Computational medical imaging: from model-based approaches to machine learning

Adrian Basarab

Université de Toulouse, IRIT, Université Paul Sabatier Toulouse 3

STSIVA 2019, Bucaramanga, Santander, Colombia

Outline of the talk

What is medical imaging ?

- Visualization of body parts, tissues or organs, for use in clinical diagonsis, treatment and disease monitoring

- Anatomical vs Functional

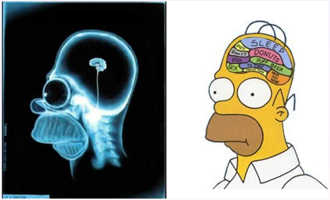

Medical imaging modalities

- Nuclear medicine (SPECT, PET)
- Radiology techniques (X-ray radiography, CT, MRI, Ultrasound)

- Scanners

Computational medical imaging

- Data inversion

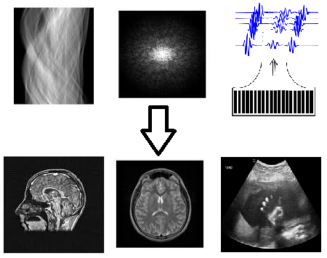

- Incomplete data, non-traditional sensing, etc.

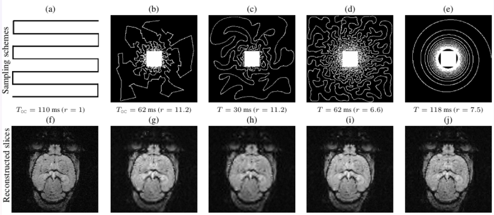

Incomplete data in ultrasound

- Computational methods to compensate for the lack of data

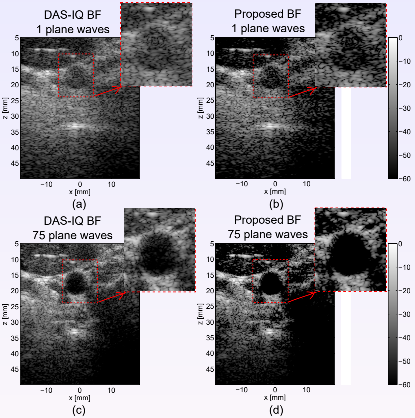

. T. Szasz, A. Basarab, D. Kouamé, Beamforming through regularized inverse problems in ultrasound medical imaging, *IEEE TUFFC*, 2016.

Outline of the talk

Schematic view

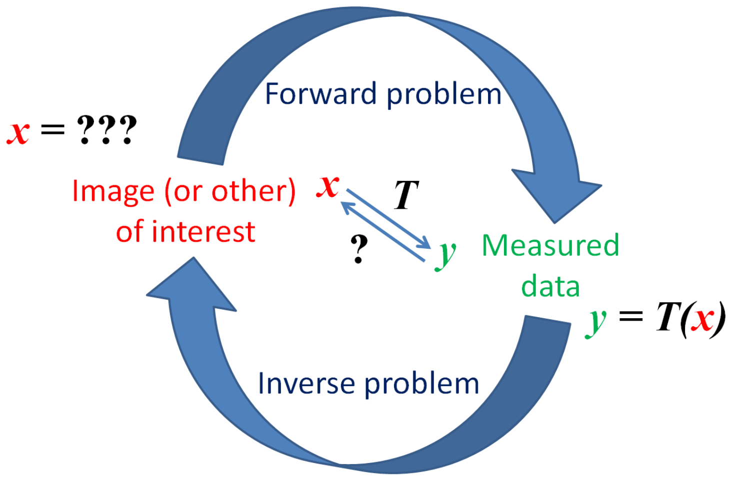

Ill-posedness

$$\mathbf{y} = T(\mathbf{x}) + \mathbf{n}$$

- $\mathbf{y} \in \mathbb{C}^M$ is the observed data (image)
- $\mathbf{x} \in \mathbb{C}^N$ is the image of interest (not observed)
- $\mathbf{n} \in \mathbb{C}^M$ is the noise

T is the observation (forward) operator

- known : estimate $\mathbf{x}$ from $\mathbf{y}$
- unknown : estimate $\mathbf{x}$ and T from $\mathbf{y}$
  - Prior information on T (linear, parametric,...)

Inverse problems in computational medical imaging are usually ill-posed

- T is not invertible
- An infinity of solutions may exist
- A small perturbation on the data may cause an important variation on the estimate (e.g. Fourier measurements)

Outline of the talk

Inversion and regularization

How to chose one (the !) solution from all the possible solutions ?

- Constrain the solution considering penalties
- Need for *a priori* information on $\mathbf{x}$ (regularization)
- Sparse regularization

Are medical images sparse ?

- Contain only a reduced number of non-zero pixels

Distributions promoting sparsity

$$\hat{\mathbf{x}} = \arg\max_{\mathbf{x}} p(\mathbf{x}|\mathbf{y}) = \arg\min_{\mathbf{x}} (-\log(p_y(\mathbf{y}|\mathbf{x})) - \log(p_x(\mathbf{x})))$$

Most common choice

- Laplace distribution

$$\hat{\mathbf{x}} = \arg\min_{\mathbf{x}} \|\mathbf{y} - T(\mathbf{x})\|_2^2 + \lambda\|\mathbf{x}\|_1$$

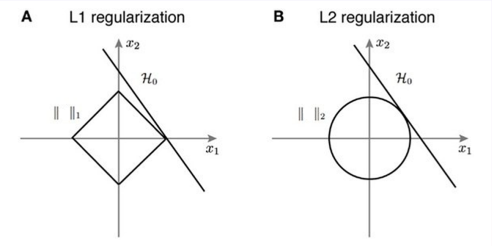

Distributions promoting sparsity

$$\hat{\mathbf{x}} = \arg\max_{\mathbf{x}} p(\mathbf{x}|\mathbf{y}) = \arg\min_{\mathbf{x}} (-\log(p_y(\mathbf{y}|\mathbf{x})) - \log(p_x(\mathbf{x})))$$

Heavy-tailed distributions

- "The tyranny of the normal distribution is that we run the world ... by attributing average levels of competence to the whole population. What really matters is what we do with the tails of the distribution rather than the middle.", R. X. Cringely, Accidental Empires, 1992
- α-stable distribution, with $\alpha < 2$

$$\hat{\mathbf{x}} = \arg\min_{\mathbf{x}} \|\mathbf{y} - T(\mathbf{x})\|_2^2 + \lambda \|\mathbf{x}\|_p^p$$

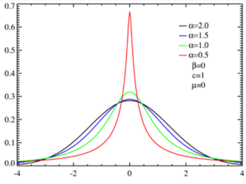

. A. Achim, A. Basarab, G. Tzagkarakis, P. Tsalakides, D. Kouamé, Reconstruction of ultrasound RF echoes modelled as stable random variables, *IEEE TCI*, 2015.

How about medical images that are not sparse ?

