Structural Gaussian Priors for Bayesian CT reconstruction of Subsea Pipes

Silja W. Christensen, Nicolai A. B. Riis, Felipe Uribe, and Jakob S. Jørgensen

Department of Applied Mathematics and Computer Science
Technical University of Denmark

CT inspection of subsea pipes

- Non-destructive inspection of subsea pipes in operation to monitor pipe integrity.
- Data acquisition is time-consuming and costly.
- Time and cost savings with fewer projections.
- Reduced quality of reconstructions using conventional methods.
- Inform the reconstruction with available structural information about the pipe geometry.

Subsea Oil Pipe

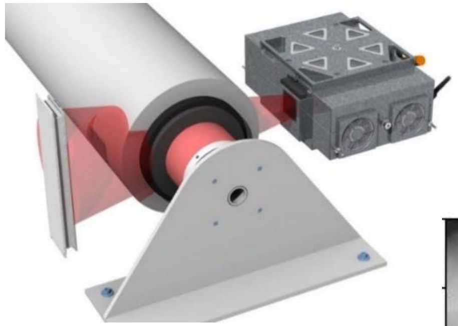

CT scanner

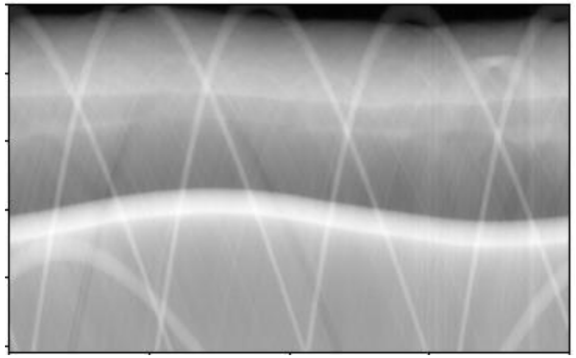

Sinogram

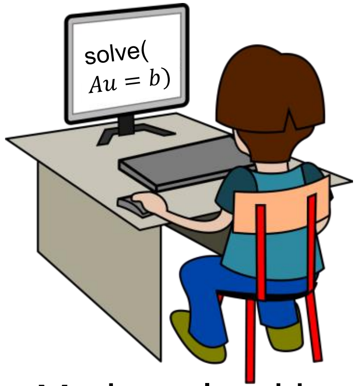

Math + algorithms

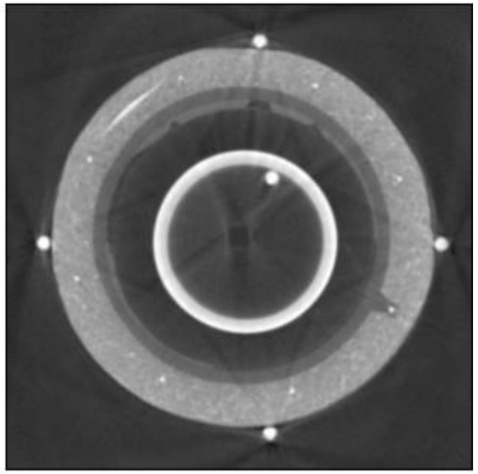

CT reconstruction

Bayesian CT problem

CT model:

$$\mathbf{d} = \mathbf{A}\mathbf{x} + \mathbf{e}, \qquad \mathbf{e} \sim \mathcal{N}(\mathbf{0}, \lambda^{-1}\mathbf{I})$$

$\mathbf{d} \in \mathbb{R}^m$	Observed X-ray absorption.
$\mathbf{x} \in \mathbb{R}^n$	Unknown linear attenuation coefficients in the reconstruction domain with $N \times N$ pixels.
$\mathbf{A} \in \mathbb{R}^{m \times n}$	Contains entries a_{ij} that describe the intersection of ray i with pixel j.
$\mathbf{e} \in \mathbb{R}^m$	Independent and identically distributed Gaussian noise with mean zero and variance λ^{-1}.

Bayesian reconstruction:

$$\pi_{\text{post}}(\mathbf{x}|\mathbf{d}) \propto \pi(\mathbf{d}|\mathbf{x})\, \pi_{\text{pr}}(\mathbf{x})$$

Posterior: Distribution of possible solutions.

Likelihood: Represents data misfit.

Prior: Contains some a-priori information.

A-priori structural information

Layered pipe structure:

- p non-overlapping regions of different materials.

Known pipe materials and X-ray beam energy:

- Linear attenuation coefficients are given: $\alpha = \kappa(E)\rho$

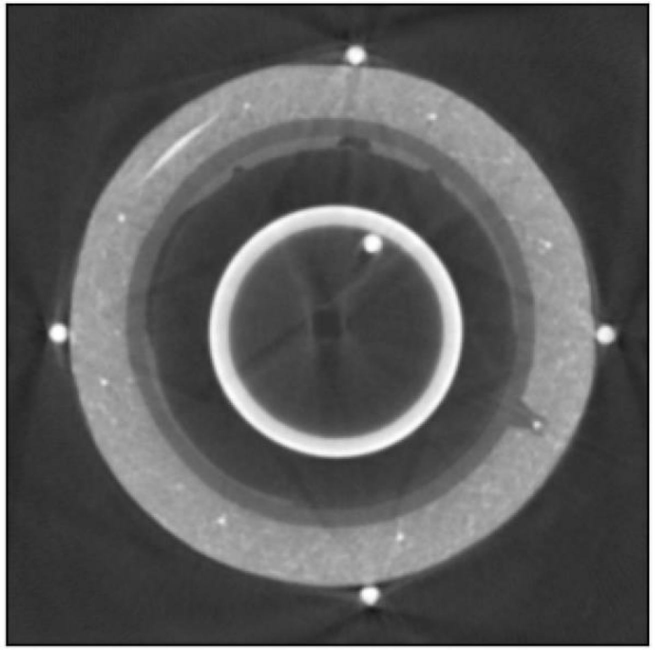

Gaussian Priors

Gaussian prior density (intrinsic)[1]:

$$\pi_i(\mathbf{x}) = \frac{1}{(2\pi)^{(n-k)/2}} \left(|\mathbf{R}_i^T \mathbf{R}_i|^*\right)^{1/2} \exp\left(-\frac{1}{2}\|\mathbf{R}_i(\mathbf{x}-\boldsymbol{\mu}_i)\|_2^2\right)$$

- With mean $\boldsymbol{\mu}_i \in \mathbb{R}^{l_i}$ and precision matrix $\left(\mathbf{R}_i^T \mathbf{R}_i\right) \in \mathbb{R}^{n\times n}$ with rank $(n-k)$.

GMRF: Gaussian Markov Random Field

$$\boldsymbol{\mu}_0 = \mathbf{0} \qquad \text{and} \qquad \mathbf{R}_0 = \sqrt{\delta_0}\begin{bmatrix}\mathbf{D}_1\\ \mathbf{D}_2\end{bmatrix}$$

- Favors neighboring pixels with the same values.
- Employed on entire image domain to promote roughly homogeneous materials.

IID: Independent and identically distributed Gaussian

$$\boldsymbol{\mu}_i = \alpha_i \mathbf{1} \qquad \text{and} \qquad \mathbf{R}_i = \sqrt{\delta_i}\mathbf{M}_i$$

- Favors posterior pixel values around the prior mean.
- Employed in each masked region to promote the expected attenuation coefficients of that material.

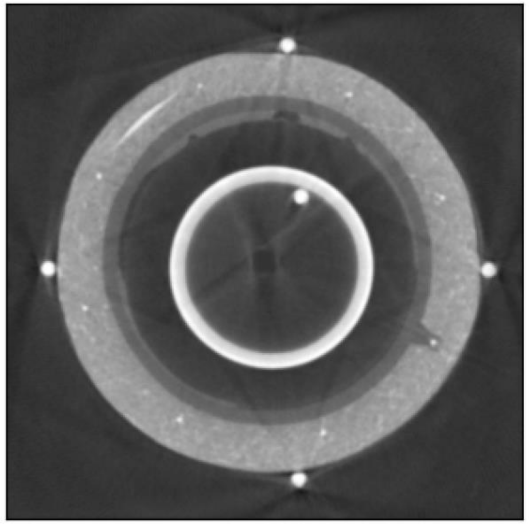

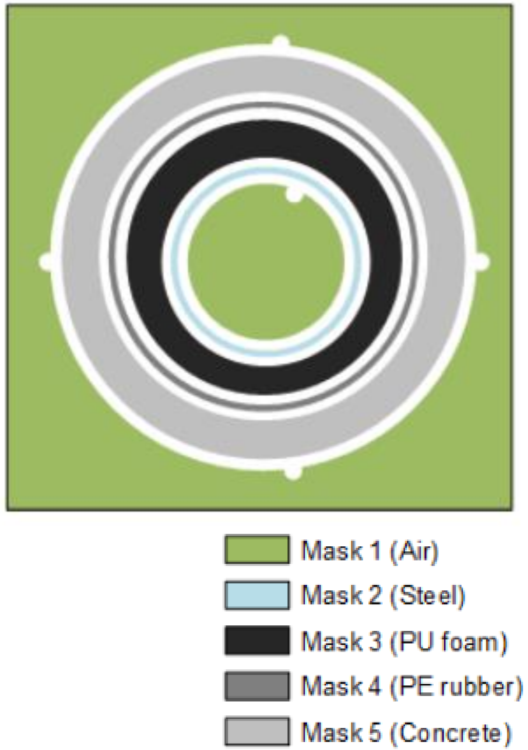

[1]Rue H and Held L Gaussian Markov Random Fields Theory and Applications (New York: Chapman & Hall/CRC) ISBN 1-54488-432-0

Structural Gaussian priors

Structural Gaussian Prior expression:

Assuming independent priors, the joint prior is Gaussian with mean and square-root precision:

$$\mathbf{R}_{\mathrm{pr}} = \begin{bmatrix} \mathbf{R}_0 \\ \vdots \\ \mathbf{R}_p \end{bmatrix} \quad \text{and} \quad \boldsymbol{\mu}_{\mathrm{pr}} = \left(\mathbf{R}_{\mathrm{pr}}^T \mathbf{R}_{\mathrm{pr}}\right)^{-1} \sum_{i=0}^{p} \mathbf{R}_i^T \mathbf{R}_i \boldsymbol{\mu}_i$$

SGP configurations:

GMRF: GMRF prior on the whole domain.

SGP-BG: GMRF prior on the whole domain + IID prior on the background.

SGP-F: GMRF prior on the whole domain + IID prior on the background + IID priors on all pipe layers.

GMRF:

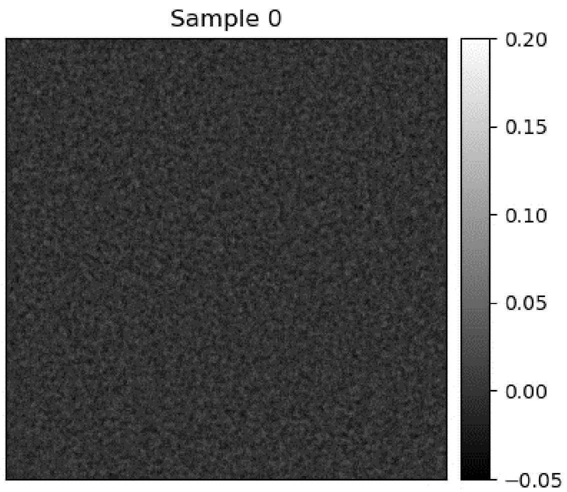

SGP-BG:

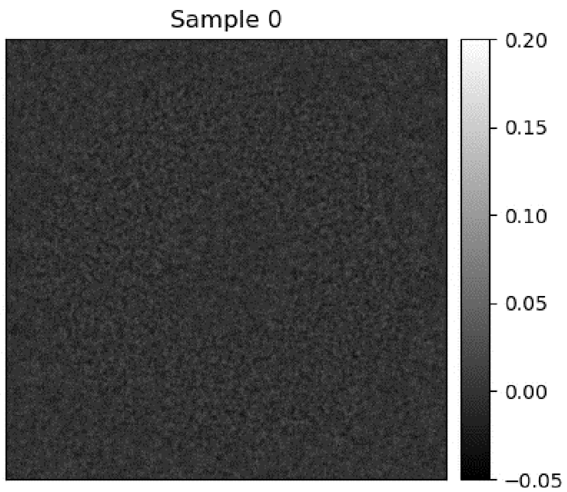

SGP-F:

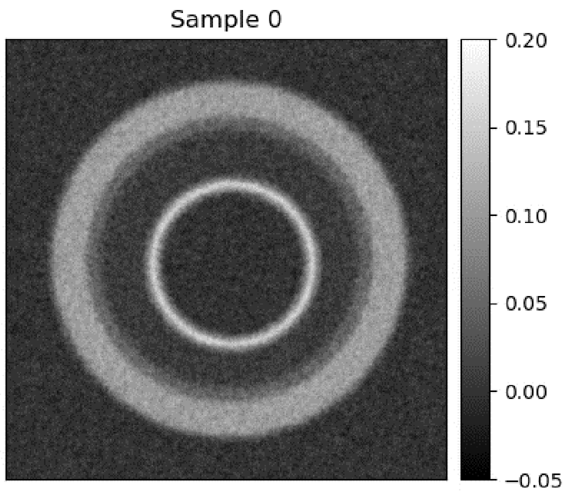

Posterior sampling

Recall:
Model: $\mathbf{d} = \mathbf{A}\mathbf{x} + \mathbf{e}, \quad \mathbf{e} \sim \mathcal{N}(\mathbf{0}, \lambda^{-1}\mathbf{I})$

Prior: $\mathbf{R}_{\text{pr}} = \begin{bmatrix} \mathbf{R}_0 \\ \vdots \\ \mathbf{R}_p \end{bmatrix}$ and $\boldsymbol{\mu}_{\text{pr}} = \left(\mathbf{R}_{\text{pr}}^T \mathbf{R}_{\text{pr}}\right)^{-1} \sum_{i=0}^{p} \mathbf{R}_i^T \mathbf{R}_i \boldsymbol{\mu}_i$

Analytic posterior:

$$\mathbf{x} \mid \mathbf{d} \sim \mathcal{N}\left(\boldsymbol{\mu}_{\text{post}}, \left(\mathbf{R}_{\text{post}}^T \mathbf{R}_{\text{post}}\right)^{-1}\right)$$

$$\mathbf{R}_{\text{post}} = \begin{bmatrix} \sqrt{\lambda}\mathbf{A} \\ \mathbf{R}_0 \\ \vdots \\ \mathbf{R}_p \end{bmatrix} \qquad \text{and} \qquad \mathbf{R}_{\text{post}}^T \mathbf{R}_{\text{post}} \boldsymbol{\mu}_{\text{post}} = \lambda \mathbf{A}^T \mathbf{d} + \sum_{i=0}^{p} \mathbf{R}_i^T \mathbf{R}_i \boldsymbol{\mu}_i$$

- The posterior mean can be found by solving the equation above.
- The posterior precision is impractical to compute because it involves forming $\mathbf{A}^T\mathbf{A}$.

