



## A Bayesian approach to CT reconstruction with uncertain geometry

**Pedersen, Frederik Hagsholm; Jørgensen, Jakob Sauer; Andersen, Martin Skovgaard**

*Published in:*  
Applied Mathematics in Science and Engineering

*Link to article, DOI:*  
[10.1080/27690911.2023.2166041](https://doi.org/10.1080/27690911.2023.2166041)

*Publication date:*  
2023

*Document Version*  
Publisher's PDF, also known as Version of record

[Link back to DTU Orbit](#)

*Citation (APA):*  
Pedersen, F. H., Jørgensen, J. S., & Andersen, M. S. (2023). A Bayesian approach to CT reconstruction with uncertain geometry. *Applied Mathematics in Science and Engineering*, 31(1), Article 216604.  
<https://doi.org/10.1080/27690911.2023.2166041>

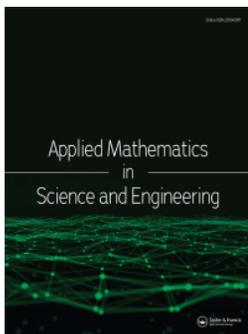
---

### General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.



## A Bayesian approach to CT reconstruction with uncertain geometry

Frederik Hagsholm Pedersen, Jakob Sauer Jørgensen & Martin Skovgaard Andersen

To cite this article: Frederik Hagsholm Pedersen, Jakob Sauer Jørgensen & Martin Skovgaard Andersen (2023) A Bayesian approach to CT reconstruction with uncertain geometry, Applied Mathematics in Science and Engineering, 31:1, 2166041, DOI: [10.1080/27690911.2023.2166041](https://doi.org/10.1080/27690911.2023.2166041)

To link to this article: <https://doi.org/10.1080/27690911.2023.2166041>



© 2023 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group



Published online: 21 Jan 2023.



Submit your article to this journal [↗](#)



Article views: 178



View related articles [↗](#)



View Crossmark data [↗](#)

# A Bayesian approach to CT reconstruction with uncertain geometry

Frederik Hagsholm Pedersen , Jakob Sauer Jørgensen  and Martin Skovgaard Andersen 

Department of Applied Mathematics and Computer Science, Technical University of Denmark, Richard Petersens Plads, Building, 324, Kongens Lyngby, Denmark

## ABSTRACT

Computed tomography (CT) is a method for synthesizing volumetric or cross-sectional images of an object from a collection of projections. Popular reconstruction methods for CT are based on the assumption that the projection geometry, which describes the relative location and orientation of the radiation source, object, and detector for each projection, is exactly known. However, in practice, these geometric parameters are estimated quantities with uncertainty. A failure to accurately estimate the geometry may lead to reconstructions with severe misalignment artifacts that significantly decrease their scientific or diagnostic value. We propose a novel reconstruction method that jointly estimates the reconstruction and the center-of-rotation offset for fan-beam tomography. The reconstruction method is based on a Bayesian approach that yields a point estimate for the reconstruction and center-of-rotation offset and, in addition, provides valuable information regarding their uncertainty. This is achieved by approximately sampling from the joint posterior distribution of the reconstruction and center-of-rotation offset using a hierarchical Gibbs sampler. Our methodology is highly flexible and can be adapted to other geometric parameters and/or scan modes. Through numerical experiments based on real tomographic data, we compare the proposed Bayesian method to two existing approaches to the problem of correcting the geometry and demonstrate that our method achieves comparable or better results under challenging conditions.

## ARTICLE HISTORY

Received 20 June 2022  
Accepted 4 January 2023

## KEYWORDS

Computed tomography; model errors; reconstruction methods; parameter estimation; uncertainty quantification

## MSC

62F15; 60J22; 65F22; 15A29; 65C20

## 1. Introduction

Computed tomography (CT) is a computational imaging technique for synthesizing volumetric or cross-sectional images of an object from a collection of projections. The projections are obtained by illuminating the object with radiation at different orientations and measuring the attenuation of the radiation due to absorption. The collection of all projections is known as the sinogram, and the process of reconstructing an image of the interior absorption of an object from the sinogram is generally an ill-posed linear inverse problem. Numerous reconstruction methods for CT have been developed, see e.g. [1] for

**CONTACT** Frederik Hagsholm Pedersen  fhape@dtu.dk  Department of Applied Mathematics and Computer Science, Technical University of Denmark, Richard Petersens Plads, Building, 324, Kongens Lyngby 2800, Denmark

© 2023 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group

This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited.

an overview including projection methods, algebraic methods, and variational methods. Computed tomography is used in many different scientific fields, the most well-known being medical imaging, where it provides a non-invasive method for inspecting the interior of patients [2]. It is also popular in geoscience, where it is an invaluable tool to analyze microstructure of materials [3,4].

Experimental CT setups are comprised of a radiation source, an object to be scanned, and a detector measuring the radiation after attenuation by the object. The projection geometry is the relative position of the source, object, and detector. Popular reconstruction methods rely on the assumption that the projection geometry is known. However, in practice, the projection geometry is estimated and therefore uncertain. Mismatches between the actual projection geometry from the data acquisition, and the estimated projection geometry used in the reconstruction model may lead to reconstructions with severe misalignment artifacts that degrade their quality, and hence also their scientific or diagnostic value.

To mitigate misalignment artifacts, correction of the projection geometry for CT is essential. Many different methods have been proposed to address various errors in the projection geometry. One class of methods is based around doing a calibration scan with fiducial markers, and then correcting the geometry based on the acquired scan [5]. However, taking additional measurements may be impossible or very costly depending on the specific application or experimental setup. For this reason, it is of great interest to have marker-free methods, that do not rely on additional measurements. One class of such methods corrects the center-of-rotation based solely on the measured sinogram without a need to compute additional reconstructions. These methods typically rely on inherent symmetry features in the recorded sinogram. For instance, [6] corrects the center-of-rotation based on a center-of-mass calculation, and [7,8] identifies the center-of-rotation by computing cross-correlation between projections. The main advantage of these methods is their simplicity and low computational cost. However, they are highly dependent on the quality of the data, and thus they may suffer if provided with highly noisy data. Another class of methods is based on image metrics in the reconstruction space [9,10]. The correction is done based on computed reconstructions, which makes it more robust against noise (due to the possibility of regularization). The downside of these reconstruction-based methods is increased computational cost, and they may also suffer if the reconstruction contains other artifacts, e.g. ring artifacts due to uncertainty in the source intensity. Finally, we have projection matching methods [11–13] that determine the projection geometry and reconstruction jointly. Projection matching methods work by framing the CT reconstruction problem as a joint nonlinear least-squares optimization problem for the unknown attenuation *and* the projection geometry. Due to the nonlinear coupling of the reconstruction and the projection geometry, the joint optimization problem is often solved using an alternating approach, where either the reconstruction or projection geometry is fixed while solving for the other quantity. This alternating approach is framed into an outer iterative scheme, where each iteration updates the reconstruction and projection geometry. However there is no guarantee for convergence using this alternating scheme, and many reconstructions may need to be computed to accurately estimate the projection geometry.