Non-adaptive dictionaries to sparsify **x**

- Fourier transform : only adapted to stationary signals
- Short-time Fourier transform, wavelet transform
- Overcomplete representations : curvelet, ridgelet

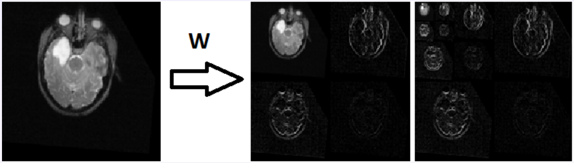

$$\hat{\mathbf{x}} = \arg\min_{\mathbf{x}} \|\mathbf{y} - T(\mathbf{x})\|_2^2 + \lambda \|\mathbf{W}\mathbf{x}\|_1$$

Learned dictionaries

Redundancy and patch self-similarity

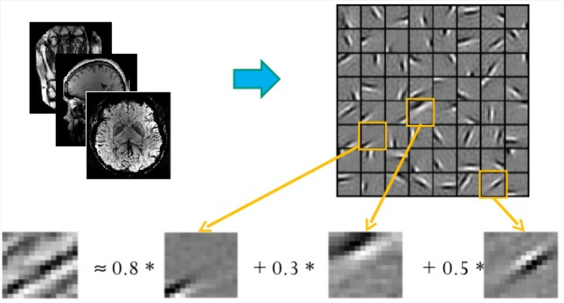

. I. Tosic and P. Frossard, Dictionary learning, IEEE Signal Processing Magazine, 2011.

. Image source : https ://www.slideshare.net/zukun/p02-sparse-coding-cvpr2012-deep-learning-methods-for-vision

Outline of the talk

General path

From the forward model to its inversion

- Establish the forward model T linking the unknown (image) to the data
  - Balance between fidelity to physics and computational tractability
- Define proper prior information about $\boldsymbol{x}$ and the noise
  - Important impact on the solution's pertinence
- Formalize the inverse problem as a cost function minimization
- Stochastic simulation or numerical optimization to find the minimizer
  - Convexity of the cost function
  - Form of the forward operator T
  - Continuous and/or discrete variables
- Is the solution reliable ?

Outline of the talk

Image restoration models

$$\mathbf{y} = \mathbf{SHx} + \mathbf{n}$$

- $\mathbf{x} \in \mathbb{R}^N$: image to reconstruct
- $\mathbf{y} \in \mathbb{R}^M$: observable data
- $\mathbf{H} \in \mathbb{R}^{N \times N}$: 2D convolution matrix

Deconvolution

- $\mathbf{S}$: identity matrix ($M = N$)

Super-resolution

- $\mathbf{S}$: subsampling matrix ($M = d^2 N$)

Compressed deconvolution

- $\mathbf{S}$: random subsampling matrix ($M << N$)

Outline of the talk

Application to ultrasound - Forward model

Approximations

- Linear image formation model, under the first order Born approximation
- Spatially invariant PSF
- Circulant boundary conditions $\rightarrow$ **H** is diagonalizable via Fourier transform

$$y = h \otimes x + n \Leftrightarrow \mathbf{y} = \mathbf{Hx} + \mathbf{n}$$

Example - 3D printed phantom

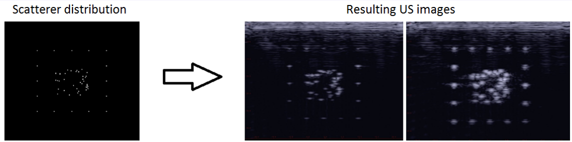

. K. Füzesi, A. Basarab, G. Cserey, D. Kouamé, M. Gyöngy, Validation of image restoration methods on 3D-printed ultrasound phantoms, *IEEE IUS*, 2017.

Hierarchical Bayesian model

Bayesian law

posterior $\propto$ likelihood $\times$ prior

Likelihood and parameter priors

- **Likelihood** : additive white Gaussian noise

$$p(\mathbf{y}|\mathbf{x}, \sigma_n^2) = \frac{1}{(2\pi\sigma_n^2)^{N/2}} \exp\left(-\frac{1}{2\sigma_n^2}\|\mathbf{y} - \mathbf{Hx}\|_2^2\right)$$

- **Priors** :
  - Noise variance σ_n^2 : conjugate inverse gamma ($\mathcal{IG}$) prior

$$p(\sigma_n^2) \sim \mathcal{IG}(\alpha, \nu)$$

Image regularization

- Generalized Gaussian distributions (*GGD*)

$$p(x_i) = \sum_{k=1}^{K} w_k GGD(\xi_k, \gamma_k) \ \text{ with } \ w_k = P(z_i = k)$$
$$\Leftrightarrow x_i | z_i = k \sim GGD(\xi_k, \gamma_k)$$

where $i \in \{1, \cdots, N\}$ and $k \in \{1, \cdots, K\}$, with N the number of pixels and K the number of statistically homogeneous regions.

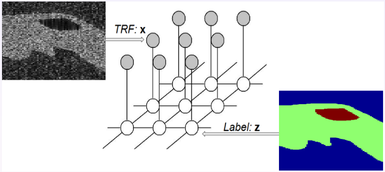

- Continuous (**x**) and discrete (**z**) variables → MCMC algorithm

Hyperparameter priors

- Label map $\mathbf{z}$

$$p(z_i|\mathbf{z}_{-i}) = p(z_i|\mathbf{z}_{\mathcal{V}(i)})$$

Potts model

$$p(\mathbf{z}) = \frac{1}{C(\beta)} \exp\left[\sum_{i=1}^{N} \sum_{i' \in \mathcal{V}(i)} \beta\delta(z_i - z_{i'})\right]$$

- Shape (ξ) and scale (γ) parameters : Uniform and Jeffreys non-informative priors

$$p(\xi) = \prod_{k=1}^{K} p(\xi_k) = \prod_{k=1}^{K} \frac{1}{3}\mathcal{I}_{[0,3]}(\xi_k)$$

$$p(\gamma) = \prod_{k=1}^{K} p(\gamma_k) = \prod_{k=1}^{K} \frac{1}{\gamma_k}\mathcal{I}_{\mathbb{R}+}(\gamma_k)$$

- Posterior distribution Using Bayesian theorem, the joint posterior (target) distribution is

$$p(\mathbf{x}, \mathbf{z}, \xi, \gamma, \sigma_n^2|\mathbf{y}) \propto p(\mathbf{y}|\mathbf{x}, \sigma_n^2, \mathbf{z}, \xi, \gamma)p(\mathbf{x}|\mathbf{z}, \xi, \gamma)p(\mathbf{z}|\xi, \gamma)p(\xi, \gamma)$$

Conditional Distributions

- Noise variance σ_n^2

$$p(\sigma_n^2|\mathbf{y},\mathbf{x},\boldsymbol{\xi},\boldsymbol{\gamma},\mathbf{z}) \propto \mathcal{IG}\left(\alpha + N/2, \theta + \frac{1}{2}\|\mathbf{y}-\mathbf{Hx}\|_2^2\right)$$

- Scale γ

$$p(\gamma_k|\mathbf{x},\boldsymbol{\xi},\mathbf{z},\boldsymbol{\gamma}_{-k}) \propto \mathcal{IG}\left(\frac{N_k}{\xi_k}, \|\mathbf{x}_k\|_{\xi_k}^{\xi_k}\right)$$