Fast posterior sampling using Linear Randomize-Then-Optimize (RTO):

- Generate sample from the posterior by solving iteratively with CGLS:

$$\mathbf{R}_{\text{post}} \mathbf{x}^* = \begin{bmatrix} \sqrt{\lambda}\mathbf{A} \\ \mathbf{R}_0 \boldsymbol{\mu}_0 \\ \vdots \\ \mathbf{R}_p \boldsymbol{\mu}_p \end{bmatrix} + \boldsymbol{\xi}, \qquad \boldsymbol{\xi} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$$

Linear RTO see e.g. Bardsley: Uncertainty Quantification for Inverse Problems, SIAM, Philadelphia (2018)

Numerical experiments with synthetic data

- Pipe phantom representing the real pipe

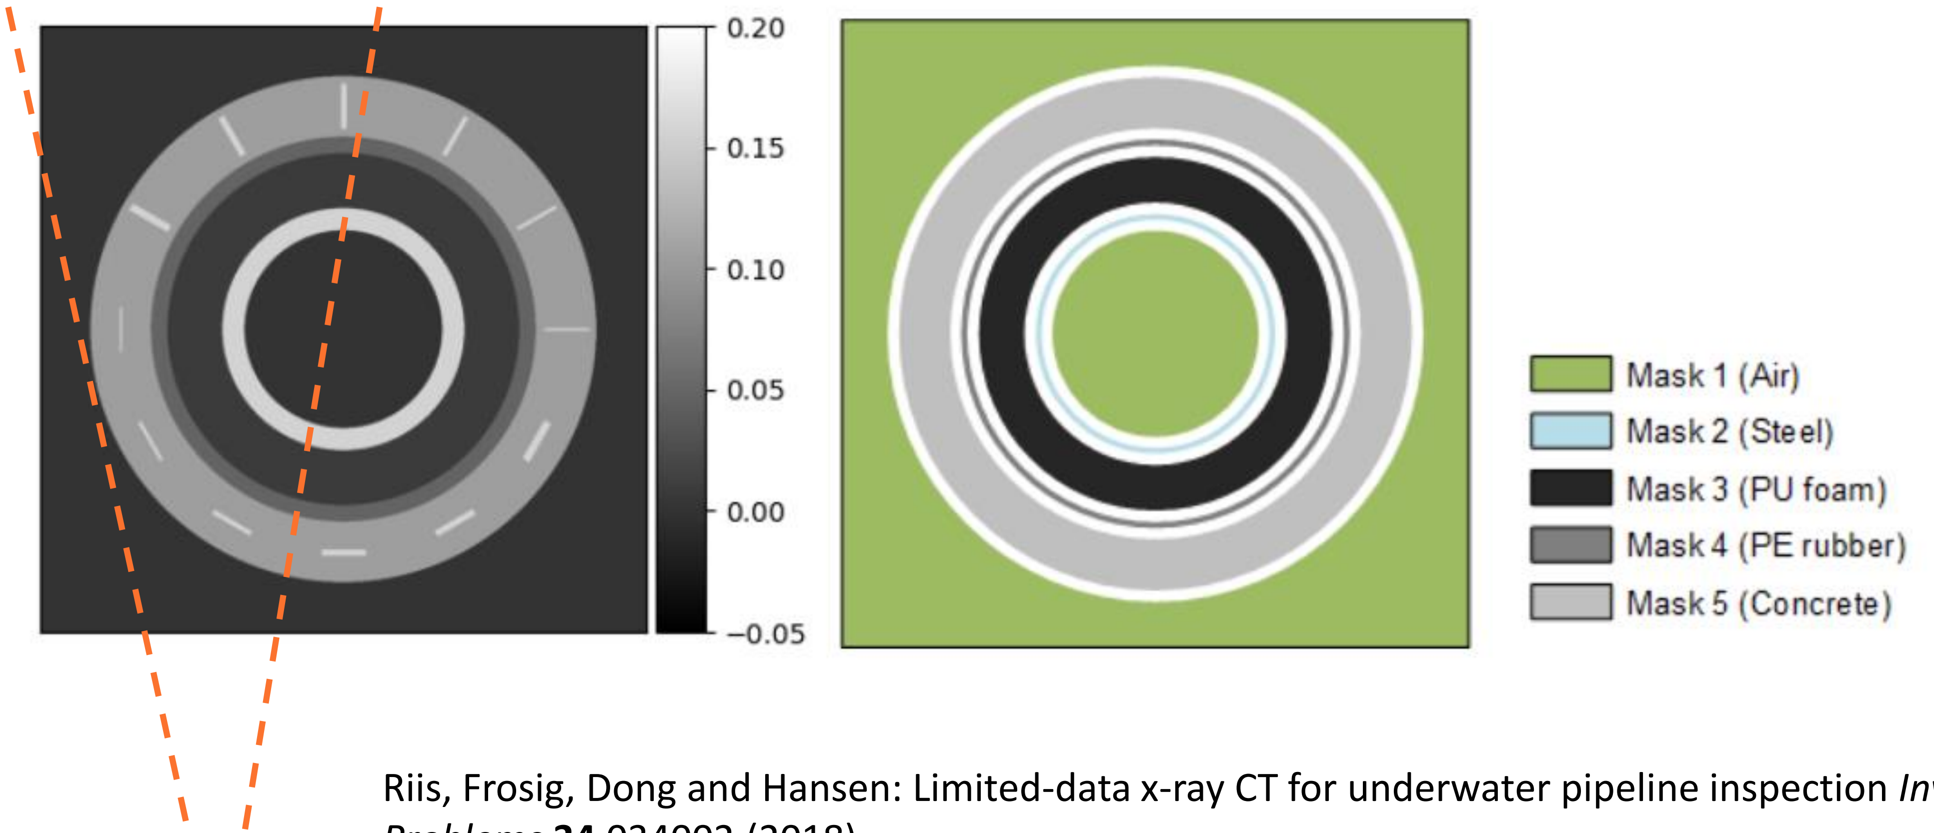

Riis, Frosig, Dong and Hansen: Limited-data x-ray CT for underwater pipeline inspection *Inverse Problems* **34** 034002 (2018)

