, t) + \nu(\theta, t))$$

In the presence of noise, $\|MWf - y_0\|_2 > 0$ and an arbitrary extension y of y_0 may fail to be in the range of W.
Hence we formulate two constraints:

$$\|My - y_0\|_2 \leq \alpha \quad \text{(data fidelity)}$$
$$\|y - Wf\|_2 \leq \beta \quad \text{(data consistency)}$$

In the presence of noise, α and β cannot be both set to 0 in general.

ROI Reconstruction problem

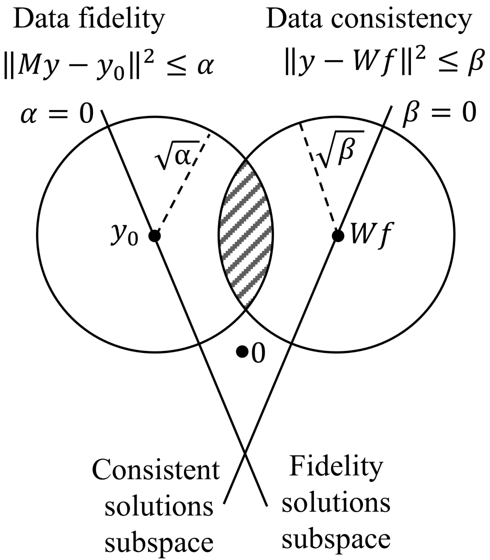

ROI Reconstruction problem

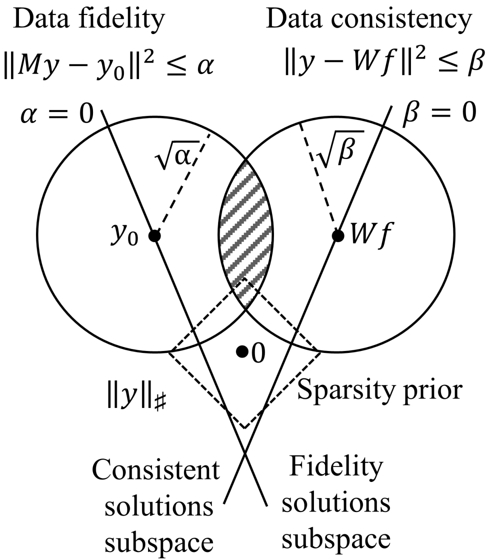

Robust Width Property

Robust Width Property (RWP) [Cahill and Mixon, 2014]

• A sort of "Generalization" of the Restricted Isometry Property (RIP) which is commonly used in compressed sensing.

• Geometric criterion to guarantee that – under the assumption that the solution space is sparse (in compressible rather than hard sense) – convex optimization yields an accurate approximate solution to an underdetermined, noise-affected linear system.

• We will apply this framework to the ROI problem and verify that this criterion holds.

• We start by defining the appropriate approximation space...

Robust Width Property - Compressed Sensing space

A **compressed sensing (CS) space** $\left(\mathcal{H}, \mathcal{A}, \|\cdot\|_\sharp\right)$ with bound L consists of a Hilbert space $\mathcal{H}$, a subset $\mathcal{A} \subseteq \mathcal{H}$ and a norm or semi-norm $\|\cdot\|_\sharp$ on $\mathcal{H}$ such that

1. $0 \in \mathcal{A}$
2. For every $a \in \mathcal{A}$ and $z \in \mathcal{H}$, there exists a decomposition

$$z = z_1 + z_2$$

such that

$$\|a + z_1\|_\sharp = \|a\|_\sharp + \|z_1\|_\sharp$$

with

$$\|z_2\|_\sharp \leq L\|z_2\|_2$$

Robust Width Property - Definition

A linear operator $\Phi : \mathcal{H} \to \tilde{\mathcal{H}}$ satisfies the (ρ, η) **robust width property (RWP)** over the ball $B_\sharp = \left\{ x \in \mathcal{H} \, : \, \|x\|_\sharp \leq 1 \right\}$ if

$$\|x\|_2 < \rho \, \|x\|_\sharp$$

for every $x \in \mathcal{H}$ such that $\|\Phi x\|_2 \leq \eta \|x\|_2$.