- Shape ξ Metropolis Hastings

$$p(\xi_k|\mathbf{x},\boldsymbol{\gamma},\mathbf{z},\boldsymbol{\xi}_{-k}) \propto a_k^{N_k} \exp\left(-\frac{\|\mathbf{x}_k\|_{\xi_k}^{\xi_k}}{\gamma_k}\right) \mathcal{I}_{[0,3]}(\xi_k)$$

Conditional Distributions

- Label map $\mathbf{z}$

$$\pi_{i,k} \propto a_k \exp\left(-\frac{|x_i|^{\xi_k}}{\gamma_k}\right) \exp\left(\sum_{n' \in \mathcal{V}(i)} \beta\delta(k - z_{n'})\right).$$

The normalized conditional probability of the label z_i is defined as

$$\tilde{\pi}_{i,k} = \frac{\pi_{i,k}}{\sum_{k=1}^{K} \pi_{i,k}}.$$

- TRF $\mathbf{x}$ Hamiltonian Monte Carlo (HMC)

$$p(\mathbf{x}|\mathbf{y}, \sigma_n^2, \xi, \gamma, \mathbf{z}) \propto \exp\left(-\frac{\|\mathbf{y} - \mathbf{H}\mathbf{x}\|_2^2}{2\sigma_n^2} - \sum_{k=1}^{K} \frac{\|\mathbf{x}_k\|_{\xi_k}^{\xi_k}}{\gamma_k}\right)$$

Hybrid Gibbs sampler

for $t = 1$ to N_{MC} **do**

▷ Sample the noise variance σ_n^2 according to its conditional distribution.

▷ Sample the scale parameter γ according to its conditional distribution.

▷ Sample the shape parameter ξ using an RWMH algorithm.

▷ Sample the labels $\mathbf{z}$ according to the normalized conditional distribution.

▷ Sample the TRF $\mathbf{x}$ using an HMC method.

end for

Basic idea of Gibbs sampler

Generate samples according to the conditional distributions of variables of interest.

. N. Zhao, A. Basarab, D. Kouamé, J.-Y. Tourneret, Joint deconvolution and segmentation of ultrasound images using a hierarchical Bayesian model based on generalized Gaussian priors, *IEEE TIP*, 2016.

Deconvolution result on *in vivo* data (1/2)

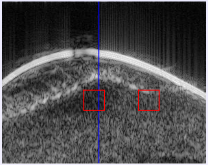

(a) Observation

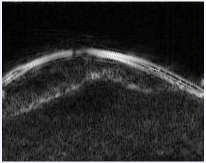

(b) ℓ_2

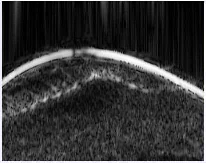

(c) ℓ_1

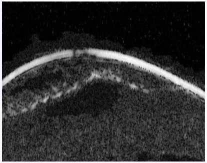

(d) Proposed $\hat{\mathbf{x}}$

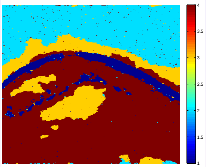

(e) Proposed $\hat{\mathbf{z}}$

Deconvolution result on *in vivo* data (2/2)

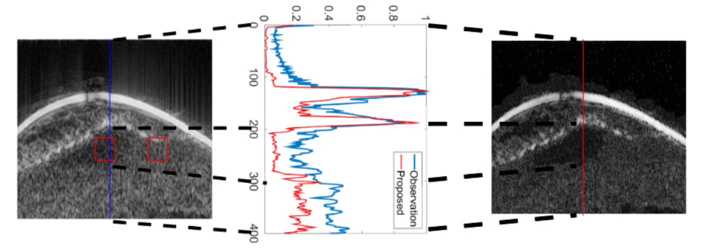

TABLE – Deconvolution Quality for the real US data

Group	Metrics	RF	ℓ_2	ℓ_1	Proposed
Skin melanoma	RG	-	3.01	4.63	10.01
	CNR	1.17	1.09	1.19	1.35
	Time(s)	-	0.007	3.53	1303.4

Outline of the talk

Forward model

$$\mathbf{y} = \mathbf{SHx} + \mathbf{n}$$

- $\mathbf{y} \in \mathbb{R}^{N_l \times 1}$: is the measured image, $N_l = m_l \times n_l$
- $\mathbf{x} \in \mathbb{R}^{N_h \times 1}$: super-resolved image to be estimated, $N_h = d^2 N_l$
- $\mathbf{n} \in \mathbb{R}^{N_l \times 1}$: Gaussian noise

Degradation operators

- $\mathbf{H} \in \mathbb{R}^{N_h \times N_h}$: 2D circulant convolution matrix (PSF of the transducer)
- $\mathbf{S} \in \mathbb{R}^{N_l \times N_h}$: subsampling operator

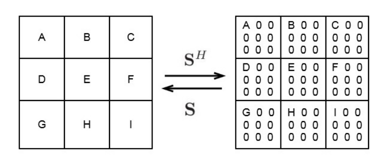

SR optimization problem

$$\min_{\mathbf{x}} \frac{1}{2}\|\mathbf{y}-\mathbf{SHx}\|_2^2 + \tau\phi(\mathbf{Ax})$$

- Total variation or ℓ_p-norm regularization
- Constrained optimization

$$\begin{aligned} \min_{\mathbf{x},\mathbf{u}} \quad & \frac{1}{2}\|\mathbf{y}-\mathbf{SHx}\|_2^2 + \tau\phi(\mathbf{u}) \\ \text{subject to} \quad & \mathbf{Ax} = \mathbf{u} \end{aligned}$$

- Associated augmented Lagrangian function

$$\mathcal{L}(\mathbf{x},\mathbf{u},\boldsymbol{\lambda}) = \frac{1}{2}\|\mathbf{y}-\mathbf{SHx}\|_2^2 + \tau\phi(\mathbf{u}) + \frac{\mu}{2}\|\mathbf{Ax}-\mathbf{u}+\boldsymbol{\lambda}\|_2^2$$

ADMM-based algorithm

Iterate

$\mathbf{x}^{k+1} = \arg\min_{\mathbf{x}} \|\mathbf{y} - \mathbf{SHx}\|_2^2 + \mu\|\mathbf{Ax} - \mathbf{u}^k + \mathbf{d}^k\|_2^2$

$\mathbf{u}^{k+1} = \arg\min_{\mathbf{u}} \tau\phi(\mathbf{u}) + \frac{\mu}{2}\|\mathbf{Ax}^{k+1} - \mathbf{u} + \mathbf{d}^k\|_2^2$

$\mathbf{d}^{k+1} = \mathbf{d}^k + (\mathbf{Ax}^{k+1} - \mathbf{u}^{k+1})$

until stopping criterion is satisfied.