Numerical experiments with synthetic data

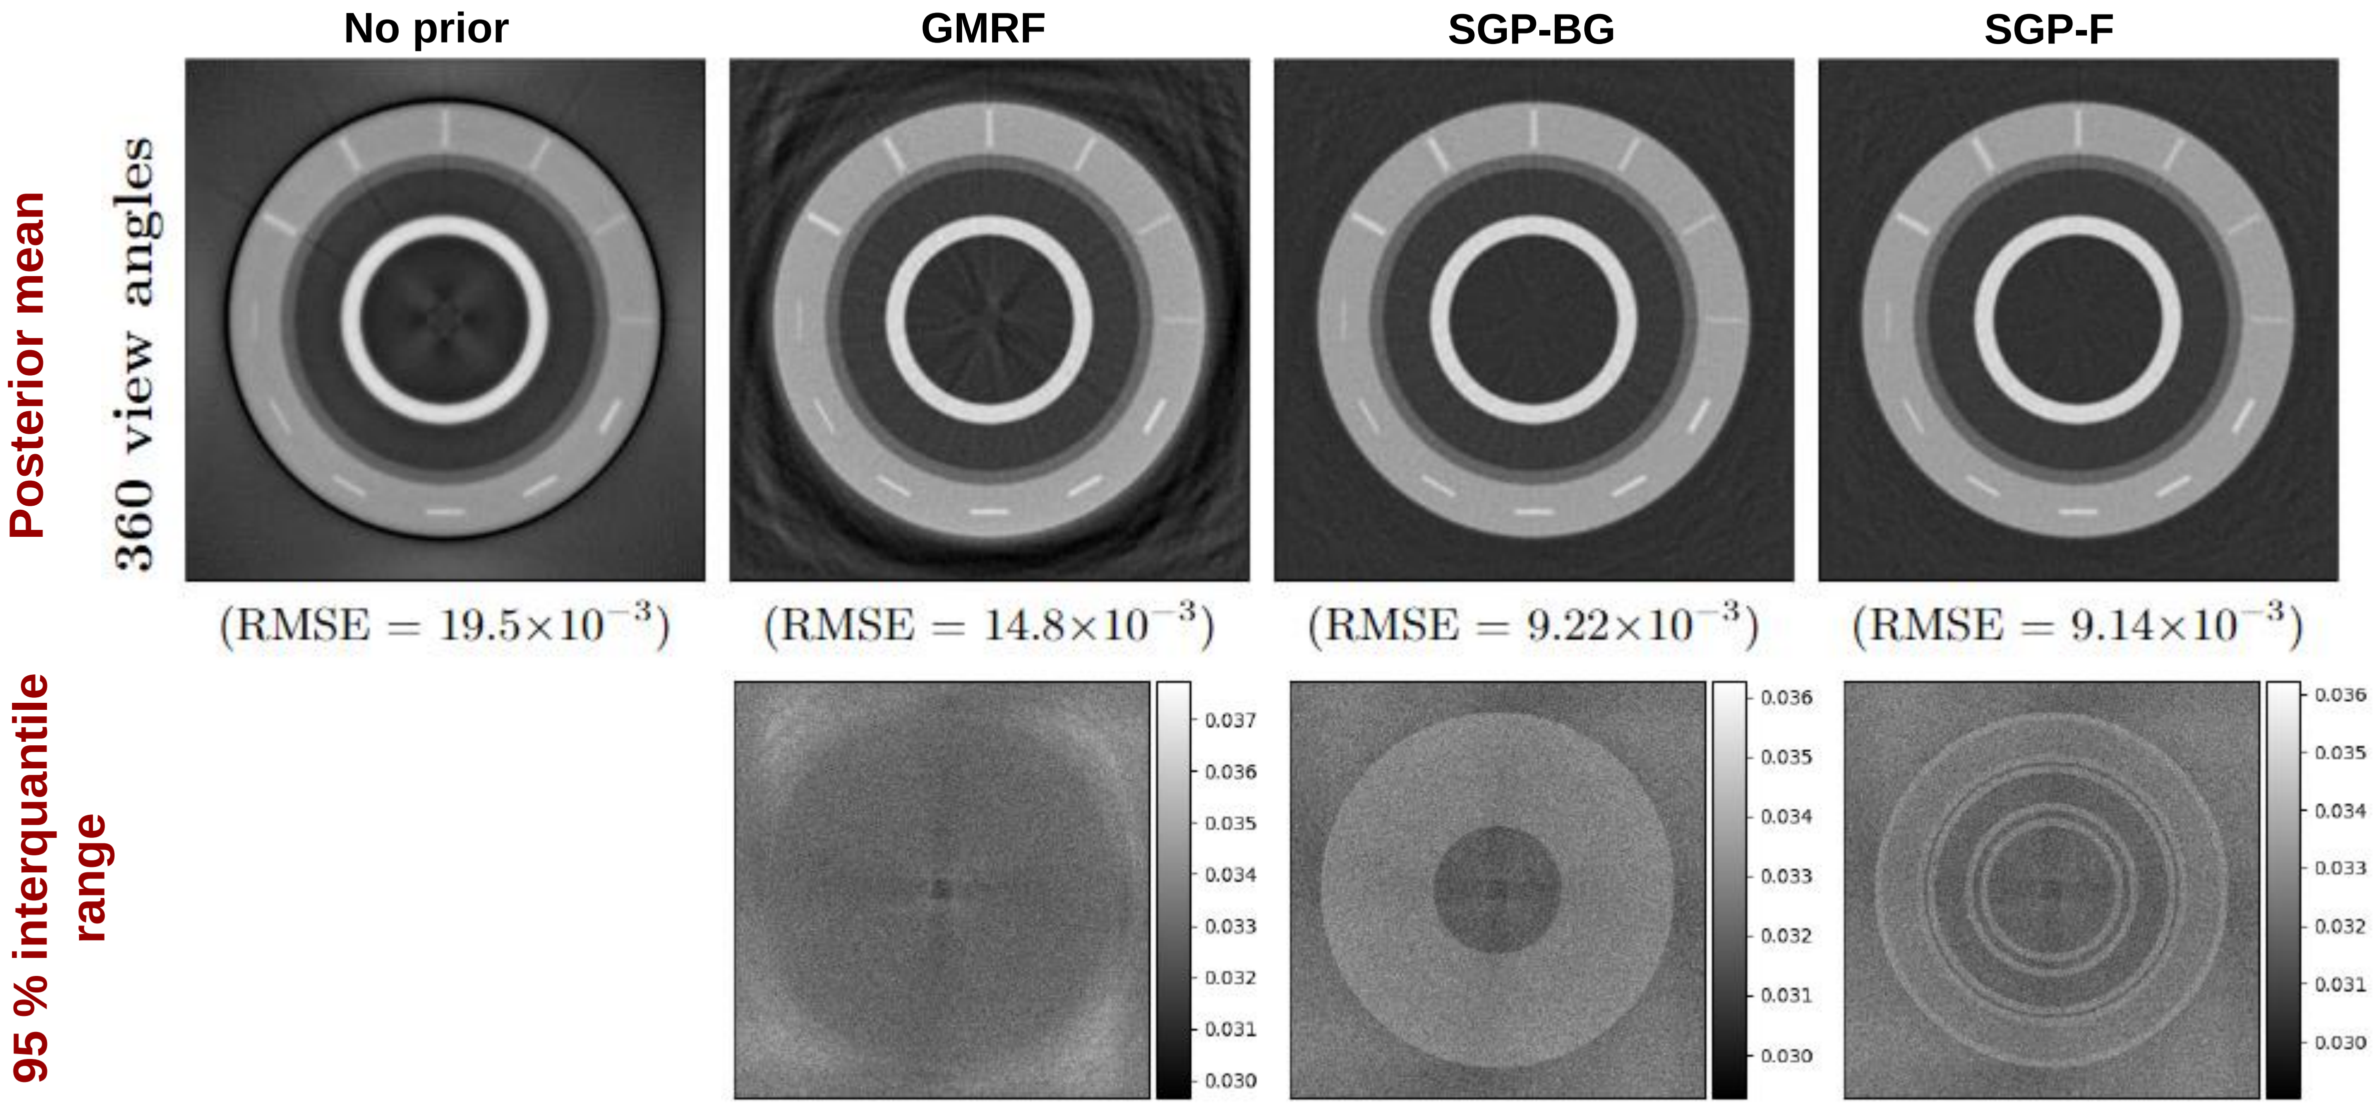

Numerical experiments with synthetic data

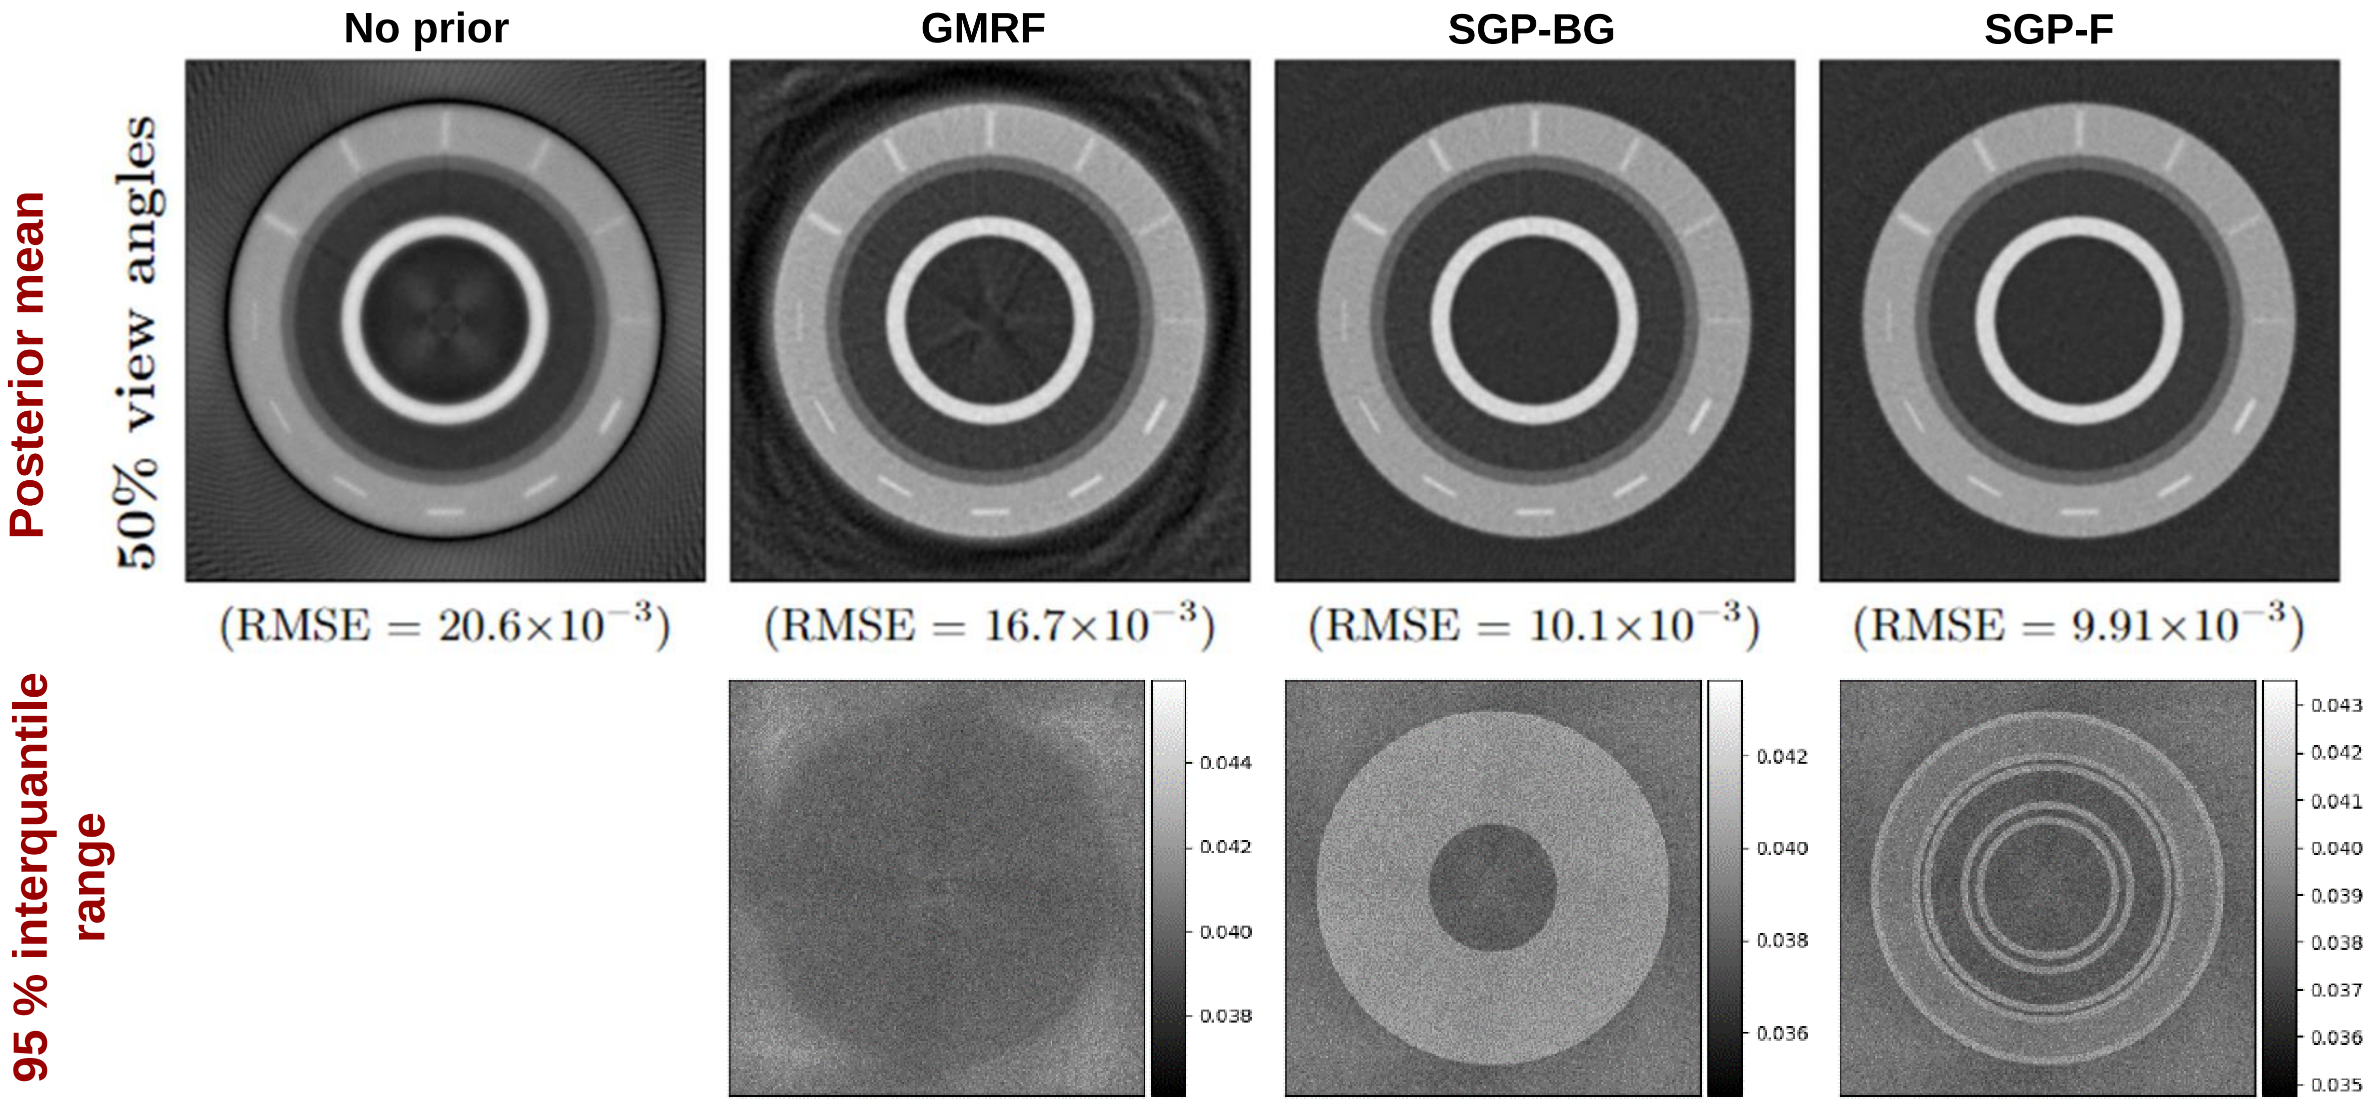

Numerical experiments with synthetic data

No prior | **GMRF** | **SGP-BG** | **SGP-F**

Posterior mean

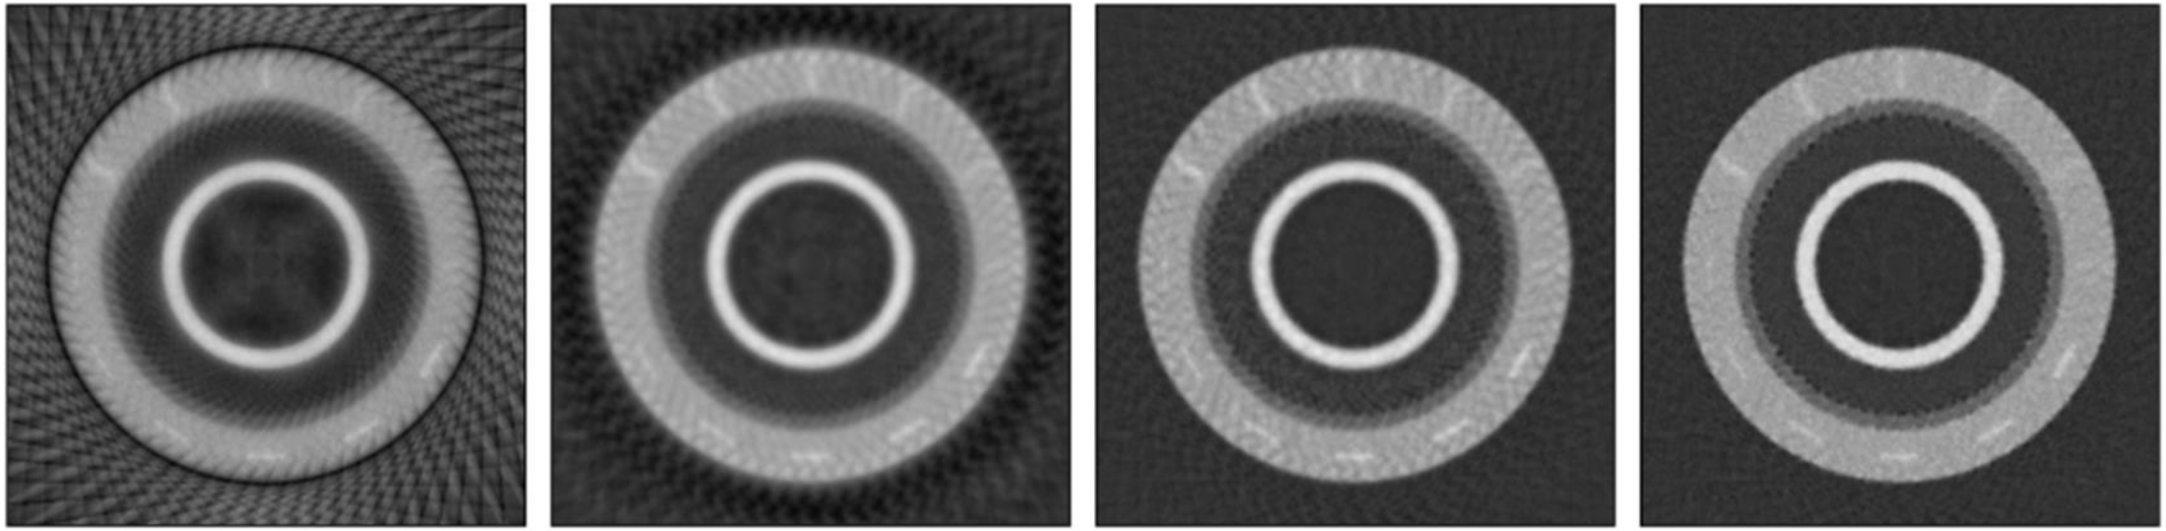

(RMSE = 25.1×10^{-3}) (RMSE = 16.5×10^{-3}) (RMSE = 12.2×10^{-3}) (RMSE = 11.6×10^{-3})

