Edge-Preserving Tomographic Reconstruction with Uncertain View Angles

CUQI

sites.dtu.dk/cuqi

Per Christian Hansen, Technical University of Denmark, `pcha@dtu.dk`

J.M. Bardsley, Univ. Montana; Y. Dong & N.A.B. Riis, DTU; F. Uribe, LUT

DTU

VILLUM FONDEN

Setting the Stage – True and Nominal View Angles

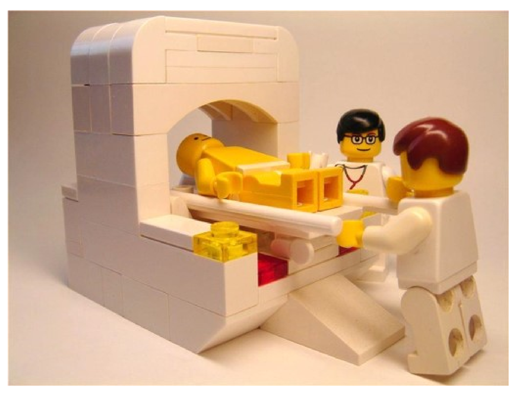

Computed Tomography (CT)

CT data consist of measurements of the attenuation of X-rays passing through an object.
We reconstruct an image of the linear attenuation coefficient of the object's interior.
For each position of the X-ray source, we measure a set of data referred to as a **view**.

The **true view angles** may differ from the assumed **nominal view angles**:

- The model for the measured data is $\boldsymbol{b} = \boldsymbol{A}_{\rm tru}\,\boldsymbol{x} + \boldsymbol{e}$, where $\boldsymbol{e}$ is the measurement noise, $\boldsymbol{x}$ represents the image, and $\boldsymbol{A}_{\rm tru}$ is the forward model for the *unknown* true angles.
- A "naive" and bad reconstruction uses the matrix $\boldsymbol{A}_{\rm nom}$ based on the nominal angles.

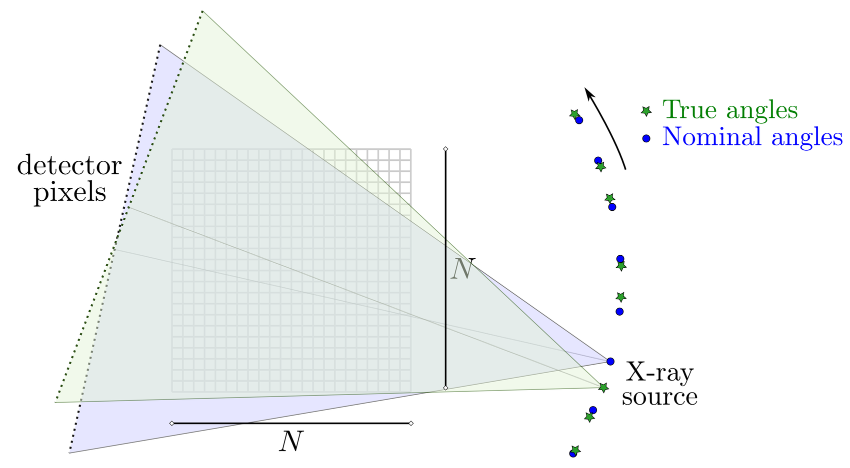

How to Handle Uncertain View Angles – The Bayesian Framework

We consider the true view angles as unknowns $\boldsymbol{\theta}$, together with the image $\boldsymbol{x}$:

$$\text{find } (\boldsymbol{x},\boldsymbol{\theta}) \text{ such that } \boldsymbol{b} = \boldsymbol{A}(\boldsymbol{\theta})\,\boldsymbol{x} + \boldsymbol{e}\ .$$

Here, $\boldsymbol{A}(\boldsymbol{\theta})$ denotes the forward model corresponding to the view angles $\boldsymbol{\theta}$.