A common theme for the described methods is that they only yield a point-estimate of the corrected projection geometry – they do not provide a statistical measure of how (un)certain the computed estimate is. One contribution of our method is to address this

limitation. Our proposed method is based on a Bayesian framework, where we determine the projection geometry and the reconstruction jointly by approximately sampling from a posterior probability distribution using a hierarchical Metropolis-within-Gibbs sampler. In this way, we naturally obtain uncertainty estimates of the projection geometry. Through numerical experiments with real tomographic data, we compare the proposed method to a center-of-mass method [6] and a cross-correlation-based method [7], where we demonstrate that the performance of the proposed method is similar for high-quality low-noise data and that we have superior performance under challenging conditions with fewer projections or highly noisy data.

In this paper, we focus on correcting the center-of-rotation offset in 2D fan-beam CT, since this setup is common for many other geometry-correcting methods studied in the literature. However, we emphasize that the proposed method is flexible by nature and that it can be generalized to include other geometric variables and/or other scan modes, including parallel-beam and 3D CT setups.

## 2. Data description

### 2.1. SparseBeads data

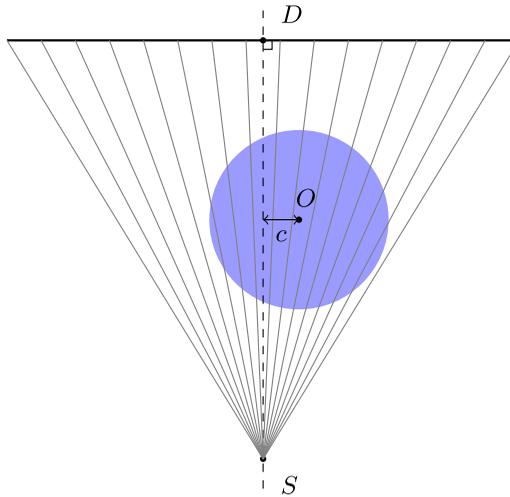
We use real tomographic data from the SparseBeads collection [14]. The data was collected using a 320/225 kV Nikon XTEK Bay micro-CT instrument, located at the Henry Mosely X-ray Imaging Facility, The University of Manchester. The scanned objects are plastic cylinders containing glass beads. The detector has  $2000 \times 2000$  pixels, and the size of a single detector pixel is  $0.2\text{mm} \times 0.2\text{mm}$ . The lab CT system employed a cone-beam setup with 2520 equidistant projections acquired over  $360^\circ$ , and the source-to-center distance and center-to-detector distance were measured as 122mm and 1400mm respectively. A single central slice was provided for each data set, corresponding to a fan-beam setup with sinograms of size  $2520 \times 2000$ . Using terminology from [14], we consider the B3L1 and Low\_B3L1 data sets in this paper. These data sets are scans of the same object, but with different exposure times per projection, 1000ms for B3L1 and 134ms for Low\_B3L1, so we have access to a high-dose and low-dose data set of the same object.

## 3. Extended reconstruction model

### 3.1. Center-of-rotation misalignment for fan-beam CT

We consider fan-beam CT setups with a linear detector. We regard the horizontal offset of the center-of-rotation from the midline as an unknown quantity, see Figure 1. We parametrize the center-of-rotation offset with the parameter  $c$ , which is defined as the signed distance from the midline to the center-of-rotation, and we define the sign such that Figure 1 illustrates a situation where  $c > 0$ . For  $c = 0$  the midline exactly passes through the center-of-rotation.

In an idealized fan-beam setup we have  $c = 0$ , but in practice the center-of-rotation may be significantly offset from the midline. A failure to accurately estimate the center-of-rotation offset in the reconstruction model may lead to reconstructions with considerable artifacts. We illustrate the effect that the center-of-rotation offset has on the reconstruction quality in Figure 2. The  $2000 \times 2000$  pixel reconstructions (pixel side length  $17.4 \mu\text{m}$ ) were



**Figure 1.** Illustration of geometric misalignment in fan-beam CT. The thin lines illustrate the radiation emitting from the source  $S$  penetrating the blue object domain with center-of-rotation  $O$ , and the horizontal black line illustrates the detector. The midline is the dotted line, which intersects with the detector perpendicularly at the point  $D$ . The center-of-rotation offset,  $c$ , is defined as the signed distance from the midline to the center-of-rotation  $O$ . The illustration depicts a configuration where  $c$  is positive.

computed using the high-dose B3L1 data, and we measure  $c$  in units of pixels such that  $c = 1$  is a center-of-rotation offset of  $17.4 \mu\text{m}$  in physical units. For a given value of  $c$ , we frame the reconstruction problem as a Tikhonov regularized least-squares optimization problem with nonnegativity constraints

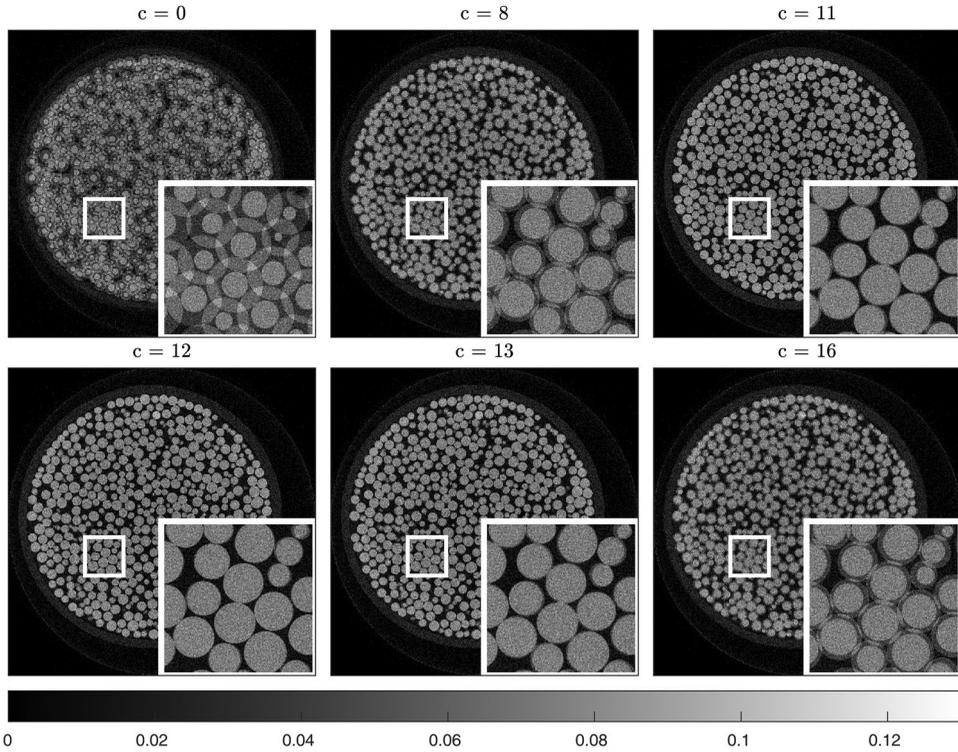
$$\min_{x \geq 0} \|A_c x - b\|_2^2 + \alpha \|x\|_2^2, \quad (1)$$