- **Geometric interpretation**

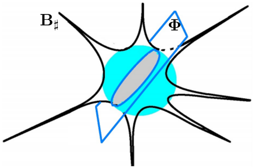

- **width property**: nullspace of Φ intersects $B_\sharp$ with small width.
- **robust width property**: any slight perturbation of the nullspace of Φ satisfies the width property, ensuring stability of the minimizer.

Main Theorem [Goossens et al., 2019]

Let $\mathcal{H} = \left\{(y, f) : \|(y,f)\|_H = \|f\|_2^2 + \|y\|_2^2 < \infty\right\}$, $\mathcal{A} \subset \mathcal{H}$

$$\mathcal{E} = \{(y, f) \in \mathcal{H} : \; y = Wf, \; My = y_0\} \text{ and } \Phi = \begin{pmatrix} I & -W \\ M & 0 \end{pmatrix}$$

Suppose $\left(\mathcal{H}, \mathcal{A}, \|\cdot\|_\sharp\right)$ is a CS space and $\Phi \colon \mathcal{H} \to \tilde{\mathcal{H}}$ satisfies the (ρ, η)-RWP over the ball $B_\sharp$. Then, for every $\left(y^\natural, f^\natural\right) \in \mathcal{E}$, a solution

$$(y^\star, f^\star) = \operatorname{argmin}_{(y,f)\in\mathcal{H}} \|(y, f)\|_\sharp \text{ s.t. } \|My - y_0 - \nu\|_2 \leq \alpha, \; \|y - Wf\|_2 \leq \beta$$

satisfies:

$$\|f^\star - f^\natural\|_2 \leq C_1\sqrt{\alpha^2 + \beta^2} + C_2\rho\gamma \text{ and } \|y^\star - y^\natural\|_2 \leq C_1\sqrt{\alpha^2 + \beta^2} + C_2\rho\gamma,$$

where $C_1 = 2/\eta$ and $C_2 = \inf_{a\in\mathbb{A}} \left\|\left(y^\natural, f^\natural\right) - a\right\|_\sharp$,

provided $\rho \leq \left(\frac{2\gamma}{\gamma-2}L\right)^{-1}$ for some $\gamma > 2$.

Main Theorem [Goossens et al., 2019]

Remarks

• If the ROI problem has a unique solution $(y^\star, f^\star) \in \mathcal{H}$, then Theorem shows that this solution is close to any $\left(y^\natural, f^\natural\right) \in \mathbb{A}$, with error controlled by $\sqrt{\alpha^2 + \beta^2}$.

• If we do not know whether the ROI problem has a unique solution, but $\left(y^\natural, f^\natural\right) \in \mathcal{E}$ ($=$ space of consistent functions satisfying data fidelity), also in this case our solution $(y^\star, f^\star)$ is close to $\left(y^\natural, f^\natural\right)$.

• In case we only obtain an approximate solution $\left(\tilde{y}^\star, \tilde{f}^\star\right)$ with $\left\|\left(\tilde{y}^\star, \tilde{f}^\star\right)\right\|_\sharp \leq \|(y^\star, f^\star)\|_\sharp + \delta$ then a refinement of this theorem gives that $\left(\tilde{y}^\star, \tilde{f}^\star\right)$ is close to $\left(y^\natural, f^\natural\right)$, with a controllable approximation error.

Sparsity-Inducing Reconstruction Algorithm (SIRA)

We construct a CS space $(\mathcal{H}, \mathcal{A}, \|\cdot\|_\sharp)$ where $\mathcal{H} = \left\{(y, f) : \; y \in \ell^2(\mathbb{Z}^2) \; f \in \ell^2(\mathbb{Z}^2)\right\}$, the sparsity norm is

$$\|(y, f)\|_\sharp = \left(\sum_j \left(\sum_i \left| (Ty)_{ij} \right| \right)^2 + \left(\sum_i \left| (TWf)_{ij} \right| \right)^2 \right)^{1/2}$$

where T is the discrete wavelet transform on $\ell^2(\mathbb{Z}^2)$.
Hence the transform basis functions are ridgelets.