We apply the Bayesian framework with a likelihood that involves both $\boldsymbol{x}$ and $\boldsymbol{\theta}$:

$$\pi_{\rm pos}(\boldsymbol{x},\boldsymbol{\theta}) \propto \pi_{\rm lik}(\boldsymbol{b}\,|\,\boldsymbol{x},\boldsymbol{\theta}) \times \pi_{\rm pri}(\boldsymbol{x}) \times \pi_{\rm pri}(\boldsymbol{\theta})\ .$$

- The distribution of $\boldsymbol{e}$ is determined by the measurements; in CT it is log-Poisson and we approximate it by a Gaussian. Hence, $\pi_{\rm lik}(\boldsymbol{b}\,|\,\boldsymbol{x},\boldsymbol{\theta})$ is a **Gaussian**.
- For $\pi_{\rm pri}(\boldsymbol{x})$ we use a **Laplace distribution of the differences of neighbour pixels** (enables sharp edges in the image; related to total variation (TV) regularization.
- For $\pi_{\rm pri}(\boldsymbol{\theta})$ we use the **von Mises distribution** (i.e., a *periodic* normal distribution).

But wait, there's more. We introduce scalar **hyperparameters:** λ in the Gaussian likelihood, δ in the Laplace-difference prior for $\boldsymbol{x}$, and κ in the von Mises prior for $\boldsymbol{\theta}$. All three have exponential distributions $\pi_{\rm hpri}(\cdot) = \beta\,\exp(-\beta\,\cdot)$ with $\beta = 10^{-4}$. Thus, the posterior takes the form

$$\boxed{\pi_{\rm pos}(\boldsymbol{x},\boldsymbol{\theta},\lambda,\delta,\kappa) \propto \pi_{\rm lik}(\boldsymbol{b}\,|\,\boldsymbol{x},\boldsymbol{\theta},\lambda) \times \pi_{\rm pri}(\boldsymbol{x}\,|\,\delta) \times \pi_{\rm pri}(\boldsymbol{\theta}\,|\,\kappa) \times \pi_{\rm hpri}(\lambda)\,\pi_{\rm hpri}(\delta) \times \pi_{\rm hpri}(\kappa)}$$

How to Sample – A Hybrid Gibbs Sampler

Performing statistical inference of the full posterior $\pi_{\rm pos}(\boldsymbol{x},\boldsymbol{\theta},\lambda,\delta,\kappa)$ is challenging: the number of pixels n is large, the forward model $\boldsymbol{A}(\boldsymbol{\theta})$ is nonlinear in the view angles $\boldsymbol{\theta}$, and the prior $\pi_{\rm pri}(\boldsymbol{x}\,|\,\delta)$ is nondifferentiable due to the 1-norm. We split the posterior and apply different samplers for each parameter, hence the sampler is hybrid.

$$\pi_1(\boldsymbol{x}\,|\,\boldsymbol{\theta},\lambda,\delta) \propto \exp\left(-\frac{\lambda}{2}\|\boldsymbol{A}(\boldsymbol{\theta})\,\boldsymbol{x}-\boldsymbol{b}\|_2^2 - \delta(\|(\boldsymbol{I}\otimes\boldsymbol{D})\,\boldsymbol{x}\|_1 + \|(\boldsymbol{D}\otimes\boldsymbol{I})\,\boldsymbol{x}\|_1)\right)$$

$$\pi_2(\boldsymbol{\theta}\,|\,\boldsymbol{x},\lambda,\kappa) \propto \exp\left(-\frac{\lambda}{2}\|\boldsymbol{A}(\boldsymbol{\theta})\,\boldsymbol{x}-\boldsymbol{b}\|_2^2 + \kappa\boldsymbol{1}^{\rm T}\cos(\boldsymbol{\theta}-\bar{\boldsymbol{\theta}})\right)$$

$$\pi_3(\lambda\,|\,\boldsymbol{x},\boldsymbol{\theta}) \propto \lambda^{m/2}\exp\left(-\lambda\left[\frac{1}{2}\|\boldsymbol{A}(\boldsymbol{\theta})\,\boldsymbol{x}-\boldsymbol{b}\|_2^2+\beta\right]\right)$$