where  $b \in \mathbb{R}^m$  is the measured sinogram,  $x \in \mathbb{R}^n$  is the absorption coefficients of the object,  $A_c \in \mathbb{R}^{m \times n}$  is the linear forward projection operator parametrized by the center-of-rotation offset  $c \in \mathbb{R}$ , and  $\alpha \geq 0$  is the regularization parameter. We solve (1) using the Fast Iterative Shrinkage–Thresholding Algorithm (FISTA), where the smooth part is the sum of the two terms in (1) and the non-smooth part is the indicator function for the non-negative orthant, see [15] for details. We use the ASTRA toolbox [16,17] for the forward and back projections needed for FISTA. As we are dealing with a real data set, we do not know the true value of the center-of-rotation offset, but based on visual inspection of the reconstructions in Figure 2, the true center-of-rotation offset is close to 12 pixels for this particular data set, and we observe that relatively small deviations from this value result in considerable reconstruction artifacts that reduce the quality of the reconstruction.

### 3.2. Bayesian model

We consider an extended reconstruction model, which explicitly takes uncertainty in the center-of-rotation offset into account

$$b = A_c x + \epsilon, \quad (2)$$



**Figure 2.** Illustration of reconstructions for high-dose data using different values of the center-of-rotation offset. Based on visual inspection the optimal value is close to  $c = 12$ .

where  $\epsilon \in \mathbb{R}^m$  is additive noise and the remaining quantities are as defined in (1). Throughout we assume that the noise is Gaussian  $\epsilon \sim N(0, \lambda^{-1}I)$ , where  $\lambda > 0$  is the precision (inverse variance). Looking at problem (2) from a Bayesian perspective, we consider  $x$  and  $c$  as independent random variables, so that the likelihood of (2) is given by

$$\pi(b|x, c) \propto \exp\left(-\frac{\lambda}{2}\|A_c x - b\|_2^2\right). \quad (3)$$

The maximum likelihood (ML) solution can be obtained by maximizing (3) over the variables  $x$  and  $c$ . However, since CT is an ill-posed inverse problem, the ML solution will be dominated by high-frequency noise. To obtain high-quality reconstructions we must impose regularization, which in the Bayesian setting is done by considering prior distributions for the reconstruction and the projection geometry, that reflect our prior knowledge of the solution. Examples of common prior distributions in computational imaging include Gaussian [18], Cauchy [19], and edge-preserving priors such as total variation [20]. The cost of using more informative priors is that increasingly sophisticated (inversion) methods must be used, which may lead to increased computational cost. The model (2) has the added complexity of a stochastic forward operator, so we choose a Gaussian prior  $x|\delta \sim N(0, \delta^{-1}\Sigma_x)$  for simplicity. Here  $\delta$  controls the intensity of the prior structure imposed by the precision matrix  $\Sigma_x^{-1}$ . The quantities  $\delta$  and  $\lambda$  (referred to as hyperparameters going

forward) serve the role in the Bayesian inversion that the classical regularization parameter does in the non-Bayesian setting. We consider the hyperparameters as stochastic with conjugate Gamma hyperpriors  $\lambda \sim \text{Gamma}(\alpha_\lambda, \beta_\lambda)$  and  $\delta \sim \text{Gamma}(\alpha_\delta, \beta_\delta)$ , where we assume that the shape parameters,  $\alpha_\delta, \alpha_\lambda > 0$ , and inverse scale parameters,  $\beta_\delta, \beta_\lambda > 0$ , are known. Treating  $\delta$  and  $\lambda$  as stochastic hyperparameters imposes a natural hierarchy in the model and for this reason models with this structure are often denoted as hierarchical Bayesian models, see e.g. [21]. For the center-of-rotation offset prior, we choose a Gaussian  $c \sim N(\mu_c, \sigma_c^2)$ . The Gaussian prior allows us to incorporate knowledge of our best initial guess through the mean and account for the uncertainty through the variance. The prior distribution for the center-of-rotation offset should be seen as a modeling choice, and other options are also possible, e.g. a uniform distribution as a completely uninformative prior.

In the Bayesian setting, the solution to the inverse problem (2) is a high-dimensional posterior probability distribution  $\pi(x, c, \lambda, \delta|b)$ . Using Bayes' law and assuming independence of the variables, the posterior distribution is given by

$$\pi(x, c, \lambda, \delta|b) \propto \pi(b|x, c, \lambda)\pi(x|\delta)\pi(c)\pi(\delta)\pi(\lambda),$$

which for our concrete likelihood and priors become

$$\pi(x, c, \lambda, \delta|b) \propto \lambda^{m/2+\alpha_\lambda-1} \delta^{n/2+\alpha_\delta-1} \exp\left(-\lambda\left(\frac{1}{2}\|A_c x - b\|_2^2 + \beta_\lambda\right) - \delta\left(\frac{1}{2}\|x\|_{\Sigma_x^{-1}}^2 + \beta_\delta\right) - \frac{1}{2\sigma_c^2}(c - \mu_c)^2\right), \quad (4)$$

where  $\|x\|_{\Sigma_x^{-1}}^2 = x^T \Sigma_x^{-1} x$ . The benefits of using a Gaussian prior for the reconstruction are simplicity and computational efficiency since we have a closed-form expression for the posterior distribution. However, it is also a highly uninformative prior that does not reflect much of our prior knowledge of the solution. For this reason, we also consider an augmentation to the Gaussian prior, where we add an implicit nonnegativity assumption for the reconstruction. Nonnegativity is a natural prior in CT, since the absorption coefficients are nonnegative by nature, and taking this into account in the inversion often results in higher-quality reconstructions. To incorporate nonnegativity in the prior, we consider a mixture between two distributions: a truncated Gaussian defined on the positive orthant, and a distribution on the boundary of the positive orthant, which is obtained by appropriately transporting probability mass from outside the positive orthant. The resulting posterior distribution using this prior no longer has a simple closed-form expression, however [22,23] show that posterior distribution samples using this prior can be obtained by solving a constrained least-squares optimization problem, see next section for more details.