72 view angles

95 % interquantile range

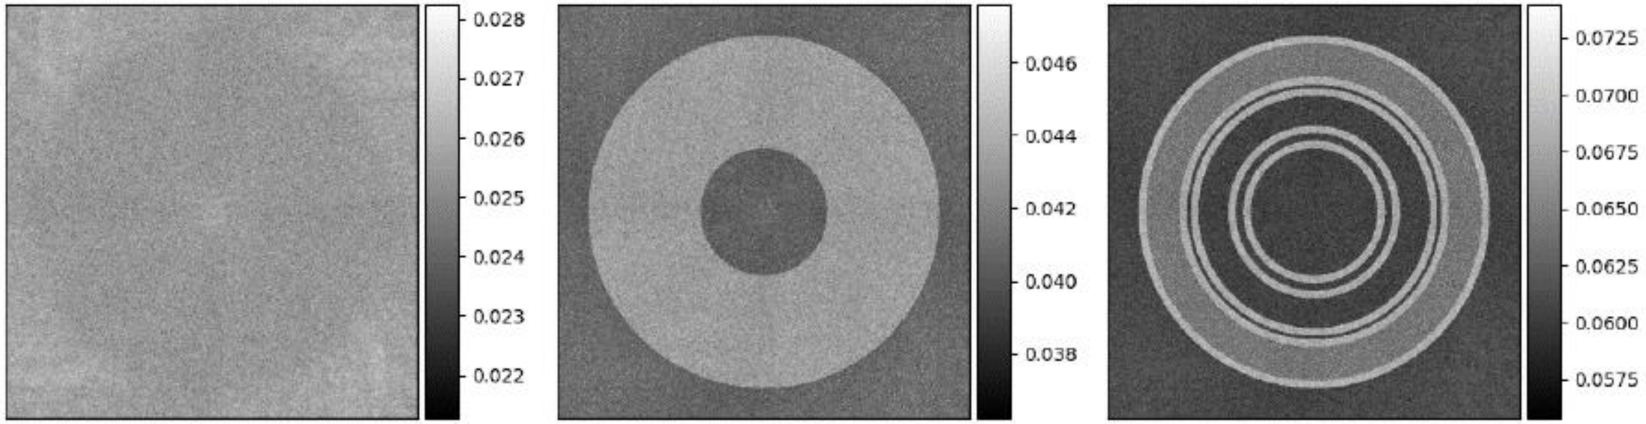

Numerical experiments with synthetic data

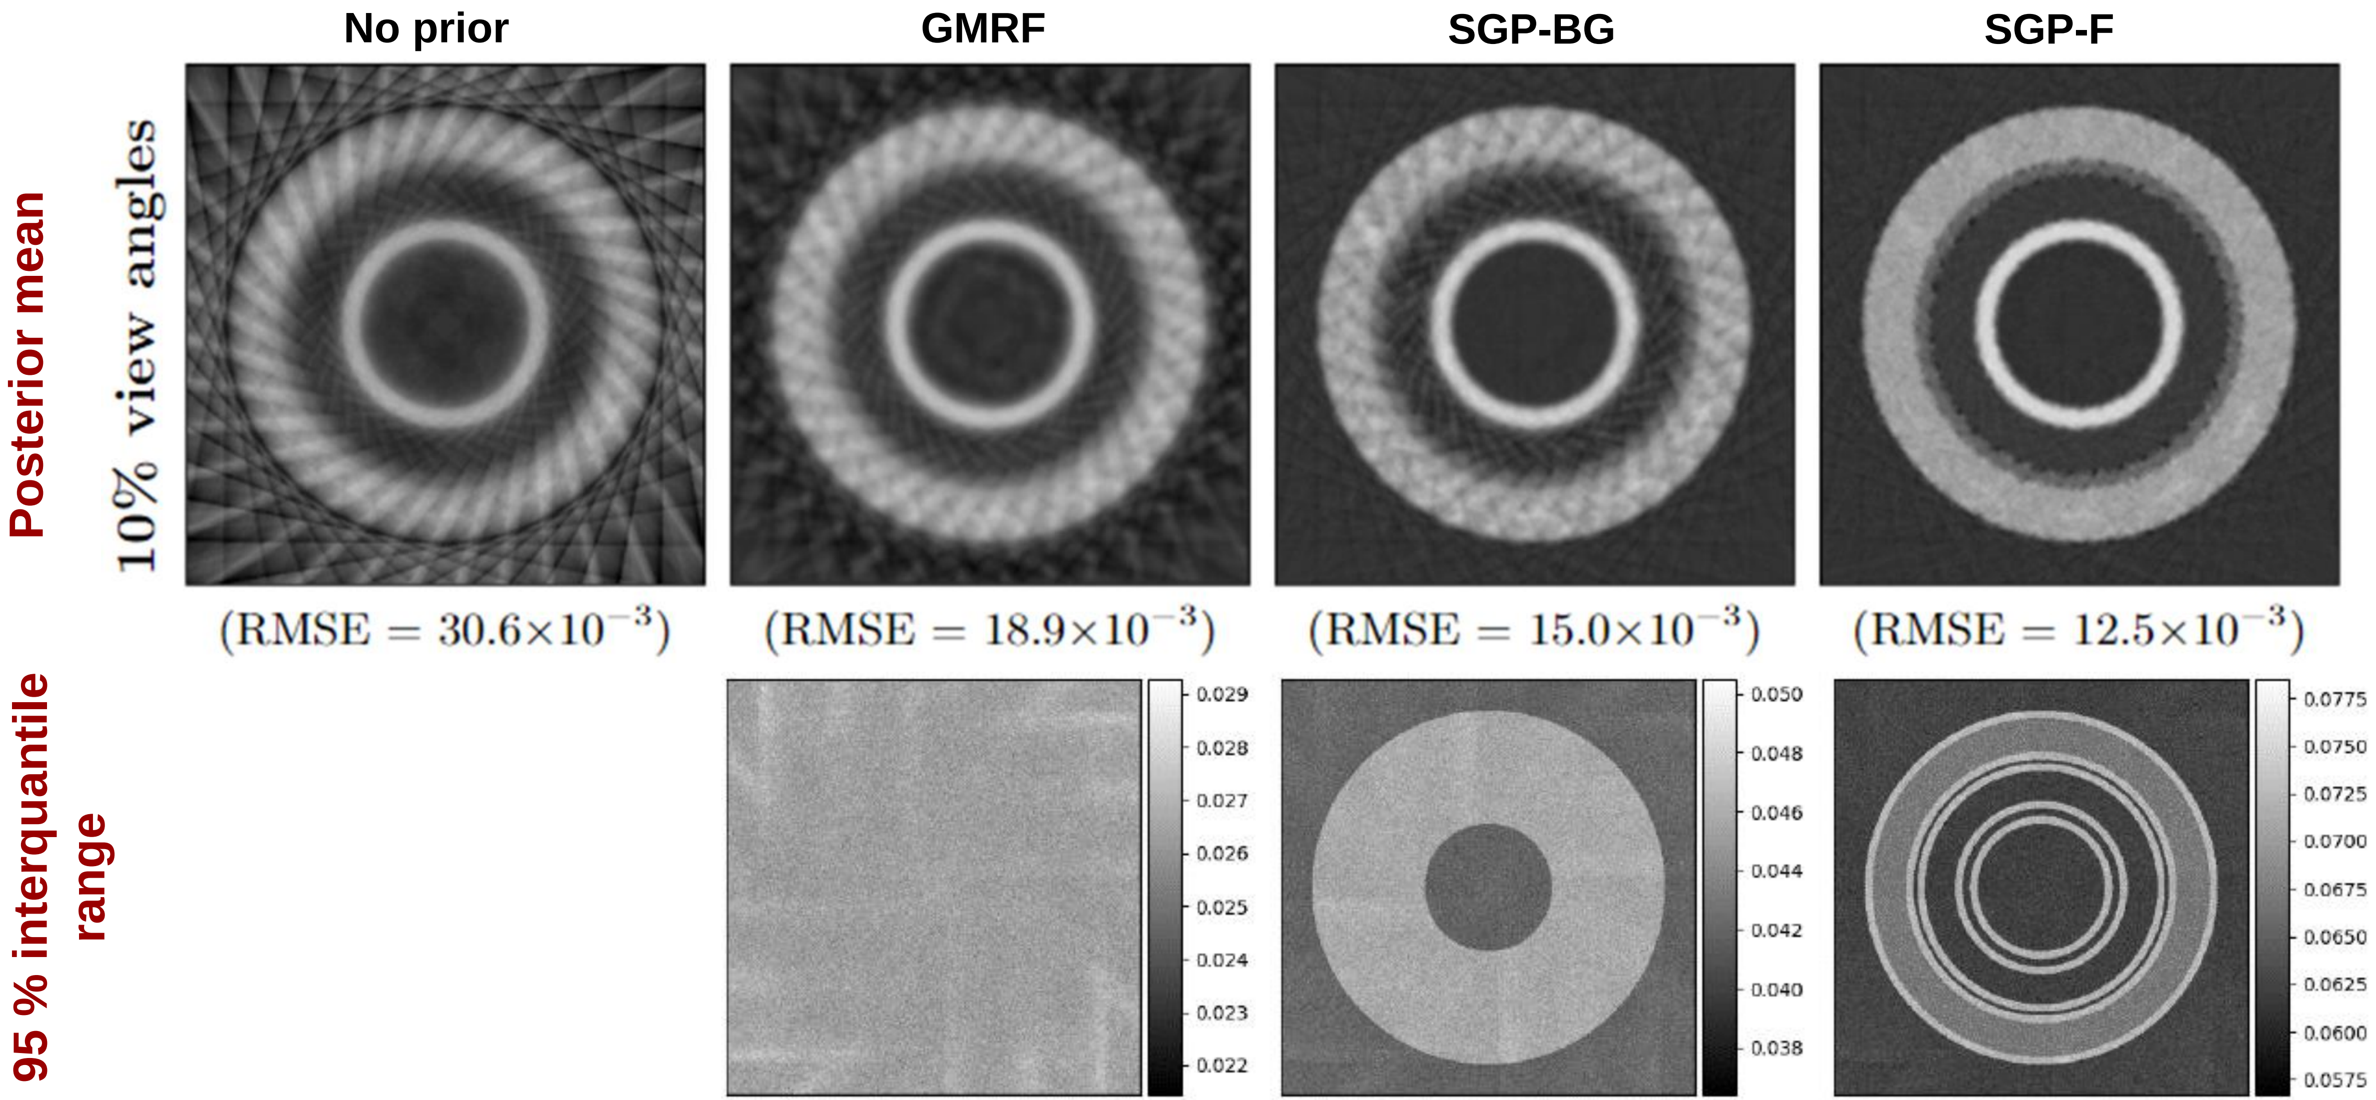

Numerical experiments with real data

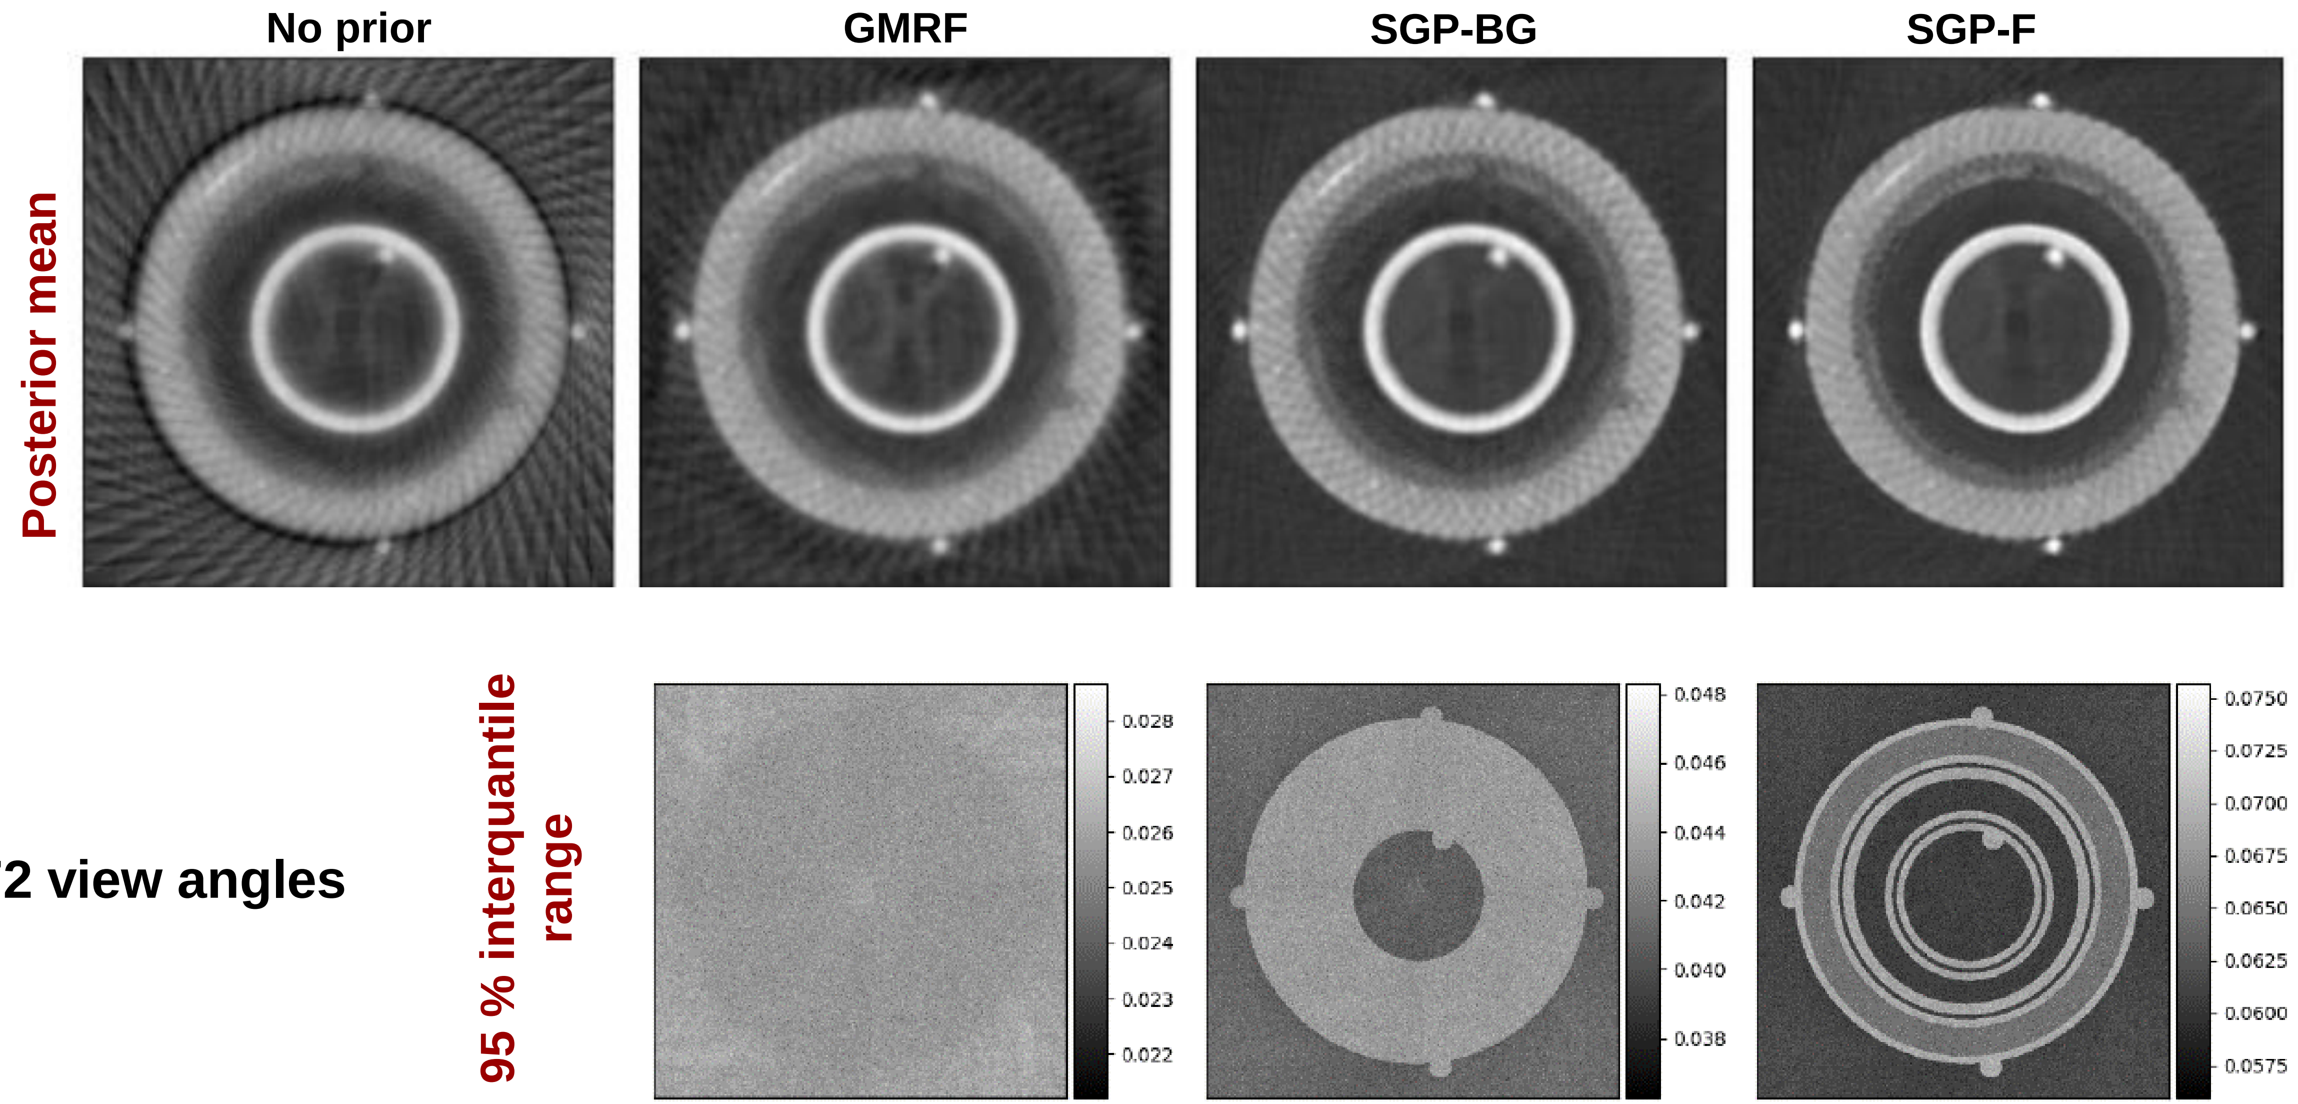

72 view angles

A new method for separate reconstruction and UQ of pipe and defects

- Separation of pipe and defect
- Defect uncertainty quantification
- Likelihood of pixels being defect (positive or negative)
- Implemented using Gibbs sampling from CUQIpy
- GPU-accelerated CT forward model using CIL

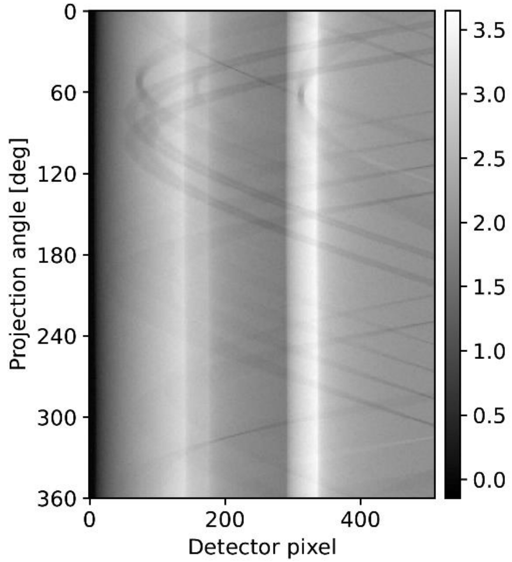

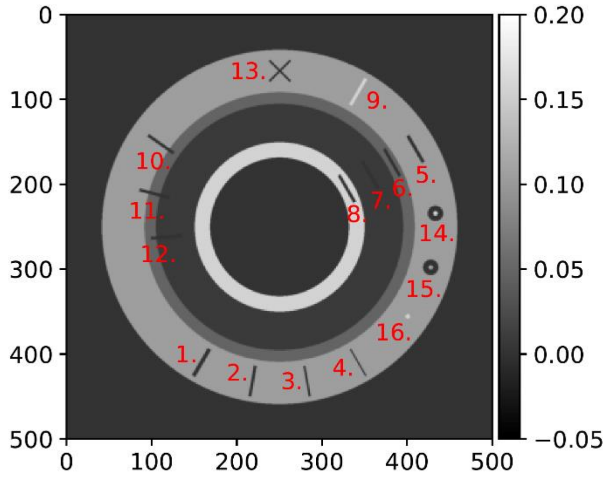

Forward model

$$\mathbf{b} = \mathbf{A}(\mathbf{z} + \mathbf{d}) = \mathbf{A}\mathbf{z} + \mathbf{A}\mathbf{d}$$

Pipe structure (z) **Small defects** (d)