The solution space is

$$\mathcal{A} = \left\{ (y, f) \in \ell^2(\mathbb{Z}^2) \,|\, \forall j \in \mathbb{Z} : \; \left((Ty)_{:,j}, (TWf)_{:,j} \right) \text{ is a } K\text{-sparse vector} \right\}$$

Hypotheses of Theorem, including RWP, can be satisfied.

Ridgelets: example

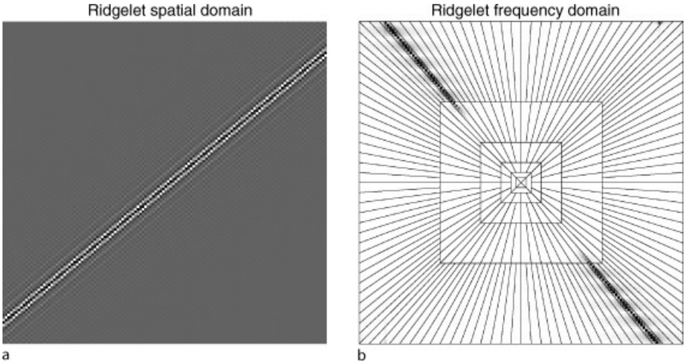

Picture taken from [Fadili and Starck, 2012]

Sparsity-Inducing Reconstruction Algorithm (SIRA)

To solve our constrained optimization problem on (y, f) we use an algorithmic procedure called **Sparsity-Inducing Reconstruction Algorithm (SIRA)** that relies on the Bregman iteration and Bregman divergence.

• We prove that SIRA reaches an approximately sparse solution in a **finite number of steps**, within a predictable distance from the ideal noiseless solution of the ROI problem.

• We found experimentally that the relationship between the ROI radius and the RWP parameters (α, β, L) suggests that the ROI reconstruction performance (in the projection space and image spaces) depends on the ROI radius.

•Accurate reconstruction is guaranteed even for relatively small ROI radii.

Numerical results

An X-O CT system was used to obtain in vivo preclinical data

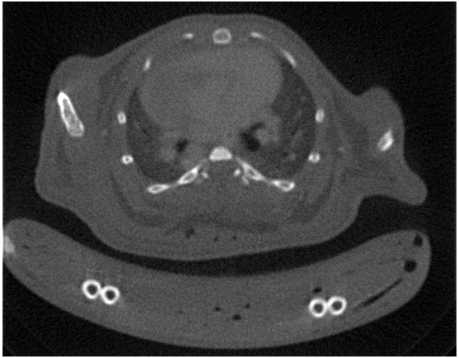

Preclinical - lungs

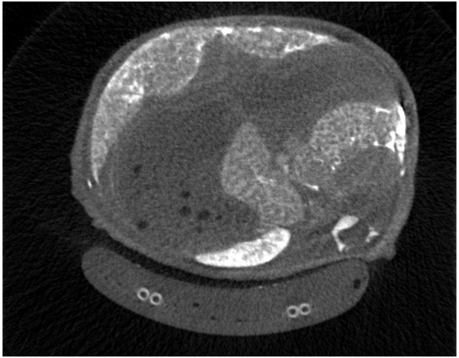

Preclinical - abdomen.

Numerical results

For benchmark comparison, we have considered

- Least-squares conjugate gradient (LSCG) [Hestenes and Stiefel, 1952, Kawata and Nalcioglu, 1985], restricted to the projection ROI $P(S)$.
- Differentiated back-projection (DBP) [Noo et al., 2004], where the Hilbert inversion is performed in the image domain using the 2D Riesz transform.
- Maximum Likelihood Estimation Method (MLEM) [Shepp and Vardi, 1982], restricted to the projection ROI $P(S)$
- Compressed Sensing with respectively TV [Kudo et al., 2013] and ridgelet-based regularization