#### 4. Hierarchical Metropolis-within-Gibbs sampler

One way to use the posterior distribution to solve the inverse problem is to compute a maximum a posteriori (MAP) reconstruction by solving a nonlinear optimization problem. However, this does not provide us with any estimates of the uncertainty. Moreover, the fact that the forward projection operator depends explicitly on the center-of-rotation offset  $c$  complicates the resulting optimization problem significantly since  $A_c$  is a highly nonlinear function of  $c$ . For these reasons, we take a sampling approach, and we will develop a

hierarchical Metropolis-within-Gibbs sampler to approximately sample from the full posterior distribution. A Gibbs sampler samples from the posterior distribution by iteratively sampling from all the conditional distributions. The overall structure of the sampler we propose is inspired by the approach presented in [21].

#### 4.1. Conditional distributions

To efficiently use a Gibbs sampler to sample from the posterior distribution, we need efficient ways of sampling from the conditionals of the posterior distribution (denoted henceforth as conditional posterior distributions). The conditional posterior distributions for the hyperparameters are

$$\pi(\lambda|x, b, c) \propto \lambda^{m/2+\alpha_\lambda-1} \exp\left(-\lambda\left(\frac{1}{2}\|A_c x - b\|_2^2 + \beta_\lambda\right)\right), \quad (5a)$$

$$\pi(\delta|x, b) \propto \delta^{n/2+\alpha_\delta-1} \exp\left(-\delta\left(\frac{1}{2}\|x\|_{\Sigma_x^{-1}}^2 + \beta_\delta\right)\right). \quad (5b)$$

We immediately recognize these as Gamma distributions

$$\begin{aligned} \lambda|x, b, c &\sim \text{Gamma}\left(m/2 + \alpha_\lambda, \frac{1}{2}\|A_c x - b\|_2^2 + \beta_\lambda\right), \\ \delta|x, b &\sim \text{Gamma}\left(n/2 + \alpha_\delta, \frac{1}{2}\|x\|_{\Sigma_x^{-1}}^2 + \beta_\delta\right). \end{aligned}$$

The conditional posterior distribution for the reconstruction is given by

$$\pi(x|\lambda, \delta, b, c) \propto \exp\left(-\frac{1}{2}\left(\lambda\|A_c x - b\|_2^2 + \delta\|x\|_{\Sigma_x^{-1}}^2\right)\right).$$

Algebraic manipulations show that this is a multivariate Gaussian

$$x|\lambda, \delta, b, c \sim N\left(\left(\lambda A_c^T A_c + \delta \Sigma_x^{-1}\right)^{-1} \lambda A_c^T b, \left(\lambda A_c^T A_c + \delta \Sigma_x^{-1}\right)^{-1}\right). \quad (6)$$

Sampling from a high-dimensional multivariate Gaussian distribution is generally expensive, however, the distribution coincides with the family of solutions to a stochastic regularized least-squares problem

$$x(\xi_m, \xi_n) = \underset{x}{\operatorname{argmin}} \left\{ \frac{\lambda}{2} \|A_c x - b - \lambda^{-1/2} \xi_m\|_2^2 + \frac{\delta}{2} \|\Sigma_x^{-1/2} x - \delta^{-1/2} \xi_n\|_2^2 \right\}, \quad (7)$$

where  $\xi_m \in \mathbb{R}^m$  and  $\xi_n \in \mathbb{R}^n$  are Gaussian *i.i.d.* samples, see [21] for details. This suggests that approximate samples from the multivariate Gaussian distribution can be obtained by inexact solutions of the stochastic optimization problem.

If we instead want to obtain samples using the implicit nonnegativity prior, we follow [22,23] and define the conditional posterior for the reconstruction as the family of solutions

to the constrained stochastic regularized least-squares problem

$$x(\xi_m, \xi_n) = \underset{x \geq 0}{\operatorname{argmin}} \left\{ \frac{\lambda}{2} \|A_c x - b - \lambda^{-1/2} \xi_m\|_2^2 + \frac{\delta}{2} \|\Sigma_x^{-1/2} x - \delta^{-1/2} \xi_n\|_2^2 \right\}. \quad (8)$$

The only difference compared to the model without nonnegativity is the addition of the constraint. Again following [22,23], the conditional prior for the hyperparameter  $\delta$  when using the implicit nonnegativity prior takes the form

$$\pi(\delta|x, b) = \delta^{\bar{n}/2 + \alpha_\delta - 1} \exp\left(-\delta \left(\frac{1}{2} \|x\|_{\Sigma_x^{-1}} + \beta_x\right)\right), \quad (9)$$

where  $\bar{n}$  is the number of non-zero entries in  $x$ .

Finally, the conditional posterior distribution for  $c$  is

$$\pi(c|x, \lambda, b) \propto \exp\left(-\frac{\lambda}{2} \|A_c x - b\|_2^2 - \frac{1}{2\sigma_c^2} (c - \mu_c)^2\right). \quad (10)$$

This is not a simple closed-form distribution due to the stochastic forward operator  $A_c$ , so we must use a Markov Chain Monte Carlo (MCMC) method to obtain samples from this conditional distribution.

## 4.2. Implementation

To obtain a sample from the conditional posterior of the reconstruction, we must solve (7) (or (8) in the case of nonnegativity constraints) with a high level of accuracy. However, this is too costly for high-dimensional problems, where we may need thousands of samples. To address this issue and increase computational efficiency, we use FISTA with a (small) fixed number of iterations to obtain approximate solutions to (7) (or (8) in case of nonnegativity constraints). Early termination of FISTA incurs an approximation error in the conditional sample of the reconstruction, which implies that we cannot guarantee the stationarity of the full Metropolis-within-Gibbs sampler. To mitigate the error caused by truncated iterations, we use the previous sample as the initial guess for the FISTA iterations as a warm-start procedure. Truncation of an iterative method to sample from a multivariate Gaussian resembles the Gradient Scan method presented in [24]. The difference is that the Gradient Scan method uses a fixed number of conjugate gradient iterations instead of FISTA iterations. Using properties of the conjugate gradient iterates, the authors of [24] prove asymptotic convergence. However, their proof does not extend to the case with nonnegativity constraints. Using FISTA instead of the conjugate gradient method has the benefit of being easily adapted to solve the nonnegativity case (8).