Christensen, Uribe, Riis and Jørgensen: Structural Gaussian priors for Bayesian CT reconstruction of subsea pipes, AMSE, 31, 1, 2023 (doi.org/10.1080/27690911.2023.2224918)

Christensen, Riis, Pereyra and Jørgensen (in review)

Core Imaging Library for CT and other inverse problems (CIL)

```
data = ZEISSDataReader(filename).read()
data = TransmissionAbsorptionConverter()(data)
show_geometry(data.geometry)
recon = FDK(data).run()
show2D(recon)
```

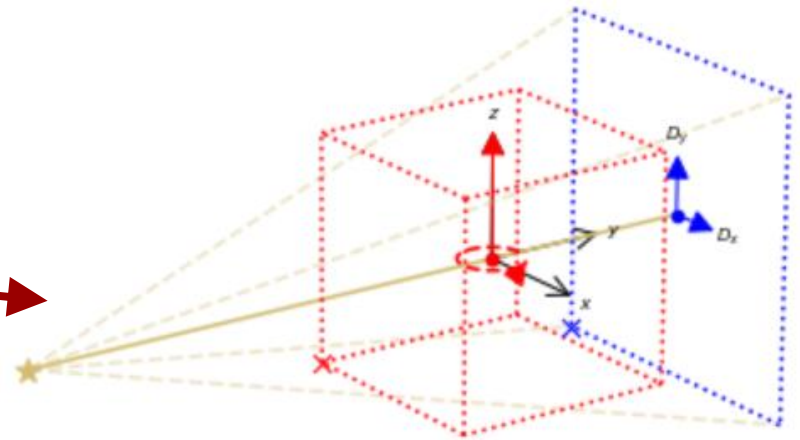

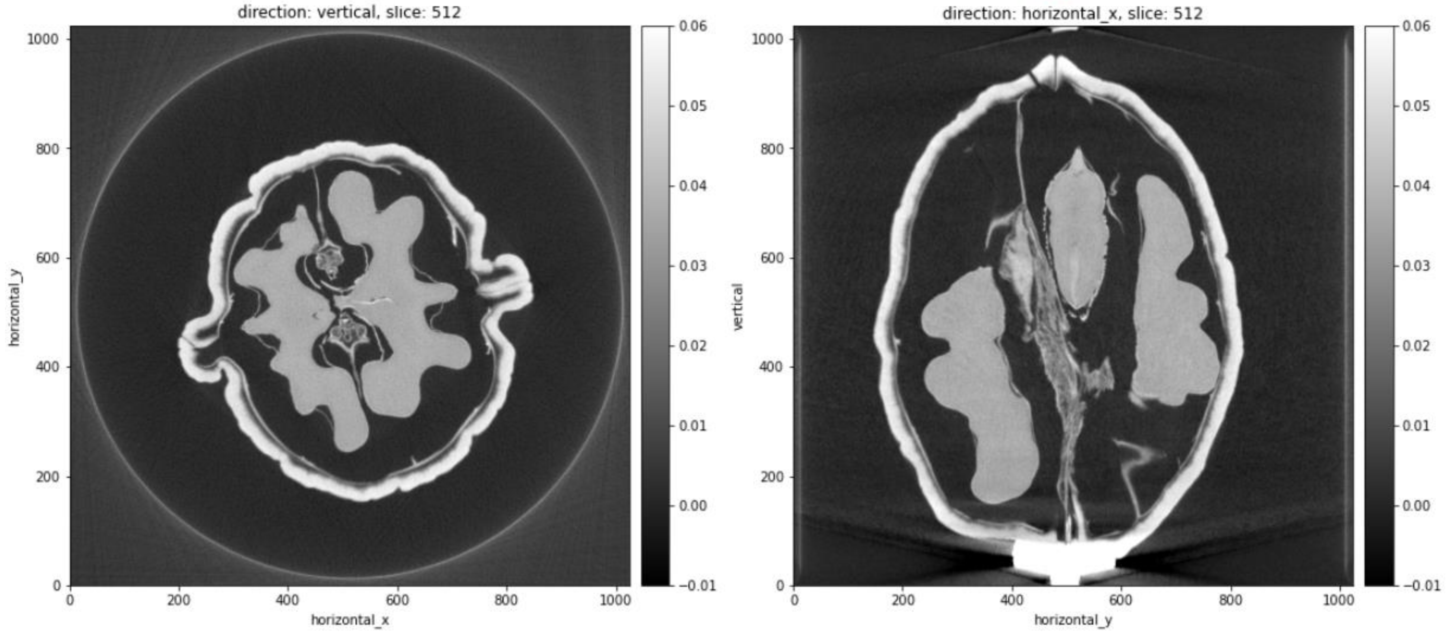

- Data readers/writers
- Pre-processing tools
- TIGRE and ASTRA backend
- 2D, 3D and 4D data
- *Near math* optimisation syntax
- Visualisation

ccpi.ac.uk/CIL

Jørgensen et al. 2021: *Core Imaging Library - Part I: a versatile Python framework for tomographic imaging*, Phil. Trans. R. Soc. A, **379**, 20200193: https://doi.org/10.1098/rsta.2020.0192