$$\pi_4(\delta\,|\,\boldsymbol{x}) \propto \delta^n\exp\left(-\delta[\|(\boldsymbol{I}\otimes\boldsymbol{D})\,\boldsymbol{x}\|_1 + \|(\boldsymbol{D}\otimes\boldsymbol{I})\,\boldsymbol{x}\|_1+\beta]\right)$$

$$\pi_5(\kappa\,|\,\boldsymbol{\theta}) \propto I_0(\kappa)^{-p}\exp\left(-\kappa\left[-\boldsymbol{1}^{\rm T}\cos(\boldsymbol{\theta}-\bar{\boldsymbol{\theta}})+\beta\right]\right)$$

$\boldsymbol{b}\in\mathbb{R}^m$, $\boldsymbol{x}\in\mathbb{R}^n$, $\boldsymbol{\theta}\in\mathbb{R}^p$, I_0 = 0-order mod. Bessel funct., $\boldsymbol{D}$ = bidiag$(-1,1)$, $\bar{\boldsymbol{\theta}}$ = nominal angles.

Initial states $\boldsymbol{x}^{(0)}, \boldsymbol{\theta}^{(0)}, \lambda^{(0)}, \delta^{(0)}, \kappa^{(0)}$

For $j = 1, 2, \ldots, N_{\rm samp}$

  Sample attenuation coefficients

  $\boldsymbol{x}^{(j)} \sim \pi_1(\cdot\,|\,\boldsymbol{\theta}^{(j-1)}, \lambda^{(j-1)}, \delta^{(j-1)})$

  Sample view angles

  $\boldsymbol{\theta}^{(j)} \sim \pi_2(\cdot\,|\,\boldsymbol{x}^{(j)}, \lambda^{(j-1)}, \kappa^{(j-1)})$

  Sample hyperparameters

  $\lambda^{(j)} \sim \pi_3\left(\cdot\,|\,\boldsymbol{x}^{(j)},\boldsymbol{\theta}^{(j)}\right)$, $\delta^{(j)} \sim \pi_4\left(\cdot\,|\,\boldsymbol{x}^{(j)}\right)$, $\kappa^{(j)} \sim \pi_5\left(\cdot\,|\,\boldsymbol{\theta}^{(j)}\right)$

End

Simulation Results – View Angles

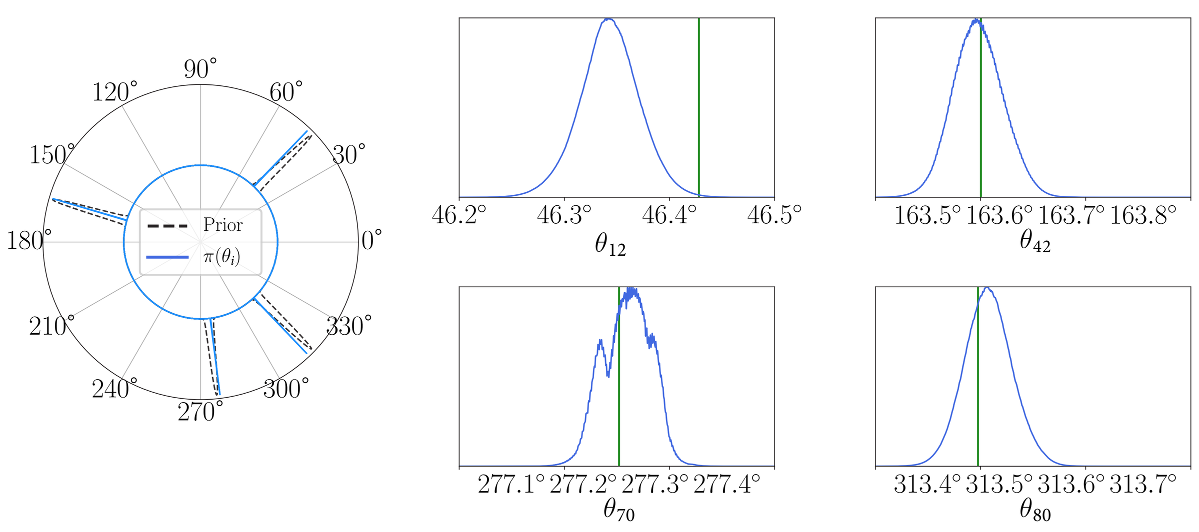

Left: von Mises prior with the respective densities for selected angles in $\boldsymbol{\theta}$.
Right: some component densities and true angles shown as vertical green lines.