Update **u** *using the proximal operator*

$$\mathrm{prox}_{\lambda,\phi}(\nu) = \arg\min_{x} \phi(x) + \frac{1}{2\lambda}\|x - \nu\|^2.$$

. Parikh & Boyd, Proximal algorithms, *Foundations and Trends in Optimization*, 2014.

. Combettes & Pesquet, Fixed-Point Algorithms for Inverse Problems in Science and Engineering, chapter Proximal splitting methods in signal processing, Springer Optimization and Its Applications, 2011.

ADMM-based algorithm

Iterate
$\mathbf{x}^{k+1} = \arg\min_{\mathbf{x}} \|\mathbf{y} - \mathbf{SHx}\|_2^2 + \mu\|\mathbf{Ax} - \mathbf{u}^k + \mathbf{d}^k\|_2^2$
$\mathbf{u}^{k+1} = \arg\min_{\mathbf{u}} \tau\phi(\mathbf{u}) + \frac{\mu}{2}\|\mathbf{Ax}^{k+1} - \mathbf{u} + \mathbf{d}^k\|_2^2$
$\mathbf{d}^{k+1} = \mathbf{d}^k + (\mathbf{Ax}^{k+1} - \mathbf{u}^{k+1})$
until stopping criterion is satisfied.

Update **x**

- $\ell_2 - \ell_2$ minimization problem

$$\min_{\mathbf{x}} \frac{1}{2}\|\mathbf{y} - \mathbf{SHx}\|_2^2 + \tau\|\mathbf{Ax} - \mathbf{v}\|_2^2$$

- Classical solution $\mathcal{O}(N_h^3)$

$$\hat{\mathbf{x}} = (\mathbf{H}^H\mathbf{S}^H\mathbf{SH} + 2\tau\mathbf{A}^H\mathbf{A})^{-1}(\mathbf{H}^H\mathbf{S}^H\mathbf{y} + 2\tau\mathbf{A}^H\mathbf{v})$$

- **SH** is not diagonalisable in the frequency domain

Proposed closed-form solution

$$\min_{\mathbf{x}} \frac{1}{2}\|\mathbf{y}-\mathbf{SHx}\|_2^2+\tau\|\mathbf{Ax}-\mathbf{v}\|_2^2$$

- Lemma

$$\mathbf{FS}^H\mathbf{SF}^H=\frac{1}{d^2}(\mathbf{J}_d\cdot\mathbf{I}_{m_l})\cdot(\mathbf{J}_d\cdot\mathbf{I}_{n_l})$$

where $\mathbf{J}_d\in\mathbb{R}^{d\times d}$ is a matrix of ones, $\mathbf{I}_{m_l}$ and $\mathbf{I}_{n_l}$ are identity matrices and $\cdot$ is the Kronecker product.

- Proposed solution $\mathcal{O}(N_h\log N_h)$

$$\hat{\mathbf{x}}=\frac{1}{2\tau}\mathbf{F}^H\Psi\mathbf{Fr}-\frac{1}{2\tau}\mathbf{F}^H\Psi\underline{\Lambda}^H\left(2\tau d\mathbf{I}_{N_l}+\underline{\Lambda}\Psi\underline{\Lambda}^H\right)^{-1}\underline{\Lambda}\Psi\mathbf{Fr}$$

where $\mathbf{r}=\mathbf{H}^H\mathbf{S}^H\mathbf{y}+2\tau\mathbf{A}^H\mathbf{v}$, $\Psi=\mathbf{F}\left(\mathbf{A}^H\mathbf{A}\right)^{-1}\mathbf{F}^H$ and $\underline{\Lambda}\in\mathbb{C}^{N_l\times N_h}$ is block diagonal

. N. Zhao, Q. Wei, A. Basarab, N. Dobigeon, D. Kouamé, J.-Y. Tourneret, Fast Single Image Super-resolution using a New Analytical Solution for ℓ_2-ℓ_2 Problems, *IEEE TIP*, 2016

Super-resolution result on *in vivo* data (1/2)

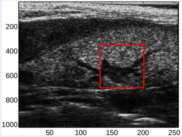

(a) Mouse kidney

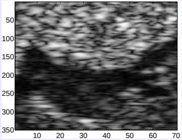

(b) Observation

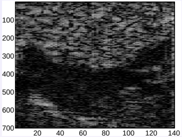

(c) Proposed (ℓ_1)

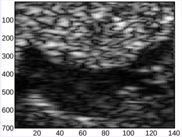

(d) Proposed (ℓ_2)

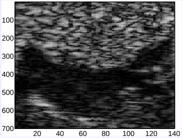

(e) Proposed ($\ell_{1.5}$)

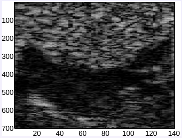

(f) Proposed ($\ell_{4/3}$)

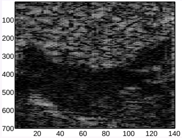

(g) Classical (ℓ_1)

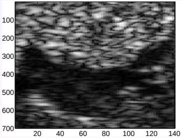

(h) Classical (ℓ_2)

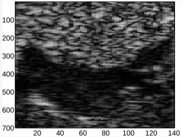

(i) Classical ($\ell_{1.5}$)

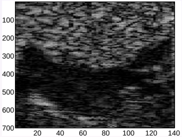

(j) Classical ($\ell_{4/3}$)

Super-resolution result on *in vivo* data (2/2)

TABLE – Numerical results

ℓ_p	Method	RG	Time (s)	Iters.
$p = 2$	Proposed	1.78	0.009	-
	Classical	1.78	0.53	55
$p = 1$	Proposed	16.26	2.42	190
	Classical	16.50	2.58	199
$p = \frac{4}{3}$	Proposed	9.72	0.76	28
	Classical	10.04	1.12	37
$p = \frac{3}{2}$	Proposed	5.55	0.31	14
	Classical	5.72	0.75	33

. M. K. Ng et. al., Solving Constrained Total-variation Image Restoration and Reconstruction Problems via Alternating Direction Methods, *SIAM J. Sci. Comput.*.

. Morin et. al., Alternating direction method of multipliers framework for super-resolution in ultrasound imaging, *Proc. ISBI*, 2012

Outline of the talk

Spatially variant image deconvolution

In ultrasound imaging (but not only), the PSF is not stationary

- Attenuation effects the amplitude of the PSF
- Wave focusing (focused, diverging, plane waves) influence the shape of the PSF

Example for one focus point

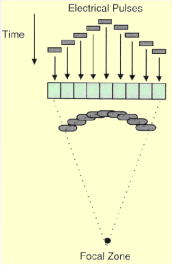

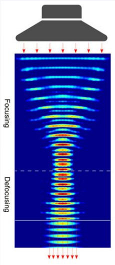

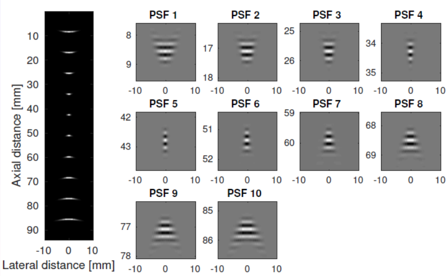

Forward model

One PSF per image pixel

$$\mathbf{y} = \mathbf{HPx} + \mathbf{n}$$

Zero padding operator **P**

- Sparse matrix
- Operator $\mathbf{P} : \mathbb{R}^{m_t \times n_t} \to \mathbb{R}^{m_p \times n_p}$ pads an image with a boundary of width n_r and height m_r, yielding an image of size $m_p = m_t + 2m_r$ times $n_p = n_t + 2n_r$.
- Example of matrix form of 1D padding operator $\mathcal{P}(10, 3)$