CUQIpy

Forward model & data

$\mathbf{A}$ $\mathbf{b}$

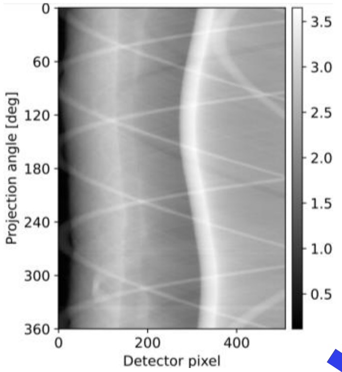

Parameter assumptions

$$\mathbf{b} \sim \mathcal{N}(\mathbf{A}\mathbf{x}, \sigma^2_{\text{noise}}\mathbf{I})$$

$$\mathbf{x} \sim \mathcal{N}(\mathbf{0}, \delta^2\mathbf{I})$$

$$\delta \sim \text{Gamma}(a, b)$$

Bayesian inverse problem modelling framework

$$p(\mathbf{x}, \delta|\mathbf{b}) \propto p(\mathbf{b}|\mathbf{x})\, p(\mathbf{x}|\delta)\, p(\delta)$$

Posterior — Likelihood — Priors

Computational UQ engine

- Efficient posterior sampling using structure
- High-level interface for non-experts
- Full control for experts
- Hierarchical problem support
- Catalogue of test problems, priors, …

https://cuqi-dtu.github.io/CUQIpy/

Parameter estimates

$\mathbf{x}$

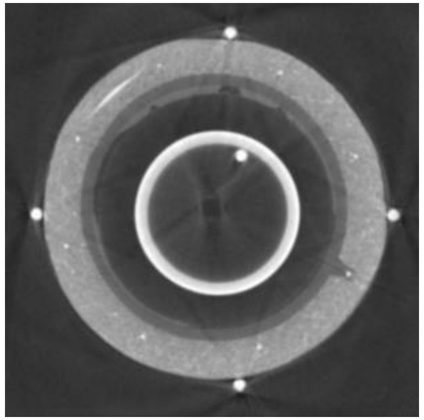

δ

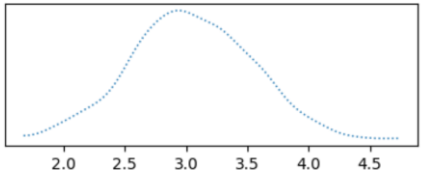

UQ information

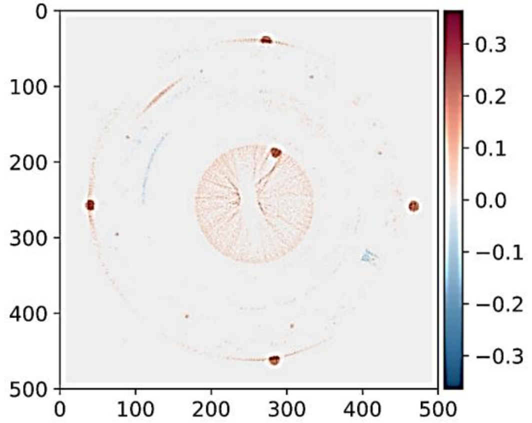

X-ray CT with the CUQIpy plugin for CIL

CIL CORE IMAGING LIBRARY

ccpi.ac.uk/CIL

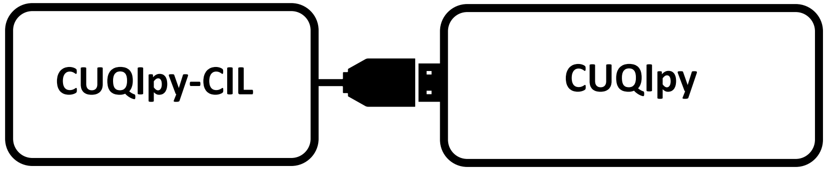

- **Set up forward model:**

```
A = FanBeam2DModel(det_count=560,
                   det_spacing=0.2,
                   angles=-np.linspace(0, 2*np.pi, 360),
                   source_object_dist=410.66,
                   object_detector_dist=143.08,
                   domain=(83.09, 83.09),
                   im_size=(500,500))
```

- **Load data as CUQIarray with Image2D geometry:**

```
y_obs = CUQIarray(sinogram, geometry=Image2D(im_shape=(360, 560)))
```

Riis et al. *CUQIpy - Part I: computational uncertainty quantification for inverse problems in Python*, https://arxiv.org/abs/2305.16949

X-ray CT with the CUQIpy plugin for CIL

- **Bayesian inverse problem:**

```
d = Gamma(1, 1e-4)
s = Gamma(1, 1e-4)
x = LMRF(1/d, geometry=A.domain_geometry)
y = Gaussian(A @ x, 1/s)
```

- **Specify posterior incl. observed data**

```
posterior = JointDistribution(d, s, x, y)(y=y_obs)
```

- **Gibbs sampling, exploiting conjugacy:**

```
sampling_strategy = {'d': ConjugateApprox,
                     's': Conjugate,
                     'x': UGLA}
samples = Gibbs(posterior, sampling_strategy).sample(500, 100)
```

X-ray CT with the CUQIpy plugin for CIL: Helsinki Tomography Challange data

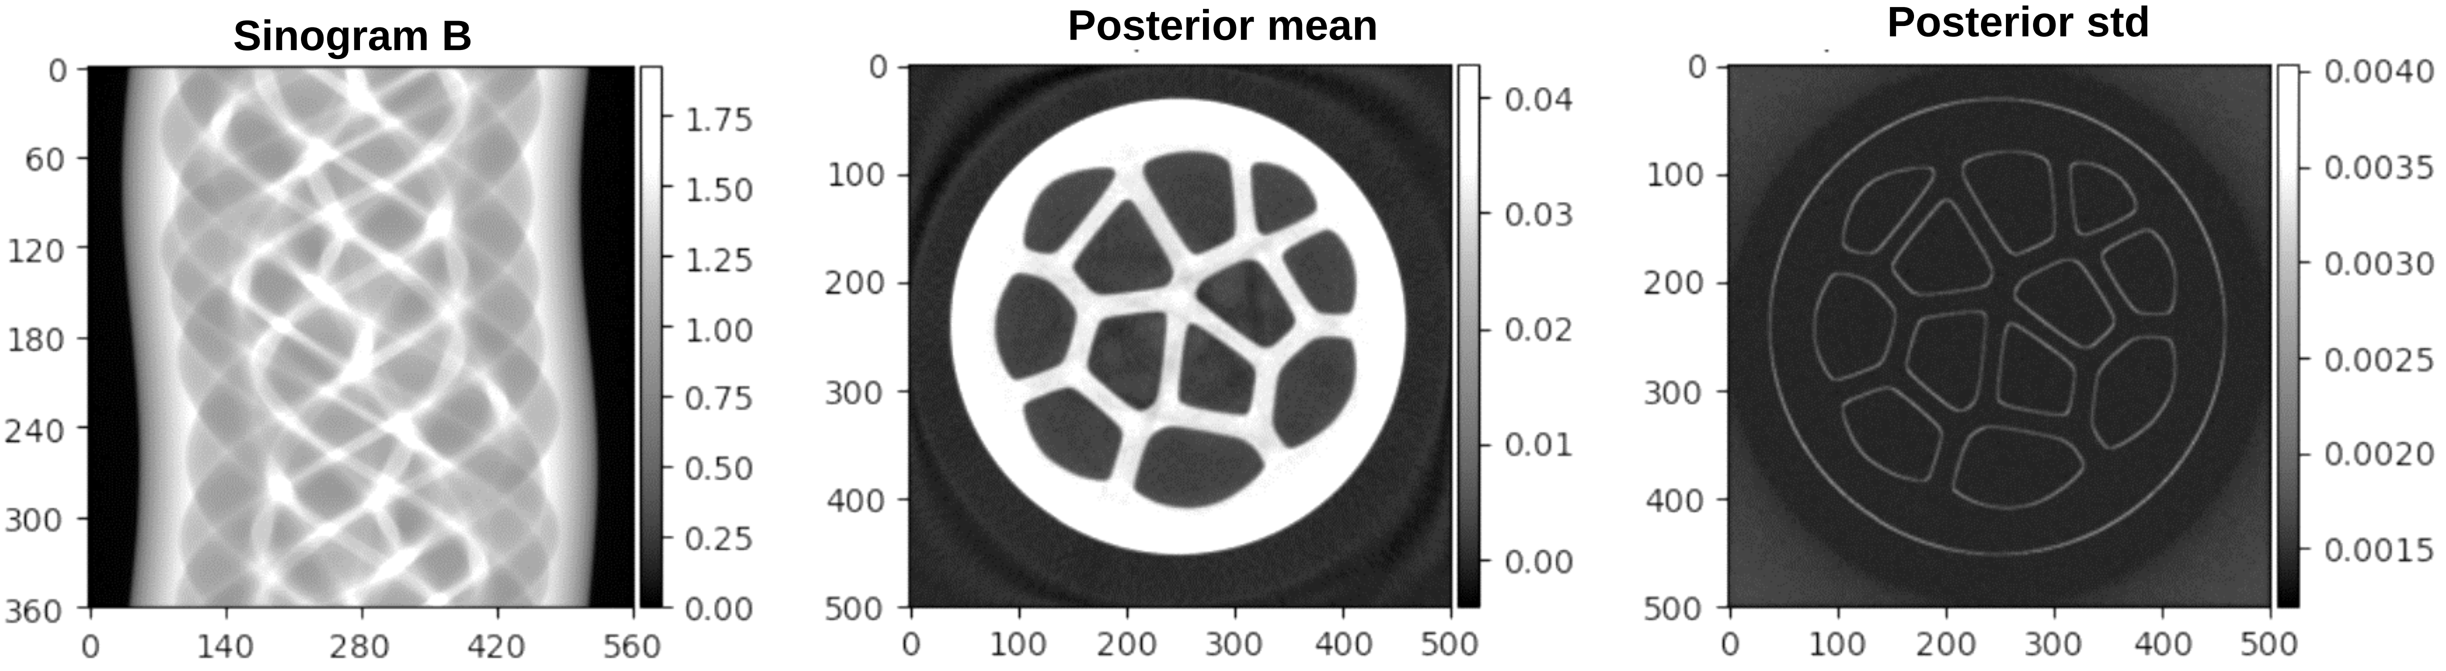

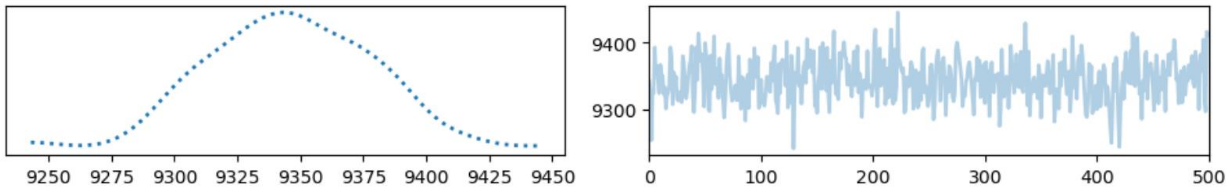

(a) Data precision s

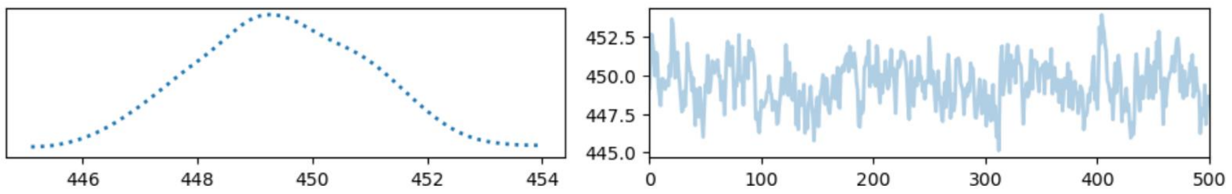

(b) Prior precision d

Helsinki Tomography Challenge 2022: https://fips.fi/HTC2022.php

Riis et al. *CUQIpy - Part I: computational uncertainty quantification for inverse problems in Python*, https://arxiv.org/abs/2305.16949