To obtain samples from the conditional posterior distribution for the center-of-rotation offset, we use Metropolis-Hastings, see e.g. [25], with a symmetric Gaussian random walk proposal distribution

$$q(\hat{c}|c) \propto \exp\left(-\frac{1}{2s^2} (\hat{c} - c)^2\right),$$

where  $s > 0$  is the *step size* in the Metropolis-Hastings method, which can be tuned to obtain a reasonable acceptance rate in the Metropolis-Hastings sampling procedure. The

optimal acceptance rate is highly dependent on the problem, but we generally aim for around 25% as is also discussed in [26]. The Metropolis-Hastings algorithm is shown in Algorithm 1.

---

**Algorithm 1** Metropolis-Hastings for center-of-rotation offset
 

---

**Input:** reconstruction,  $x \in \mathbb{R}^n$ ; initial value of center-of-rotation offset,  $c^0 \in \mathbb{R}$ ; center-of-rotation offset Gaussian mean,  $\mu_c \in \mathbb{R}$ ; center-of-rotation offset Gaussian standard deviation,  $\sigma_c > 0$ ; Metropolis-Hastings step size,  $s > 0$ ; number of Metropolis-Hastings iterations,  $k_{\text{metro}}$ ; noise hyperparameter,  $\lambda > 0$ ; sinogram,  $b \in \mathbb{R}^m$ .

**Output:** Sample from  $\pi(c|x, \lambda, b)$ .

**for**  $k = 1$  to  $k_{\text{metro}}$  **do**

Sample from Gaussian proposal:  $c^{\text{prop}} \sim N(c^{k-1}, s^2)$

Compute acceptance probability:  $p^{\text{accept}} = \min\{1, \frac{\pi(c^{\text{prop}}|x, \lambda, b)}{\pi(c^{k-1}|x, \lambda, b)}\}$

Set  $c^k = c^{\text{prop}}$  with probability  $p^{\text{accept}}$ . Otherwise let  $c^k = c^{k-1}$ .

**end for**

---

As we use Metropolis-Hastings to sample from one of the conditional distributions our proposed sampler for the full posterior distribution is a hierarchical Metropolis-within-Gibbs sampler, see e.g. [27,28] for details on Metropolis-within-Gibbs samplers. The full sampler using the prior without nonnegativity is presented in Algorithm 2.

---

**Algorithm 2** Hierarchical Metropolis-within-Gibbs Sampler
 

---

**Input:** initial reconstruction,  $x^0 \in \mathbb{R}^n$ ; initial value of center-of-rotation offset,  $c^0 \in \mathbb{R}$ ; center-of-rotation offset Gaussian mean,  $\mu_c \in \mathbb{R}$ ; center-of-rotation offset Gaussian standard deviation,  $\sigma_c > 0$ ; Metropolis-Hastings step size,  $s > 0$ ; number of Gibbs samples,  $k_{\text{Gibbs}}$ ; number of Metropolis-Hastings iterations,  $k_{\text{Metro}}$ ; number of FISTA iterations,  $k_{\text{FISTA}}$ ; Gamma distribution parameters,  $\gamma_\lambda, \gamma_\delta, \beta_\lambda, \beta_\delta > 0$ ; sinogram,  $b \in \mathbb{R}^m$ .

**Output:** approximate sample chain  $\{(x^k, c^k, \lambda^k, \delta^k)\}_{k=1}^{k_{\text{Gibbs}}}$  from full posterior distribution  $\pi(x, c, \lambda, \delta|b)$ .

**for**  $k = 1$  to  $k_{\text{Gibbs}}$  **do**

*Sample conditional hyperparameters from conjugate Gamma distributions*

$\lambda^k|x^{k-1}, c^{k-1}, b \sim \text{Gamma}(m/2 + \alpha_\lambda, 1/2\|A_{c^{k-1}}x^{k-1} - b\|_2^2 + \beta_\lambda)$

$\delta^k|x^{k-1}, b \sim \text{Gamma}(n/2 + \alpha_\delta, 1/2\|x\|_{\Sigma_x^{-1}}^2 + \beta_\delta)$

*Sample conditional center-of-rotation offset using Metropolis-Hastings*

*Sample  $c^k|x^{k-1}, \lambda^k, \delta^k, b$  using  $k_{\text{Metro}}$  iterations of Algorithm 1*

*Sample conditional reconstruction using truncated FISTA*

*Approximate  $x^k|c^k, \lambda^k, \delta^k, b$  by running  $k_{\text{FISTA}}$  FISTA iterations on (7) with initial condition  $x^{[0]} = x^{k-1}$ .*

**end for**

---

Using the nonnegativity prior, Algorithm 2 must be changed in the following way: the conditional  $\delta$ -sample step must be changed to (9), and the conditional reconstruction sample steps must instead run on (8). The remaining steps are unchanged.

A good choice for the initial reconstruction  $x^0 \in \mathbb{R}^n$  can be obtained by solving a (nonnegativity-constrained) regularized least-squares problem like (1) using the initial center-of-rotation offset  $c^0$ .

In theory, we only need to take a single Metropolis-Hastings iteration,  $k_{\text{Metro}} = 1$ , for each full Gibbs iteration [29]. However, we found that taking multiple Metropolis-Hastings steps speeds up the estimation of the center-of-rotation offset considerably. The reason for this is that taking multiple Metropolis-Hastings steps allows for potentially larger steps for each full Gibbs iteration, while also reducing the risk of complete rejections when sampling the center-of-rotation offset. Complete rejections are undesirable since it is considerably more expensive to obtain a sample of the reconstruction compared to a sample of the center-of-rotation offset. In practice, we found that  $k_{\text{Metro}} = 10$  yielded a nice balance between avoiding rejections, and also not using too many computational resources on the Metropolis-Hastings step. For the number of FISTA iterations, we use  $k_{\text{FISTA}} = 20$ , since we empirically found this to provide a reasonable balance between performance and computational effort.

*Computational remarks.* The computationally demanding steps of Algorithm 2 is sampling the center-of-rotation offset and sampling the reconstruction. Sampling the center-of-rotation offset requires  $k_{\text{Metro}}$  forward projections, and sampling the reconstruction requires  $k_{\text{FISTA}}$  forward and back projections. The cost of a forward projection and a back projection is approximately equal, so the total cost of the algorithm is  $k_{\text{Gibbs}}(2k_{\text{FISTA}} + k_{\text{Metro}})$  forward projections. With our specific values for  $k_{\text{FISTA}}$  and  $k_{\text{Metro}}$ , the total cost of the MCMC algorithm is  $2.5k_{\text{Gibbs}}k_{\text{FISTA}}$  forward projections. Hence the computational scaling of the algorithm is similar to computing a MAP estimate using an iterative method, although the constant  $2.5k_{\text{Gibbs}}$  may need to be large to obtain a satisfying reconstruction with a good estimate of the center-of-rotation offset. We remark that the proposed method also samples the regularization parameter, which implies that we do not need to use additional resources to tune this quantity, which must be done for classical MAP reconstruction.