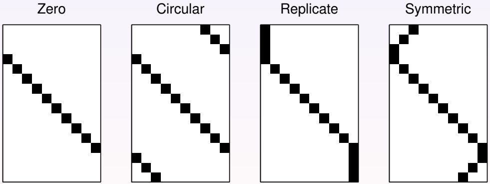

Spatially-variant convolution operator H

Valid and full convolution products

$$\mathcal{C}_1(\boldsymbol{k})\boldsymbol{a} = \boldsymbol{k} *_1 \boldsymbol{a}, \quad \mathcal{C}_2(\boldsymbol{k})\boldsymbol{a} = \boldsymbol{k} *_2 \boldsymbol{a}$$

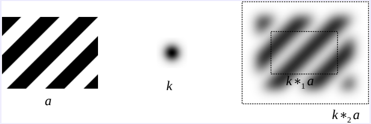

Rotation operator

$$(\mathcal{R}(\boldsymbol{k}))_{i,j} = \boldsymbol{k}_{m_k - i+1, n_k - j+1}$$

Full-width window operator

$$(\mathcal{W}_s(i_1, i_2)\boldsymbol{a})_{i,j} = \boldsymbol{a}_{i+i_1,j}, \; i \in \{0, \ldots, i_2 - i_1\}$$

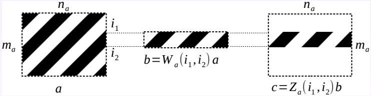

Zero padding operator

$$(\mathcal{Z}_s(i_1, i_2)\boldsymbol{a})_{i,j} = \begin{cases} \boldsymbol{a}_{i-i_1,j}, & i \in \{i_1, \ldots, i_2\}, \\ 0, & \text{otherwise} \end{cases}$$

$$\boldsymbol{H} = \sum_{i_h=1}^{m_t} \mathcal{Z}_t(i_h, i_h)\mathcal{C}_1(\boldsymbol{k}(i_h))\mathcal{W}_p(i_h, i_h + 2m_r)$$

Spatially-variant deconvolution optimization problem

Proximal splitting algorithms

- Data fidelity term $\phi(\boldsymbol{HPx} - \boldsymbol{y})$
- Employ at every iteration of the gradient of the data fidelity term

$$\nabla(\phi(\boldsymbol{HPx} - \boldsymbol{y})) = \boldsymbol{P}^T \boldsymbol{H}^T (\nabla\phi)(\boldsymbol{HPx} - \boldsymbol{y})$$

Adjoint $\boldsymbol{H}^T$ of the axially-variant convolution operator

- Matrix-free expression for the convolution and auxiliary operators

$$(\mathcal{W}_s(i_1, i_2))^T = \mathcal{Z}_s(i_1, i_2), \quad (\mathcal{C}_1(\boldsymbol{k}))^T = \mathcal{C}_2(\mathcal{R}(\boldsymbol{k}))$$

- Therefore

$$\boldsymbol{H}^T = \sum_{i_h=1}^{m_t} \mathcal{Z}_p(i_h, i_h + 2m_r)\mathcal{C}_2(\mathcal{R}(\boldsymbol{k}(i_h)))\mathcal{W}_t(i_h, i_h)$$

. M. I. Florea, A. Basarab, D. Kouamé, S. A. Vorobyov, An axially-variant kernel imaging model applied to ultrasound image reconstruction, *IEEE SPL*, 2018.

Spatially-variant deconvolution simulation result

Elastic-net regularization

- Solved with a FISTA-like algorithm

$$\min_{\boldsymbol{x}} \frac{1}{2}\|\boldsymbol{HPx} - \boldsymbol{y}\|_2^2 + \lambda_1\|\boldsymbol{x}\|_1 + \frac{\lambda_2}{2}\|\boldsymbol{x}\|_2^2.$$

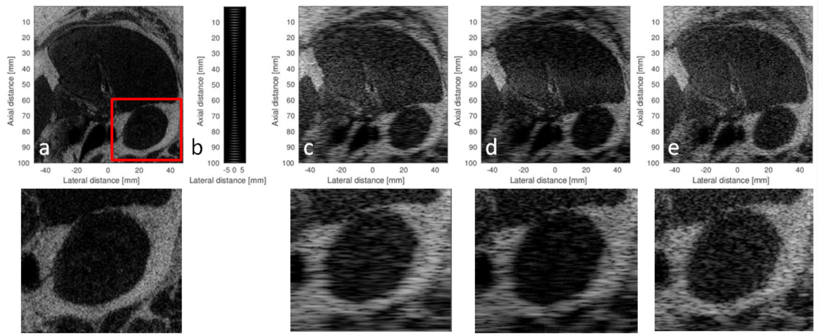

(a) TRF ground truth ; (b) PSFs for 20 depths at regularly spaced intervals of 2 mm ; (c) Observed image ; (d) Invariant deconvolution result ; (e) Variant deconvolution result.

. M. I. Florea, S. A. Vorobyov, An Accelerated Composite Gradient Method for Large-Scale Composite Objective Problems, *IEEE TSP*, 2019.

Outline of the talk

Clinical interest

Complementary medical imaging modalities

- MRI offers a large field of view but with limited spatial resolution
- High-frequency ultrasound offers a good spatial resolution but with limited field of view and poor SNR

Endometriosis diagnosis

Forward models

MRI (low spatial resolution and sampling, Gaussian noise)

$$\mathbf{y}_\mathrm{m} = \mathbf{SHx}_\mathrm{m} + \mathbf{n}_\mathrm{m}$$

Ultrasound (Rayleigh noise)

$$\mathbf{y}_\mathrm{u} = \mathbf{x}_\mathrm{u} + \mathbf{n}_\mathrm{u}$$

Different physical phenomena behind image acquisition

- No one to one correspondence between the gray levels
- US image formation is essentially based on the gradient of acoustic impedance between neighbouring tissues
- US image can thus be seen as a function of the MRI and its gradient

$$\mathbf{x}_\mathrm{u} = f(\mathbf{x}_\mathrm{m}, \nabla \mathbf{n}_\mathrm{m}^H \mathbf{u})$$

$$x_{u,i} = \sum_{p+q\leq 3} c_{pq} x_{m,i}^{p} (\nabla x_m^H u)_i^q$$

. A. Roche, X. Pennec, G. Malandain, and N. Ayache, Rigid registration of 3D ultrasound with MR images : a new approach combining intensity and gradient information, *IEEE Trans. Med. Imaging*, 2001.