Some Details of the Algorithm

π_1: non-differentiable due to $\|\cdot\|_1$ and nonlinear in $\boldsymbol{\theta}$. Use *Laplace's approximation*, i.e., a Gaussian $\pi_{\rm G} = \mathcal{N}(\boldsymbol{x};\boldsymbol{\mu},\boldsymbol{H}^{-1})$ with $\boldsymbol{H}(\boldsymbol{x}^{(j-1)}) \approx$ Hessian of $-\log\pi_1$ and

$$\boldsymbol{\mu} = \boldsymbol{\mu}(\boldsymbol{x}) = \lambda\,\boldsymbol{H}^{-1}(\boldsymbol{x})\,\boldsymbol{A}(\boldsymbol{\theta})^{\rm T}\boldsymbol{b} = \text{MAP estimator of } \pi_1,$$

Much easier to work with a Gaussian but we miss the heavy tails of π_1. We use 10 CGLS iterations to compute the LS solution that gives the sample $\boldsymbol{x}^{(j)}$.

π_2: samples from π_2 are drawn sequentially by componentwise Metropolis:

$$\theta_1^{[k+1]} \sim \pi_2\left(\boldsymbol{\theta}\,|\,\boldsymbol{x},\lambda,\kappa,\,[\theta_2^{[k]},\theta_3^{[k]},\ldots,\theta_p^{[k]}]\right),$$

$$\theta_2^{[k+1]} \sim \pi_2\left(\boldsymbol{\theta}\,|\,\boldsymbol{x},\lambda,\kappa,\,[\theta_1^{[k+1]},\theta_3^{[k]},\ldots,\theta_p^{[k]}]\right),$$

$$\vdots$$

$$\theta_p^{[k+1]} \sim \pi_2\left(\boldsymbol{\theta}\,|\,\boldsymbol{x},\lambda,\kappa,\,[\theta_1^{[k+1]},\theta_2^{[k+1]},\ldots,\theta_{p-1}^{[k+1]}]\right).$$

After 20 cycles we obtain $\boldsymbol{\theta}^{(j)} = \boldsymbol{\theta}^{[20]}$.

π_3 and π_4: can be written and approximated, respectively, in closed form.

π_5: sampled with standard random-walk Metropolis.