## 5. Numerical experiments

In this section, we present numerical experiments using the proposed MCMC method. All experiments shown were conducted using the model with the implicit nonnegativity assumption, since this yielded superior results compared to the model without nonnegativity. We compare the MCMC method with two other methods that are commonly used to correct the geometry, both of which are computationally cheap and easy to implement since they are based solely on the recorded sinogram. The first method, COM, is from [6], and it is based on a center-of-mass calculation of the sinogram. The second method, XCORR, is from [7], and it is based on auto-correlation of the sum of projections in the sinogram. Both the COM and XCORR methods make a correction based on symmetry features in the sinogram. For this reason, they both require  $360^\circ$  projection data, whereas the MCMC method makes no assumption regarding angular range of the projections.

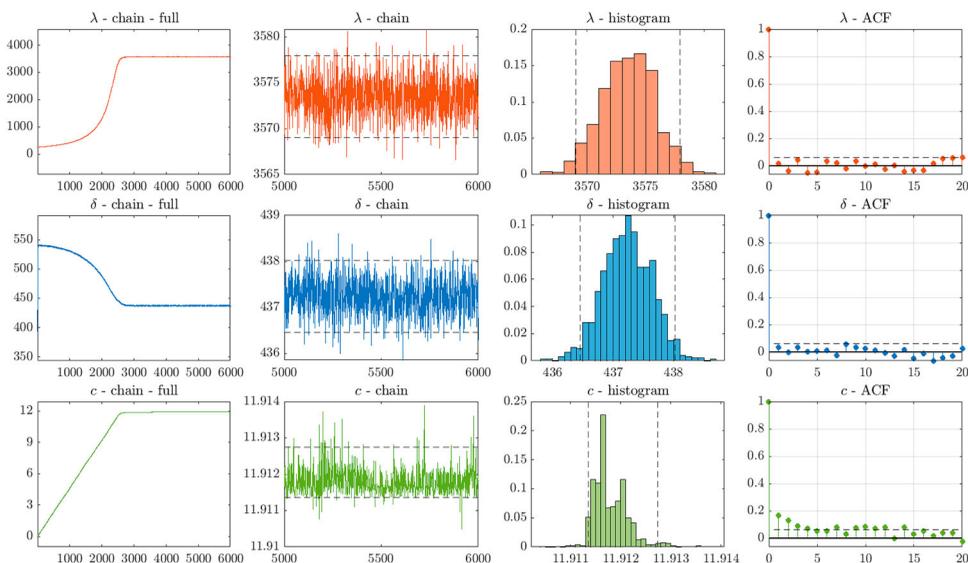
Algorithm 2 has a number of parameters that need to be set. Through our experiments, we found that a wide range of parameter choices lead to similar performance. The most sensitive parameter appears to be the Metropolis-Hastings step size  $s$ . If this value is too small, the algorithm will need an excessive number of iterations, and if it is too large, the algorithm may stall because the vast majority of proposed steps are rejected. In our implementation, we manually tune this parameter for each problem to obtain a reasonable acceptance rate, but a straightforward extension is to use a Metropolis-Hastings algorithm with adaptive step size, see [30] for more details. Another important quantity to consider is the number of samples. As is the case for general MCMC algorithms, it is difficult to determine a priori how many samples are necessary to obtain good results. Early termination may result in a reduced quality of the reconstruction and/or a larger uncertainty due to the contribution of samples from the burn-in phase. In our implementation, we generate a fixed number of samples  $k_{\text{Gibbs}}$ . Another option is to continuously monitor chain diagnostics and use these as a means to determine when to stop sampling, see [21] for more details.

We conduct numerical experiments for both the full  $360^\circ$  high-dose and low-dose data. For the high-dose data, we also conduct numerical experiments using a short-scan setup, where we limit the angular coverage to illustrate the versatility of the proposed method. The results consist of reconstructions as well as sample chains for the MCMC method to see the convergence of the hyperparameters and center-of-rotation offset.

### 5.1. High-dose data

We use  $c^0 = 0.0$ ,  $\mu_c = 0.0$ , and  $\sigma_c = 20.1$  (all measured in the unit of pixels) for the center-of-rotation initial guess and prior parameters, and for the precision matrix of the reconstruction  $\Sigma_x^{-1}$ , we choose the identity. For the Metropolis-Hastings step size, we use  $s = 1.1 \cdot 10^{-3}$  (measured in the unit of pixels). For the hyperprior parameters, we choose  $\alpha_\lambda = \alpha_\delta = 1$  and  $\beta_\lambda = \beta_\delta = 10^{-4}$  as recommended in [21], yielding weakly informative exponential hyperpriors. We use  $k_{\text{Metro}} = 10$  and  $k_{\text{FISTA}} = 20$  for the number of Metropolis iterations and FISTA iterations respectively. We take a total of  $k_{\text{Gibbs}} = 6000$  samples, and we do not do any thinning. The resulting MCMC-chains can be seen in Figure 3; the full chains have three distinct phases – a fast transient phase for the first roughly 2500 samples, where the hyperparameters and center-of-rotation offset rapidly change, a slower transient phase for the next 2500 samples, where the samples slowly approach the equilibrium, and finally an equilibrium phase for the final 1000 samples, where we obtain approximate samples from the posterior distribution. Based on visual inspection, the chains look stationary after discarding the first 5000 samples for burn-in. We note that the autocorrelation for the center-of-rotation offset chain is slightly higher than for the hyperparameters, due to the fact that we use Metropolis-Hastings to sample this. The center-of-rotation offset in the stationary phase is around 12 pixels, which aligns with the observation made in Figure 2. In addition, we observe that the variance of the  $c$ -chain is relatively low. Looking at the histogram the distribution of the center-of-rotation offset looks slightly skewed, but this is likely due to the relatively few samples taken, and we expect a bell-shaped distribution as we increase the number of samples.