MR-US fusion optimization problem

Total variation regularization

$$\hat{\mathbf{x}} = \underset{\mathbf{x}}{\operatorname{argmin}} \quad \frac{1}{2}\|\mathbf{y}_\mathrm{m} - \mathbf{SHx}\|^2 + \lambda\|\nabla\mathbf{x}\|^2 + \tau_2 \sum_{i=1}^{N} \left[\exp(y_{\mathrm{u},i} - f_i(\boldsymbol{x}, \nabla x^H u)) - \gamma(y_{\mathrm{u},i} - f_i(\boldsymbol{x}, \nabla x^H u)) \right]$$

. O. El Mansouri, A. Basarab, F. Vidal, D. Kouamé, J.-Y. Tourneret, Fusion of Magnetic Resonance and Ultrasound Images : a Preliminar Study on Simulated Data, *IEEE ISBI*, 2019.

PALM algorithm

Proximal alternating linearized minimization

$$\min_{(\mathbf{x},\mathbf{v})} \psi(\mathbf{x},\mathbf{v}) := I(\mathbf{x}) + g(\mathbf{v}) + H(\mathbf{x},\mathbf{v})$$

where I and g are continuous convex functions and H may be non-linear.

Step 1 : Take $\gamma_1 > 1$, set $c_k = \gamma_1 L_x(\mathbf{v^k})$

$$\begin{aligned}\mathbf{x^{k+1}} &= \text{prox}^{I}_{c_k}\left(\mathbf{x^k} - \frac{1}{c_k}\nabla_x H(\mathbf{x^k},\mathbf{v^k})\right) \\ &= \arg\min_{\mathbf{x}} (\mathbf{x}-\mathbf{x^k})^H \nabla_x H(\mathbf{x^k},\mathbf{v^k}) + \frac{c_k}{2}\|\mathbf{x}-\mathbf{x^k}\|^2 + I(\mathbf{x})\end{aligned}$$

Step 2 : Take $\gamma_2 > 1$, set $d_k = \gamma_2 L_v(\mathbf{x^k})$

$$\begin{aligned}\mathbf{v^{k+1}} &= \text{prox}^{g}_{d_k}\left(\mathbf{v^k} - \frac{1}{d_k}\nabla_v H(\mathbf{x^k},\mathbf{v^k})\right) \\ &= \arg\min_{\mathbf{v}} (\mathbf{v}-\mathbf{v^k})^H \nabla_v H(\mathbf{x^k},\mathbf{v^k}) + \frac{d_k}{2}\|\mathbf{v}-\mathbf{v^k}\|^2 + g(\mathbf{v})\end{aligned}$$

. J. Bolte, S. Sabach, and M. Teboulle, “Proximal alternating lin-earized minimization or nonconvex and nonsmooth problems,”Mathematical Programming,2014.

MR-US image fusion using PALM

$$\min_{(\mathbf{x},\mathbf{v})} \psi(\mathbf{x},\mathbf{v}) := l(\mathbf{x}) + g(\mathbf{v}) + H(\mathbf{x},\mathbf{v})$$

Auxiliary variable $\mathbf{v}$

$$\begin{aligned}
l(\mathbf{x}) &= \frac{1}{2}\|\mathbf{y}_{\mathrm{m}} - \mathbf{SHx}\|_2^2 + \tau_1\|\nabla\mathbf{x}\|^2 \\
g(\mathbf{v}) &= \tau_2\sum_i[\exp(y_{\mathrm{u},i} - v_i) - \gamma(y_{\mathrm{u},i} - v_i)] + \tau_3\|\nabla\mathbf{v}\|^2 \\
H(\mathbf{x},\mathbf{v}) &= \tau_4\sum_{i=1}^{N}\left(v_i - \sum_{p+q\leq 3} c_{pq} x_i^p (\nabla x^H u)_i^q\right)
\end{aligned}$$

MR-US image fusion using PALM

Input $\mathbf{y_u}$, $\mathbf{y_m}$, $\mathbf{S}$, $\mathbf{H}$, τ_1, τ_2, τ_3, τ_4, γ

– **Estimate** the coefficients of the polynomial function f from $\mathbf{y}_\mathrm{m}$ and $\mathbf{y}_\mathrm{u}$ using the least-square method

REPEAT

1 - Estimate the Lipschitz constant L_{k+1} of $\mathbf{x} \mapsto \nabla_x H(\mathbf{x}, \mathbf{v}^\mathbf{k})$

using the backtracking step size rule as in [6]

Update $\mathbf{x}$ using the analytical solution of [7]

$$\begin{aligned}
\mathbf{x}^{\mathbf{k+1}} &= \mathrm{prox}^{l}_{L_{k+1}}\left(\mathbf{x}^\mathbf{k} - \frac{1}{L_{k+1}}\nabla_x H(\mathbf{x}^\mathbf{k}, \mathbf{v}^\mathbf{k})\right) \\
&= \underset{\mathbf{x}}{\mathrm{argmin}} \;\; \frac{1}{2}\|\mathbf{SHx} - \mathbf{y}_\mathrm{m}\|^2 + \tau_1\|\nabla\mathbf{x}\|^2 \\
&\quad + \frac{L_{k+1}}{2}\|\mathbf{x} - (\mathbf{x}^\mathbf{k} - \frac{1}{L_{k+1}}\nabla_x H(\mathbf{x}^\mathbf{k}, \mathbf{v}^\mathbf{k}))\|^2
\end{aligned}$$

2 - Set $d_k = 2\tau_4$ and update $\mathbf{v}$ using the gradient descent

$$\begin{aligned}
\mathbf{v}^{\mathbf{k+1}} &= \mathrm{prox}^{g}_{d_k}\left(\mathbf{v}^\mathbf{k} - \frac{1}{d_k}\nabla_v H(\mathbf{x}^{\mathbf{k+1}}, \mathbf{v}^\mathbf{k})\right) \\
&= \underset{\mathbf{v}}{\mathrm{argmin}} \;\; \tau_2 \textstyle\sum_i [\exp(y_{\mathrm{u},i} - v_i) - \gamma(y_{\mathrm{u},i} - v_i)] \\
&\quad + \tau_3\|\nabla\mathbf{v}\|^2 + \frac{d_k}{2}\|\mathbf{v} - (\mathbf{v}^\mathbf{k} - \frac{1}{d_k}\nabla_v H(\mathbf{x}^{\mathbf{k+1}}, \mathbf{v}^\mathbf{k}))\|^2
\end{aligned}$$

Until stopping criterion is satisfied

Output: Fused image $\mathbf{x}$

Fusion result on phantom data (1/3)

Homemade phantom mimicking pelvic anatomy

3D represenation

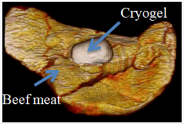

MRI on phantom

MRI *in vivo*

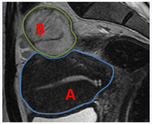

Ultrasound

US on phantom

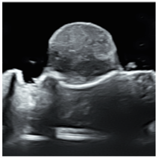

Endometrioma

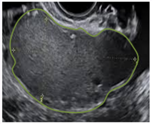

Uterus

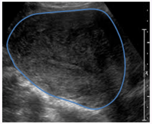

Fusion result on phantom data (2/3)