Simulation Results – Metallic Grains Phantom

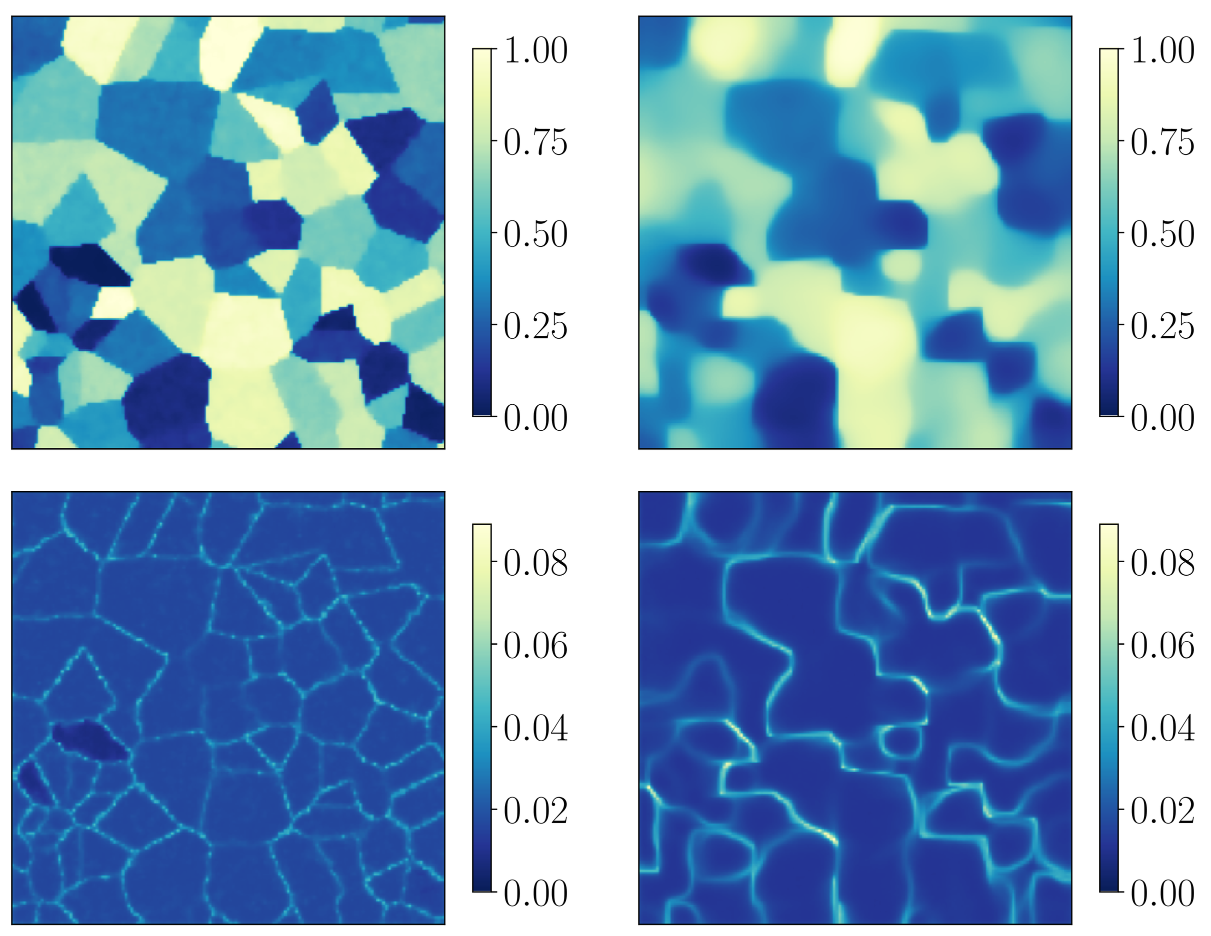

Left: our method.
Right: using the incorrect nominal angles.
Top: posterior mean
Bottom: st. dev.

Nominal angles give a blurry image with uncertain boundaries.

We compute a sharper image with uncertainty confined to pixels on grain boundaries.

Appendix – Definition of Priors

$$\pi_{\rm pri}(\boldsymbol{x}\,|\,\delta) = \left(\frac{\delta}{2}\right)^n\exp\left(-\delta(\|(\boldsymbol{I}\otimes\boldsymbol{D})\,\boldsymbol{x}\|_1 + \|(\boldsymbol{D}\otimes\boldsymbol{I})\,\boldsymbol{x}\|_1)\right)$$

$$\pi_{\rm pri}(\boldsymbol{\theta}\,|\,\kappa) = \left(\frac{1}{2\pi I_0(\kappa)}\right)^p\exp\left(\kappa\,\boldsymbol{1}^{\rm T}\cos(\boldsymbol{\theta}-\bar{\boldsymbol{\theta}})\right)$$

Reference and Funding

F. Uribe, J. M. Bardsley, Y. Dong, P. C. Hansen, N. Riis, *A hybrid Gibbs sampler for edge-preserving tomographic reconstruction with uncertain view angles*, SIAM/ASA J. Uncertain. Quantific., 10 (2022), pp. 1293–1320, doi: 10.1137/21M1412268.

A part of the project Computational Uncertainty Quantification for Inverse problems (CUQI), funded by Villum Investigator grant no. 25893 from The Villum Foundation.