Numerical results

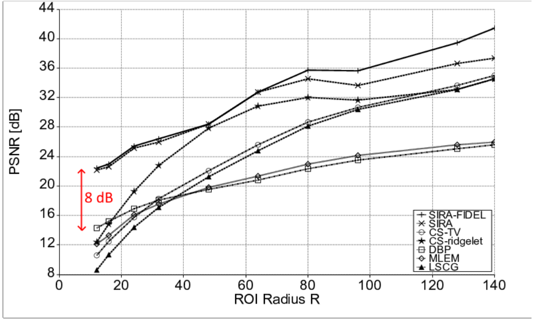

Preclinical - lungs

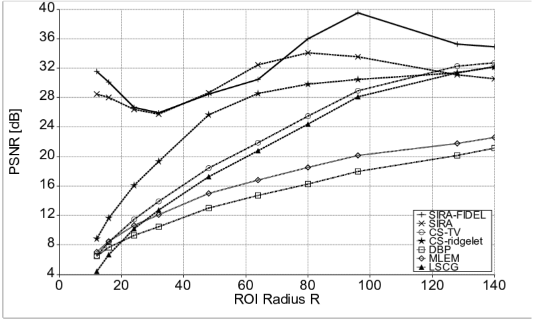

Preclinical - abdomen

Numerical results

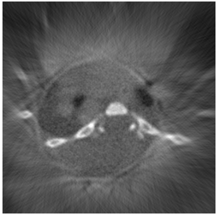

(a) SIRA-FIDEL $\alpha = 0,\ \beta = 0.25u$
$\mathrm{PSNR}_{ROI} = 30.47\mathrm{dB}$

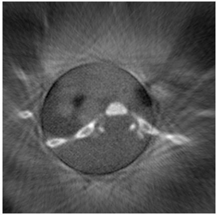

(b) SIRA $\alpha = \frac{1}{30}u,\ \beta = \frac{4}{3}u$
$\mathrm{PSNR}_{ROI} = 35.49\mathrm{dB}$

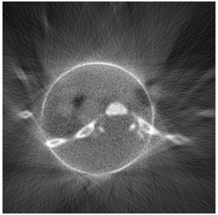

(c) LSCG
$\mathrm{PSNR}_{ROI} = 20.80\mathrm{dB}$

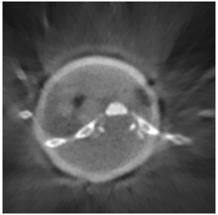

(d) CS-TV
$\mathrm{PSNR}_{ROI} = 21.33\mathrm{dB}$

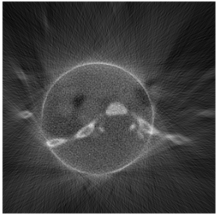

(e) CS-ridgelet
$\mathrm{PSNR}_{ROI} = 28.60\mathrm{dB}$

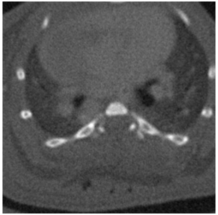

(f) Full view LSCG reconstruction

Numerical results (different ROI radius)