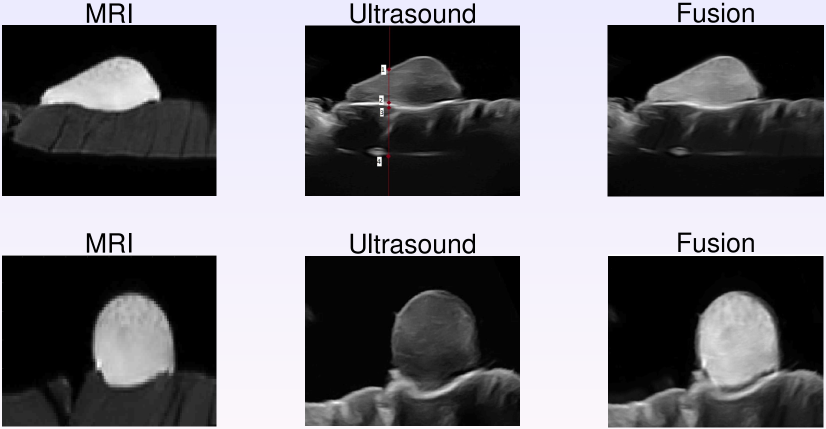

Fusion result on phantom data (3/3)

MRI

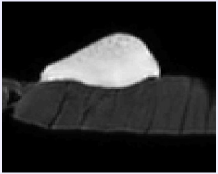

Ultrasound

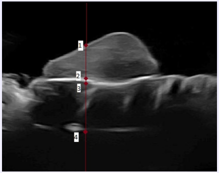

Fusion

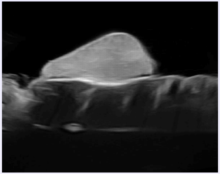

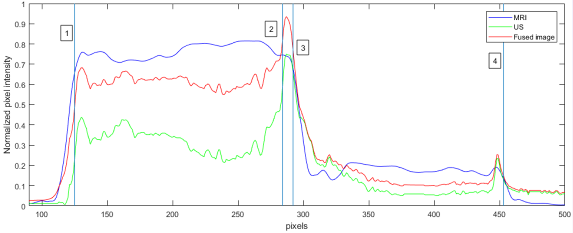

Fusion result on phantom data (3/3)

MRI

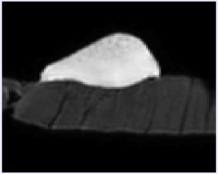

Ultrasound

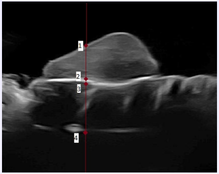

Fusion

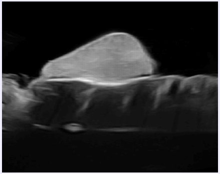

	MRI	US	Fused image
CNR	48.76 dB	20.64 dB	37.73 dB
Interface 1 slope	$2.89e^{-2}$	$7.42e^{-2}$	$7.42e^{-2}$
Interface 2 slope	$-0.10e^{-2}$	$8.89e^{-2}$	$6.86e^{-2}$
Interface 3 slope	$3.57e^{-2}$	$5.47e^{-2}$	$6.61e^{-2}$
Interface 4 slope	$-1.35e^{-2}$	$-1.95e^{-2}$	$-2.05e^{-2}$

Outline of the talk

Basics on acoustic microscopy

Single-element transducer

- Very high frequency : 250 and 500 MHz
- Transmits short ultrasound pulses
- Receives the RF echo signals reflected from the sample

Sample

- Thin section of soft tissue (12 μm) affixed to a microscopy slide

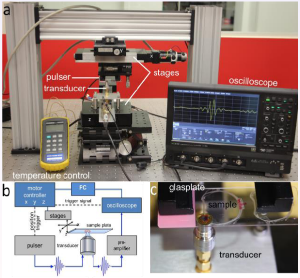

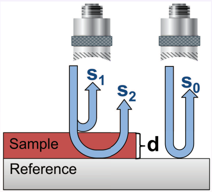

Spatial resolution in acoustic microscopy

Dependent on the central frequency

- Increasing the frequency comes with
  - Increased costs associated with the transducer and then ecessary electronics
  - Experimental difficulties also arise (e.g., sensitivity to nm scale vibrations and temperature)