A new method for separate reconstruction and UQ of pipe and defects

- Separation of pipe and defect
- Defect uncertainty quantification
- Likelihood of pixels being defect (positive or negative)
- Implemented using Gibbs sampling from CUQIpy
- GPU-accelerated CT forward model using CIL

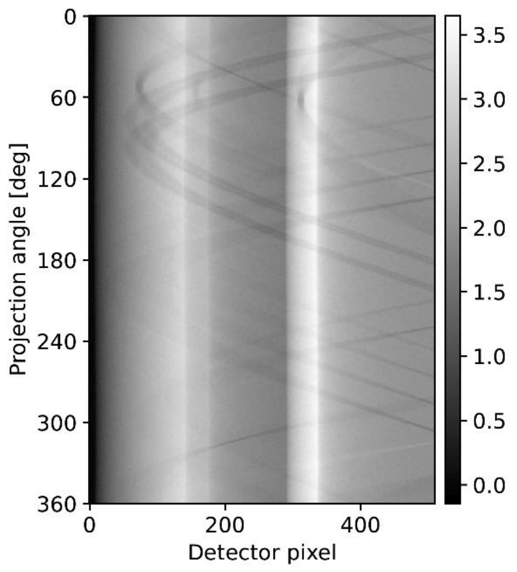

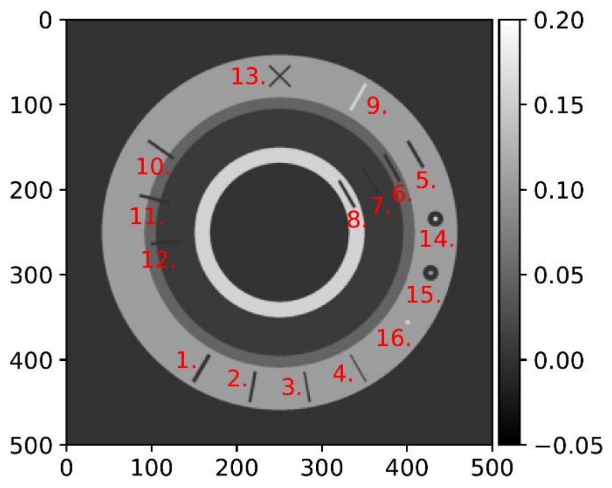

Forward model

$$\mathbf{b} = \mathbf{A}(\mathbf{z} + \mathbf{d}) = \mathbf{A}\mathbf{z} + \mathbf{A}\mathbf{d}$$

Pipe structure **Small defects**

```
A@z + A@d
```

Christensen, Uribe, Riis and Jørgensen: Structural Gaussian priors for Bayesian CT reconstruction of subsea pipes, AMSE, 31, 1, 2023 (doi.org/10.1080/27690911.2023.2224918)

Christensen, Riis, Pereyra and Jørgensen (in review)

Gamma Markov Random Field prior for defects

The gamma MRF is expressed in a hierarchical manner:

$$\tilde{d}_{i,j} \mid \tilde{s}_{i,j} \sim \mathcal{N}\left(0, \tilde{s}_{i,j}\right),$$
$$\tilde{s}_{i,j} \mid \tilde{\boldsymbol{w}} \sim \mathcal{IG}\left(\omega, \omega \tilde{g}_{i,j}(\tilde{\boldsymbol{w}})\right),$$
$$\tilde{w}_{i,j} \mid \tilde{\boldsymbol{s}} \sim \mathcal{G}\left(\omega, 1/\left(\omega \tilde{h}_{i,j}(\tilde{\boldsymbol{s}})\right)\right),$$

where $i = 1, \ldots, N$, $j = 1, \ldots, N$, ω is a scalar parameter, and

$$\tilde{g}_{i,j}(\tilde{\boldsymbol{w}}) = \left(\tilde{w}_{i,j} + \tilde{w}_{i+1,j} + \tilde{w}_{i,j+1} + \tilde{w}_{i+1,j+1}\right)/4,$$
$$\tilde{h}_{i,j}(\tilde{\boldsymbol{s}}) = \left(\tilde{s}_{i,j}^{-1} + \tilde{s}_{i-1,j}^{-1} + \tilde{s}_{i,j-1}^{-1} + \tilde{s}_{i-1,j-1}^{-1}\right)/4.$$

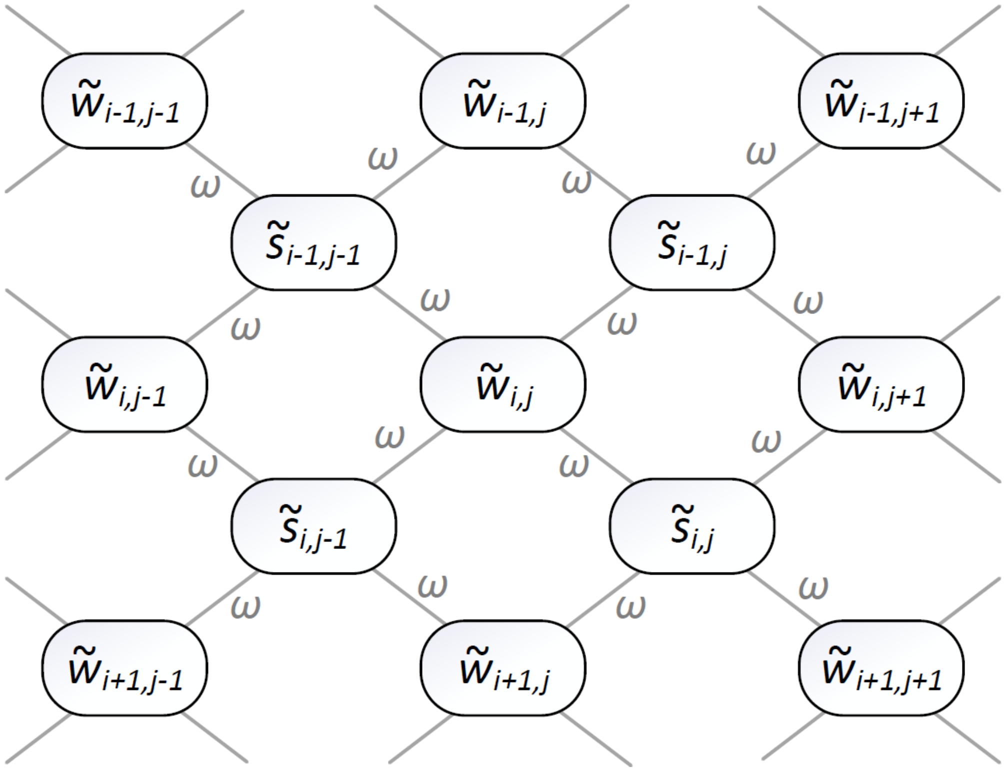

Dikmen and Cemgil. Gamma markov random fields for audio source modeling. IEEE Transactions on Audio, Speech and Language Processing, 18(3):589–601, 2010.
Altmann, Pereyra, and McLaughlin. Bayesian nonlinear hyperspectral unmixing with spatial residual component analysis. IEEE Transactions on Computational Imaging, 1(3):174–185, 2015.

Defect Detection in Subsea Pipes Bayesian model in CUQIpy

Likelihood

$$\mathbf{b} \mid \mathbf{z}, \mathbf{d} \sim \mathcal{N}(\mathbf{A}\mathbf{z} + \mathbf{A}\mathbf{d}, \lambda^{-1}\mathbf{I})$$

```
b = Gaussian(A@z + A@d, 1/lambda_noise)
```

Pipe Prior z: Structural Gaussian

$$\mathbf{z} \sim \mathcal{N}\left(\boldsymbol{\mu}_{SGP}, \left(\mathbf{R}_{SGP}^T \mathbf{R}_{SGP}\right)^{-1}\right)$$

```
z = JointGaussianSqrtPrec(mean=[mu1, mu2, …],
    prec=[R0, R1, …], … )
```

Defect Prior d: Sparsity + spatially coherent Gamma MRF

$$\mathbf{d} \mid \mathbf{s} \sim \mathcal{N}\left(0, f_d(\mathbf{s})\right)$$

$$\mathbf{s} \mid \mathbf{w} \sim \mathcal{IG}\left(\omega_0, f_s(\mathbf{w})\right)$$

$$\mathbf{w} \mid \mathbf{s} \sim \mathcal{G}\left(\omega_0, f_w(\mathbf{s})\right)$$

```
d = Gaussian(0, f_d(s))
s = InverseGamma(w0, f_s(w))
w = Gamma(w0, f_w(s))
```

Posterior

$$p(\mathbf{z}, \mathbf{d}, \mathbf{s}, \mathbf{w} \mid \mathbf{b}) \propto p(\mathbf{b} \mid \mathbf{z}, \mathbf{d}) p(\mathbf{z}) p(\mathbf{d} \mid \mathbf{s}) p(\mathbf{s}, \mathbf{w})$$

```
post = JointDistribution(b, z, d, s, w)(b=b_data)
```

Christensen, Uribe, Riis and Jørgensen: Structural Gaussian priors for Bayesian CT reconstruction of subsea pipes, AMSE, 31, 1, 2023 (doi.org/10.1080/27690911.2023.2224918)

Christensen, Riis, Pereyra and Jørgensen (in review)

Defect Detection in Subsea Pipes
Reconstruction simulated phantom

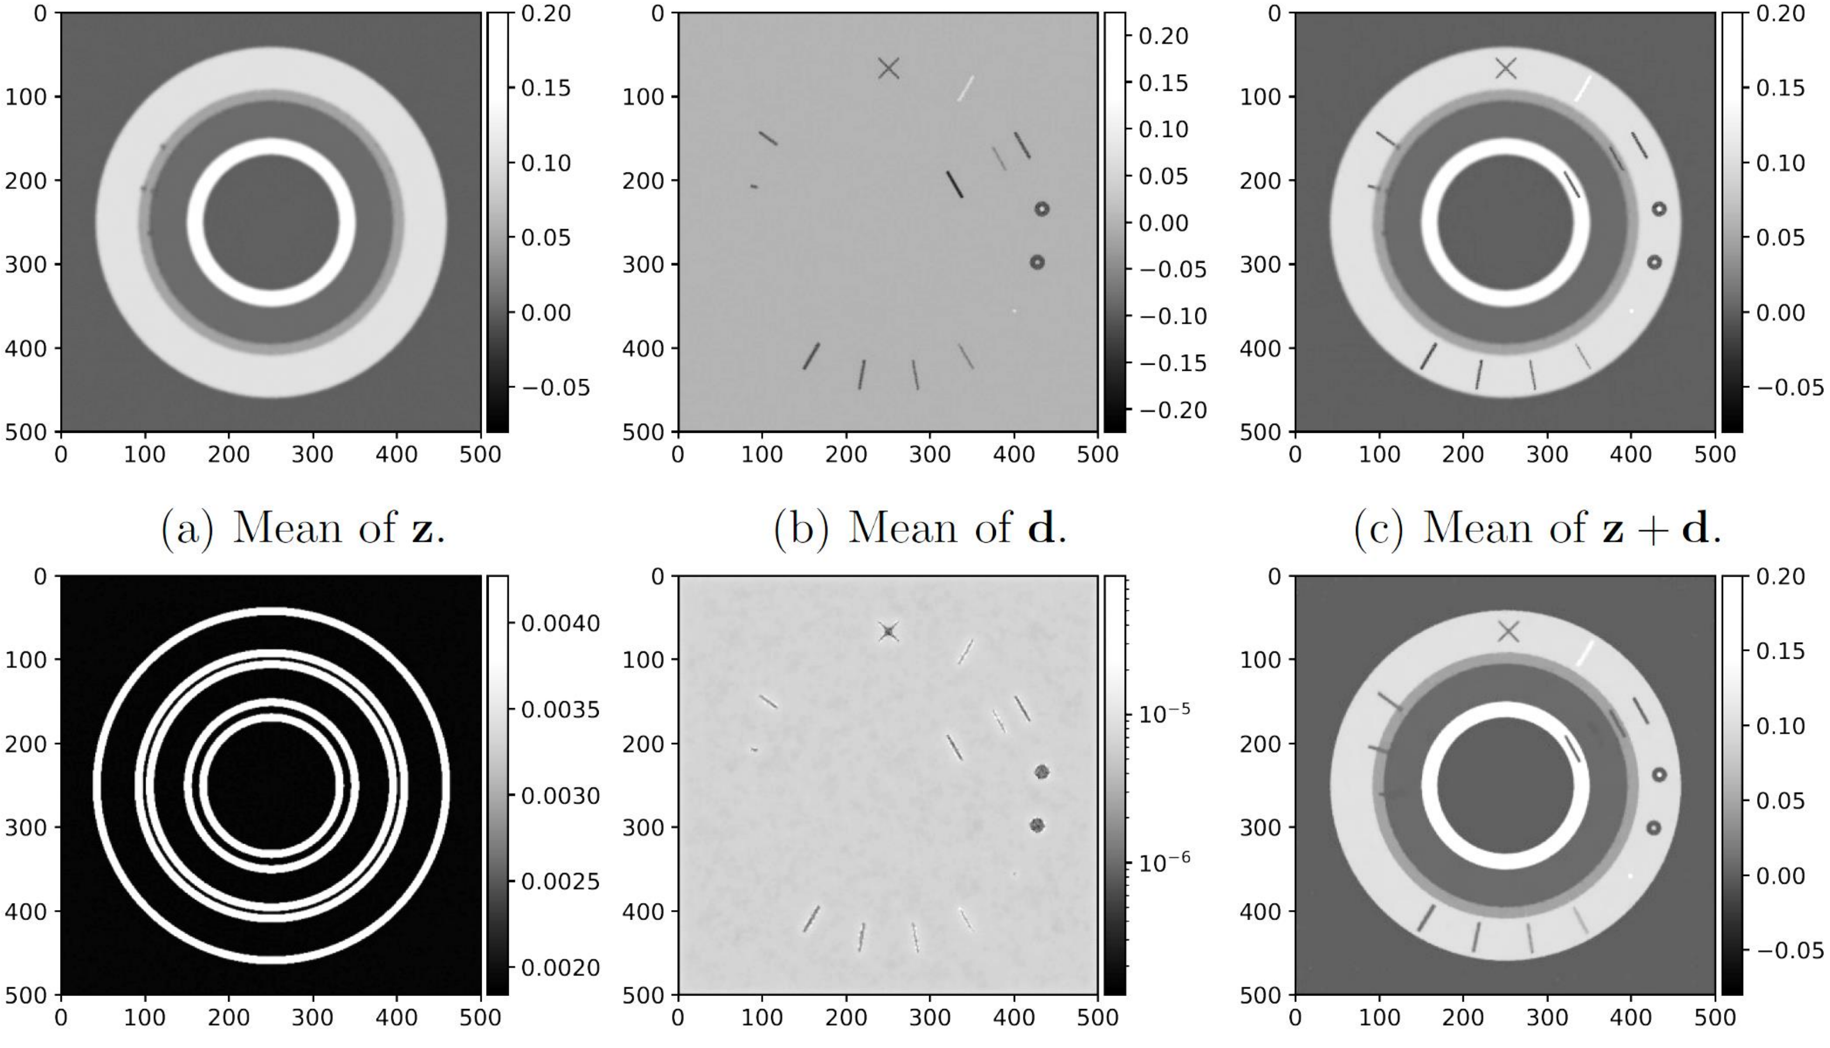

(a) Mean of $\mathbf{z}$.