The first column of Figure 6 shows the reconstructions using the initial, the MCMC corrected, the COM corrected, and the XCORR corrected center-of-rotation offset for the high-dose data. The MCMC reconstruction is the posterior sample mean after the burn-in

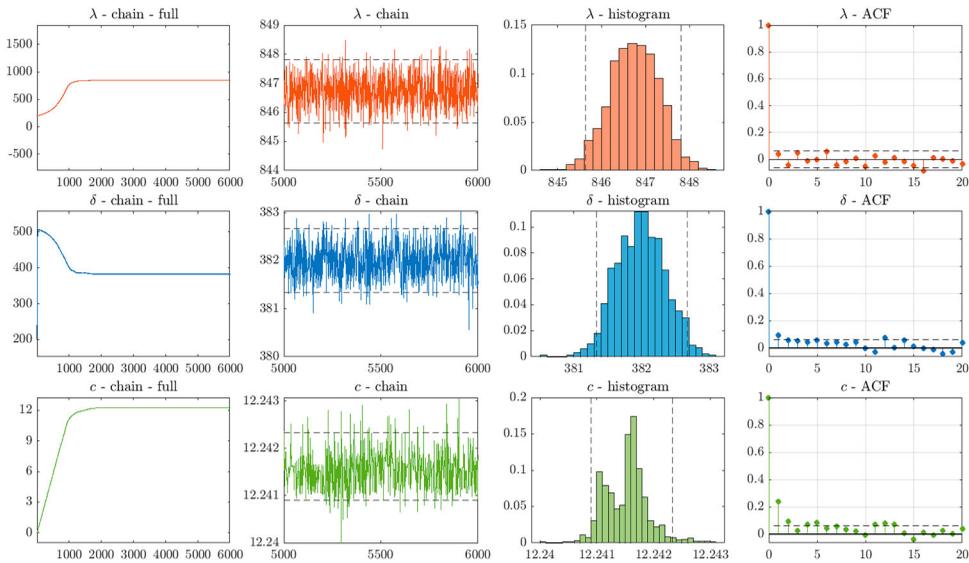


**Figure 3.** MCMC chains for high-dose SparseBeads data. The first column shows the full MCMC chains, the second column shows the chains after discarding the first 5000 samples as burn-in, the third column shows histograms after burn-in, and the final column is the autocorrelation function of the samples after burn-in. The dotted black lines indicate empirical 95% credibility intervals. Based on visual inspection all chains look stationary after burn-in.

phase. The reconstructions for the initial, the COM, and the XCORR center-of-rotation offsets were obtained by solving (1) to a high level of accuracy, where the regularization parameter was determined from the means of the hyperparameter chains after the burn-in phase. This is a fair comparison, since the solution to (1) is a MAP estimate for the simpler reconstruction problem, where  $c$ ,  $\lambda$ , and  $\delta$  are assumed known, and you let  $\alpha := \delta/\lambda$ . We see a vast increase in reconstruction quality from all the correction methods compared to the initial estimate  $c = 0$ . The MCMC and XCORR reconstructions are slightly sharper compared to the COM method, but all methods produce high-quality reconstructions. It takes approximately 4 seconds to obtain a single sample for the MCMC method, so the total running time is on the scale of hours, whereas simply running the COM and XCORR method on the sinogram takes on the order of seconds. However, the MCMC method also produces a reconstruction and an estimate of the regularization parameter, which must be computed separately for the simpler geometry correcting methods.

## 5.2. Low-dose data

We use the same algorithm parameters as for the high-dose experiment, except we use  $s = 2.9 \cdot 10^{-3}$ . The MCMC chains are shown in Figure 4. The biggest difference between the high-dose and the low-dose MCMC chains is that the sampled  $\lambda$ -values are much smaller in the low-dose case. This is expected since  $\lambda$  represents the inverse variance of the noise, and we have more noise for the low-dose data. Interestingly, the noise level does not seem to affect the credibility interval of the  $c$ -estimate in this case, which is approximately the same width in the high-dose and low-dose case.



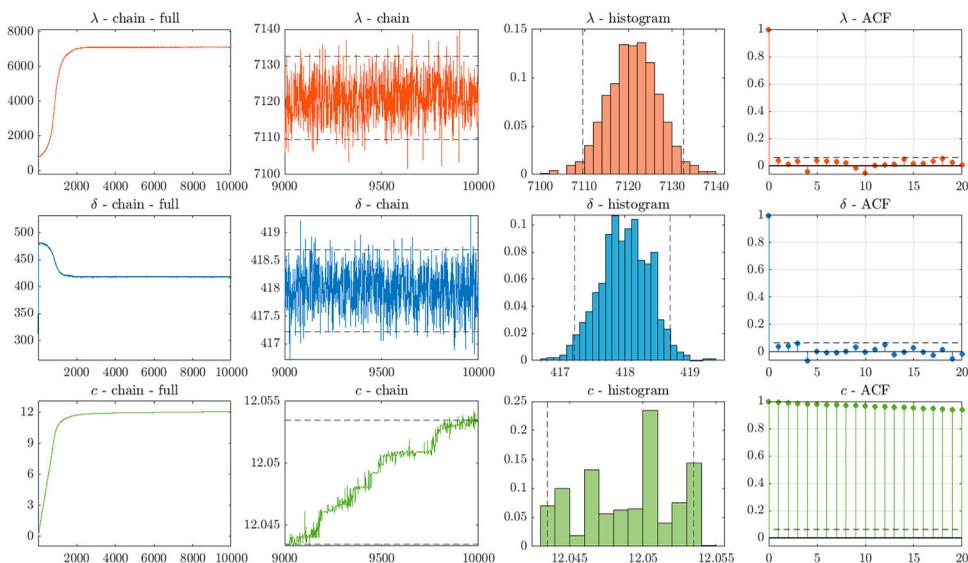
**Figure 4.** MCMC chains for low-dose SparseBeads data. Based on visual inspection the chains after burn-in look stationary. Notice that the values for the inverse noise variance  $\lambda$  is much smaller compared to the high-dose case, reflecting the higher noise levels in the low-dose data.

The second column in Figure 6 shows the reconstructions for the low-dose data using the different geometry correction techniques. In general, we see that the reconstructions are noisier compared to the high-dose reconstructions, which is expected since the signal-to-noise ratio is lower for the low-dose data. All geometry correction methods demonstrate improved reconstruction quality compared to the initial center-of-rotation offset. However, contrary to the high-dose data, we see a significant improvement in clarity and sharpness for the MCMC method compared to the COM and XCORR method. This indicates that the center-of-rotation offset from the MCMC method is closer to the actual value compared to the estimates produced by the other methods and that the MCMC method is more robust against noisy data. The run time is similar to the experiment for the high-dose data.