Example of impedance images on a section of cancerous human lymph node

- Thin section of soft tissue (12 μm) affixed to a microscopy slide

250 MHz

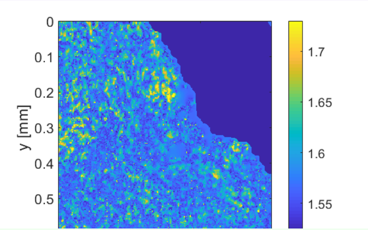

500 MHz

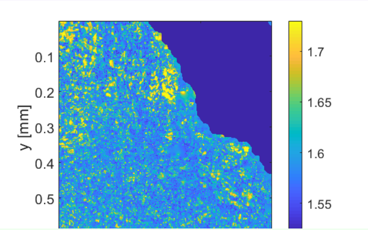

Model-based super-resolution

USAF 1951 resolution phantom

- Super-resolution factor $d = 2$

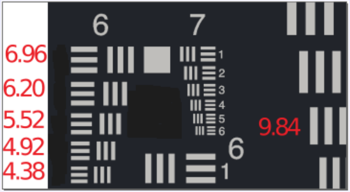

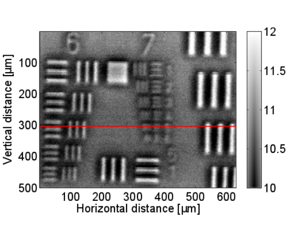

Horizontal profiles

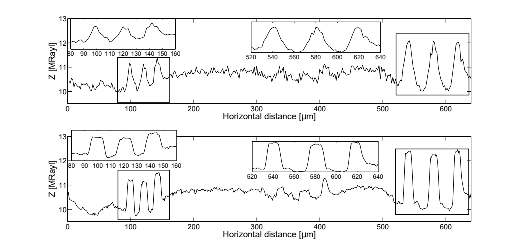

Model-based super-resolution

Fails on ex vivo samples

- Convolution with the PSF not sufficient to model the 250 MHz images

Example of impedance images on a section of cancerous human lymph node

- Thin section of soft tissue (12 μm) affixed to a microscopy slide

250 MHz | 500 MHz | Enhanced 250 MHz

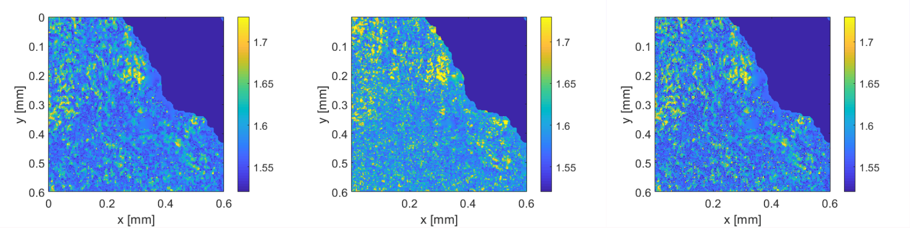

Data-driven super-resolution

Fully convolution neural network (U-net) trained on 250 and 500 MHz images

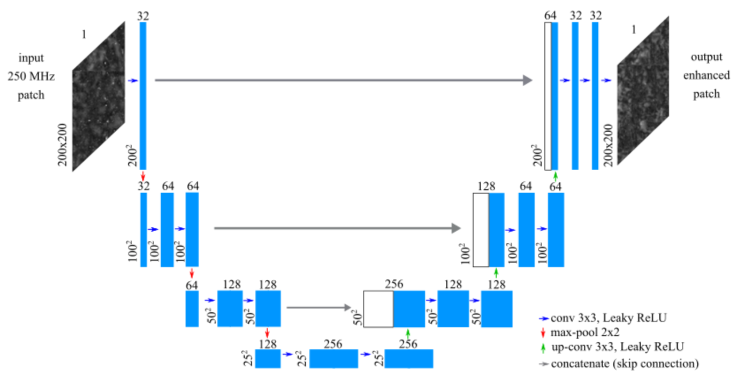

. J. Mamou, T. Pellegrini, D. Kouamé, A. Basarab, Super Resolution in Quantitative Acoustic Microscopy using a U-net like Convolution Neural Network, *IEEE ISBI*, 2019.

Results

250 MHz   500 MHz   Enhanced 250 MHz

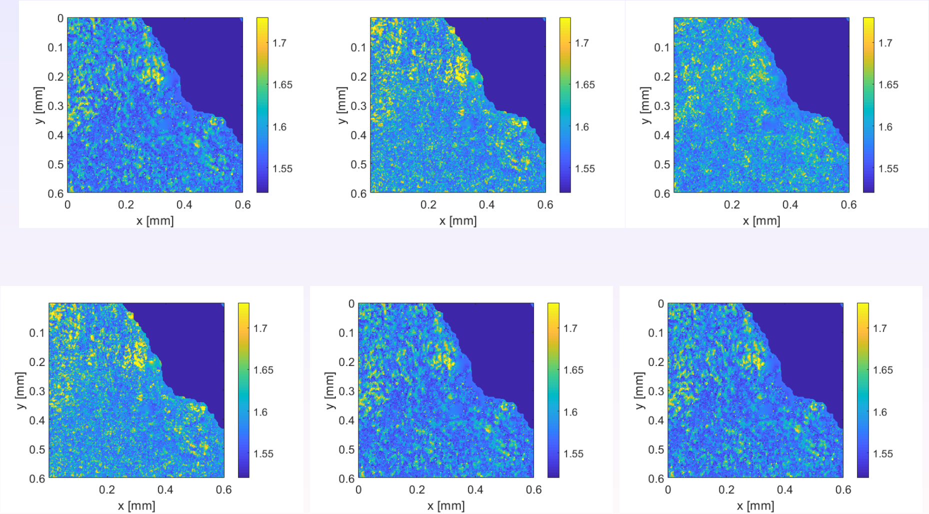

Outline of the talk

Basics on dental Computed Tomography

Cone-beam Computed Tomography

- Available in dental offices, low dose

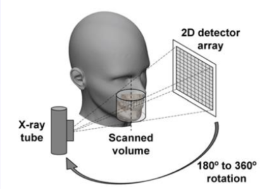

μCT

- Only for extracted teeth, high dose

Application to endodontics

Need for spatial resolution

Segmentation of the root canal

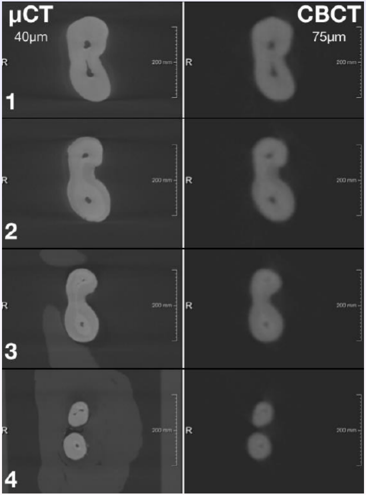

Model-based super-resolution

Fails in this particular application

- Regularization functions not appropriate

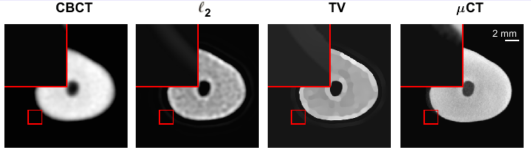

Data-driven super-resolution

Convolution neural network trained on CBCT and μCT data

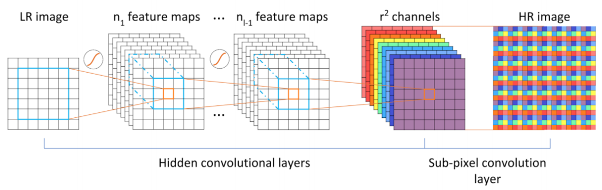

Result on one slice

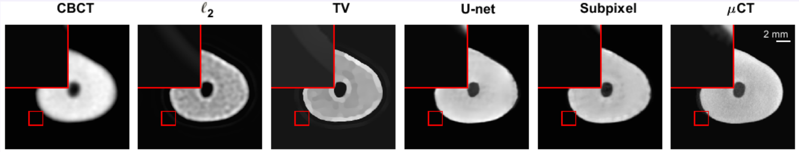

. J. Hatvani, A. Horvath, J. Michetti, A. Basarab, D. Kouamé and M. Gyöngy, Deep Learning-Based Super-Resolution Applied to Dental Computed Tomography, *IEEE TRPMS*, 2019.

Data-driven super-resolution

A more challenging slice

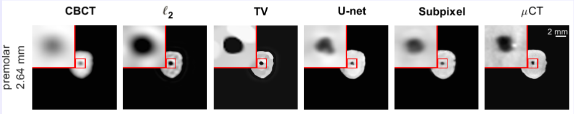

Segmentation result in 3D

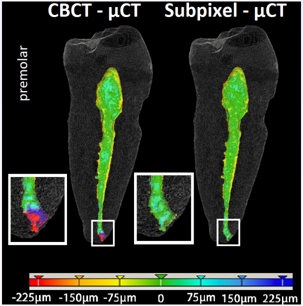

Conclusions

Computational imaging

- In most of medical applications data is not sufficient to form the image (noise, incomplete data)
- Computational methods are used to avoid the ill-posedness of the resulting inverse problem

Model-based approaches

- Models include knowledge about the physics : fidelity, tractability ?
- Regularization terms are required and usually use adaptive or non-adaptive transforms : appropriate choice ?

Data-driven approaches

- More flexibility, but usually require learning databases
- How to include knowledge about the physics ?
- Forward model, regularization, both ?

Computational medical imaging: from model-based approaches to machine learning

Adrian Basarab

Université de Toulouse, IRIT, Université Paul Sabatier Toulouse 3

STSIVA 2019, Bucaramanga, Santander, Colombia