(b) Mean of $\mathbf{d}$.

(c) Mean of $\mathbf{z} + \mathbf{d}$.

(d) Std of $\mathbf{z}$.

(e) Std of $\mathbf{d}$.

(f) TV reconstruction.

Christensen, Riis, Pereyra and Jørgensen (in review)

Real data results

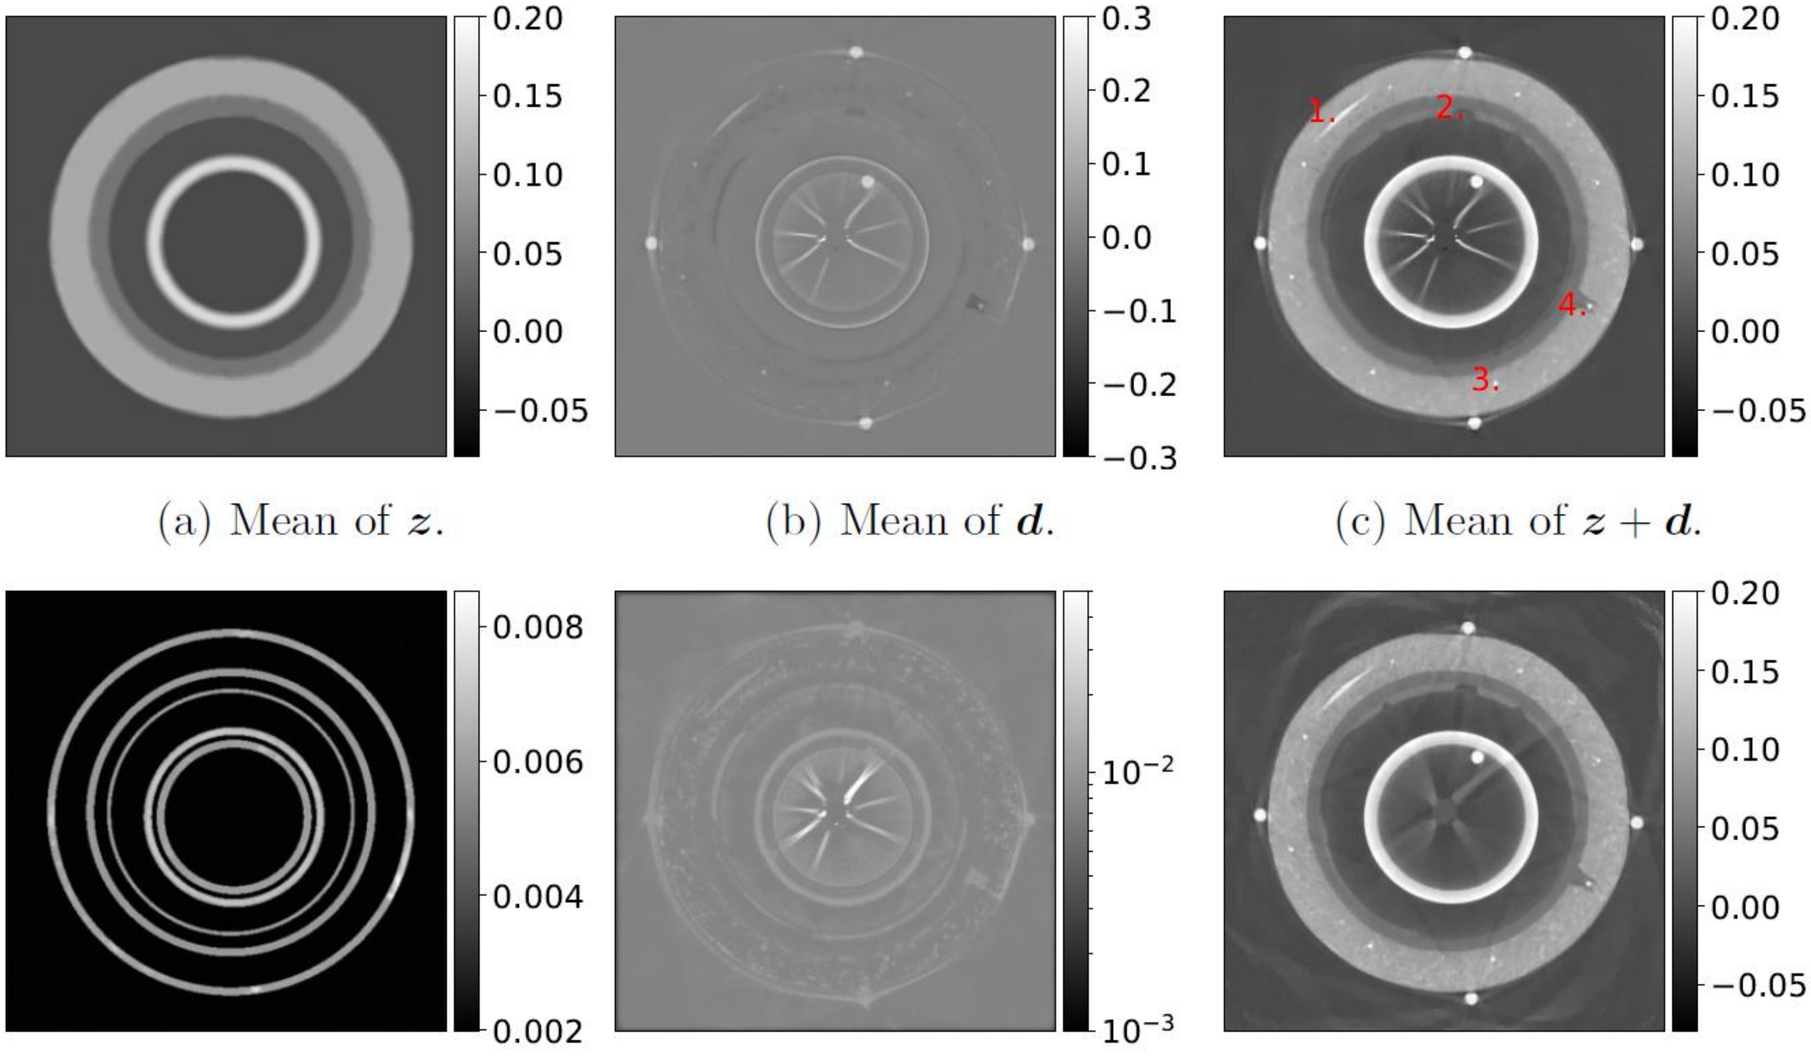

(a) Mean of $\boldsymbol{z}$.

(b) Mean of $\boldsymbol{d}$.

(c) Mean of $\boldsymbol{z} + \boldsymbol{d}$.

(d) Std of $\boldsymbol{z}$.

(e) Std of $\boldsymbol{d}$.

(f) TV reconstruction.

Christensen, Riis, Pereyra and Jørgensen (in review)

Upcoming events

CIL Training and User Meeting – 13-15 Nov 2023, Harwell, UK

CORE IMAGING LIBRARY

- Software training
- First ever CIL user meeting: Share experiences with other CIL users and developers
- Tours of Diamond Light Source synchrotron, ISIS Neutron and Muon Facility (tbc)
- Support sessions with CIL developers

CUQIpy Training

CUQIpy training @ IUQ workshop, Denmark, Sept '22

"Well documented and easy to use."

"I think the whole user-experience was very smooth […]"

"It's obvious that it is aimed towards non-experts, but it's also great that experts can really take advantage of the package and do more complex stuff."

New-ish book on CT reconstruction

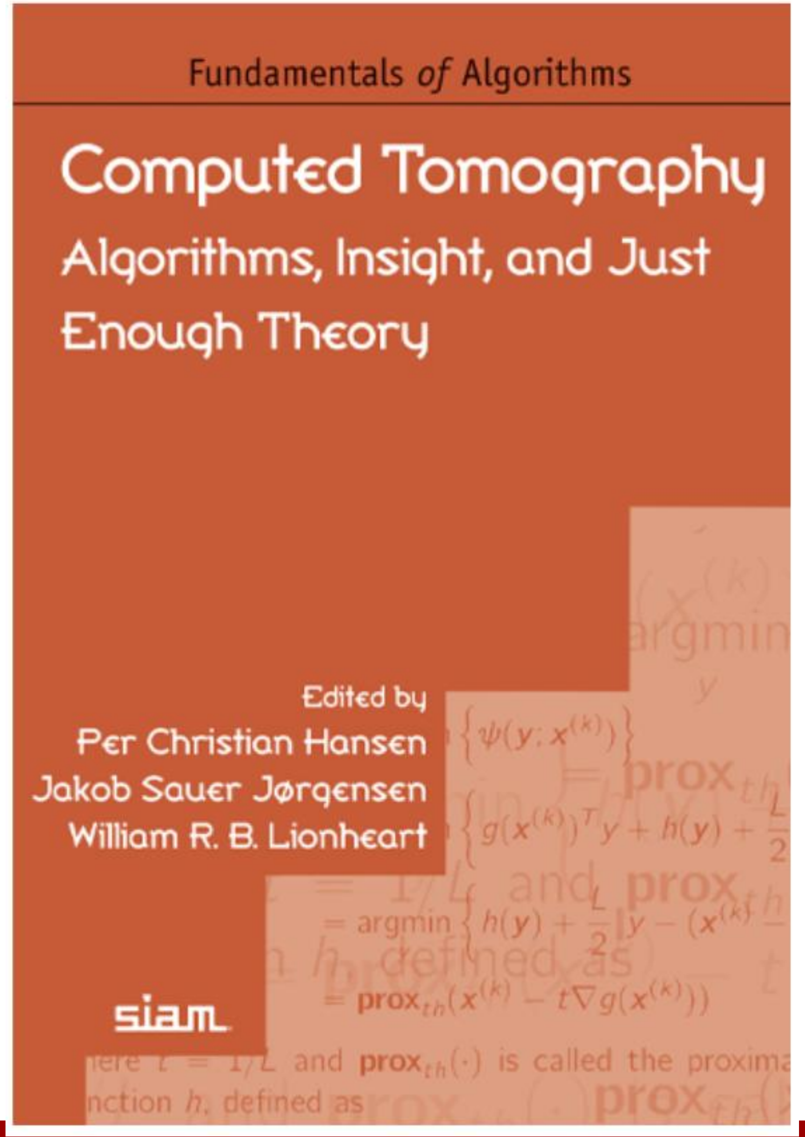

Computed Tomography: Algorithms, Insight, and Just Enough Theory

Edited by Per Christian Hansen, Jakob Sauer Jørgensen, and William R. B. Lionheart
Published: 2021
Pages: xviii + 337 pages
Softcover
ISBN: 978-1-611976-66-3
Order Code: FA18

bookstore.siam.org/fa18

Conclusions

Christensen, Uribe, Riis and Jørgensen: Structural Gaussian priors for Bayesian CT reconstruction of subsea pipes, AMSE, 31, 1, 2023 (doi.org/10.1080/27690911.2023.2224918)

- Structural Gaussian prior use well-known pipe geometry and materials to promote expected linear attenuation coefficients.
  - Reduced streak artifacts and improved layer contrast.
  - SGP-BG gives good result using only knowledge of exterior dimension.
  - SGP-F gives the best results but require knowledge of internal structure.
  - Uncertainty estimates reflect the structure of the likelihood and/or prior.

- Splitting method allows separate Bayesian estimation of pipe and defects

- Open-source Python packages for regularized and Bayesian reconstruction:

Core Imaging Library: https://ccpi.ac.uk/cil/

CUQIpy: https://cuqi-dtu.github.io/CUQIpy/