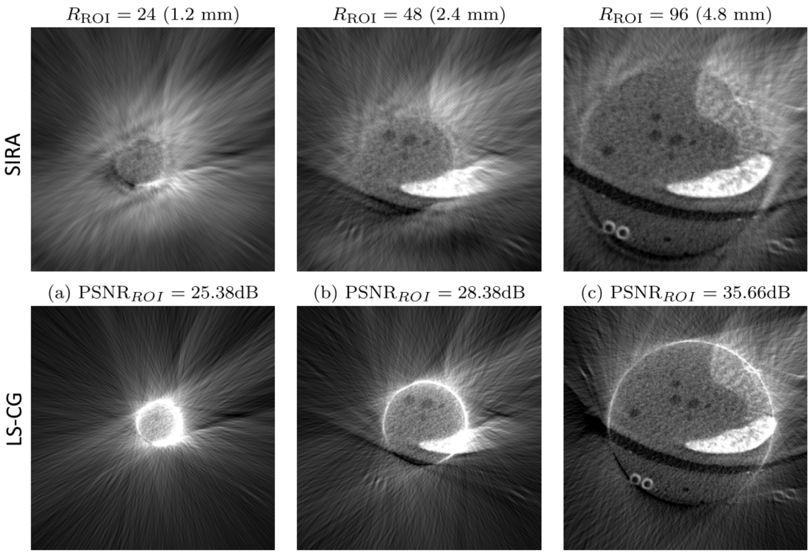

Conclusion

- After almost 50 years of CT reconstruction, it is still a challenging topic!
- The past decade: progress made by application of compressed sensing and sparsity techniques in combination with improvements in computing capabilities (e.g., GPU)
- Robust Width Property: compressed sensing framework, applied to our problem, allow achieving performance guarantees for the image reconstruction error.
- Requires treating the unknown projection data and image jointly during the reconstruction.
- Future: further improvements expected by a combination of sparsity/CS techniques and deep learning techniques.

Research supported by FWO (B. Goossens), NSF (DMS 1720487), GEAR 113491 and by a grant from the Simon Foundation (422488) (D. Labate and B. Bodmann)

Cahill, J. and Mixon, D. G. (2014).
Robust width: A characterization of uniformly stable and robust compressed sensing.
arXiv:1408.4409.

Clackdoyle, R., Noo, F., Guo, J., and Roberts, J. (2004).
Quantitative reconstruction from truncated projections in classical tomography.
IEEE Trans Nuclear Science, 51(5):2570–2578.

Fadili, J. and Starck, J.-L. (2012).
Curvelets and ridgelets.
Computational Complexity: Theory, Techniques, and Applications, pages 754–773.

Goossens, B., Labate, D., and Bodmann, B. (2019).
Robust and Stable Region-Of-Interest CT Reconstruction by Sparsity Inducing Convex Optimization.
Inverse Problems and Imaging.
In review.

Herman, G. T. and Lent, A. (1976).
Iterative reconstruction algorithms.
Computers in Biology and Medicine, 6(4):273–294.

Hestenes, M. R. and Stiefel, E. (1952).
Methods of Conjugate Gradients for Solving Linear Systems.
Journal of Research of the National Bureau of Standards, 6:409–436.

Kawata, S. and Nalcioglu, O. (1985).
Constrained iterative reconstruction by the conjugate gradient method.
IEEE Transactions on Medical Imaging, 4(2):65–71.

Klann, E., Quinto, E., and Ramlau, R. (2015).
Wavelet Methods for a Weighted Sparsity Penalty for Region of Interest Tomography.
Inverse Problems, (31).

Kudo, H., Courdurier, M., Noo, F., and Defrise, M. (2008).
Tiny a priori knowledge solves the interior problem in computed tomography.
Phys. Med. Biol., 53:2207–3923.

Kudo, H., Suzuki, T., and Rashed, E. A. (2013).
Image reconstruction for sparse-view CT and interior CT - introduction to compressed sensing and differentiated backprojection.
Quantitative imaging in medicine and surgery, 3(3):147.

Noo, F., Clackdoyle, R., and Pack, J. (2004).

A two-step Hilbert transform method for 2D image reconstruction.
Phys. Med. Biol., 49:3903–3923.

Shepp, L. A. and Vardi, Y. (1982).
Maximum likelihood reconstruction for emission tomography.
IEEE Trans Med Imaging, 1(2):113–122.

Vandeghinste, B. (2014).
Iterative reconstruction in micro-SPECT/CT: regularized sparse-view CT and absolute in vivo multi-isotope micro-SPECT quantification.
PhD thesis, Ghent University.

Yang, J., Yu, H., Jiang, M., and Wang, G. (2010).
Higher-order total variation minimization for interior tomography.
Inverse Problems, 26:035013 (29p).

Yu, D. F., Fessler, J. A., and Ficaro, E. P. (2000).
Maximum-likelihood transmission image reconstruction for overlapping transmission beams.
IEEE Transactions on Medical Imaging, 19(11):1094–1105.