### 5.3. Short-scan data

We again use the high-dose data, but we restrict the angular range to  $[0^\circ, 210^\circ]$  to give a short-scan setup. Thus the requirement regarding  $360^\circ$  projection data for the COM and XCORR method is not satisfied, so we can not expect good performance using these methods. We use the same algorithm parameters as for the low-dose experiment except that we take a total of 10000 samples and discard the first 9000 for burn-in, and we use  $s = 2.9 \cdot 10^{-3}$ . The resulting MCMC chain can be seen in Figure 5. We see that the reduced angular coverage has an adverse effect on the apparent convergence rate of the MCMC chain, and the chain is still not stationary after 10000 samples. However, we see very slow movement in the final 1000 samples, and the estimate does look close to the optimal based on earlier experiments.

The short-scan reconstructions can be seen in the third column of Figure 6. Here we see a vast increase in the quality of the reconstruction using the MCMC method compared



**Figure 5.** MCMC chains for high-dose short-scan SparseBeads data. Notice that the  $c$ -chain does not look stationary in this case.

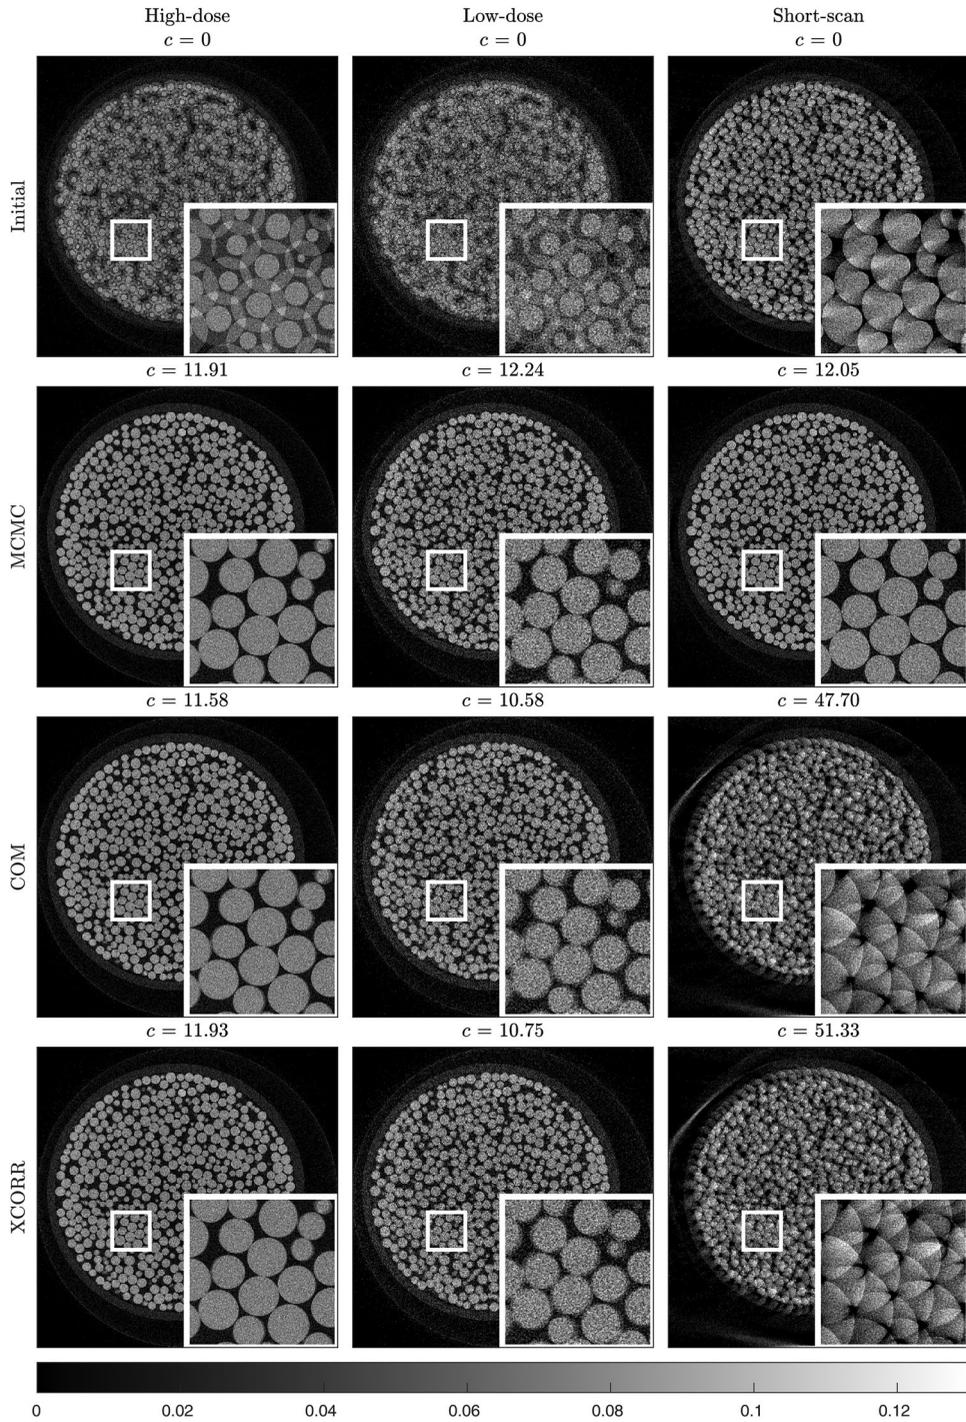
to the COM and XCORR method, both of which fail to give a meaningful estimate. One may argue that the comparison in this case is not fair, since we do not satisfy the  $360^\circ$  requirement needed for the COM and XCORR method. However, we include the results to show a challenging non-standard scenario, where a more flexible and computationally expensive approach may be needed to obtain good results. Notice that the reconstruction for the MCMC method is slightly noisier compared to the full angle case due to the fewer projections in this setup. The running time is similar to the experiment for the high-dose data.

## 6. Discussion and conclusion

We have demonstrated that explicitly taking geometric uncertainties into account in an extended reconstruction model results in high-quality reconstructions under challenging conditions with real tomographic data. Overall the results show that the MCMC method performs similarly to more conventional sinogram-based methods like COM and XCORR for the high-dose  $360^\circ$  data and that we have superior performance for the more challenging low-dose and short-scan setup.

The promising results for low-dose and short-scan setups could motivate an investigation of the robustness of the method on problems in interior or low-contrast tomography. These types of problems have a distinct lack of sharp features, which complicates the calibration of the geometric parameters. We suspect that the lack of sharp features will also make inference much harder for our proposed MCMC method since misalignments are often revealed in areas with sharp high-contrast features. However, it is an interesting perspective for future work.

The biggest concern regarding the use of MCMC methods for high-dimensional inverse problems is the computational complexity. We have addressed this by developing an



**Figure 6.** Reconstructions for different data sets and different geometry correction techniques.

efficient Metropolis-within-Gibbs sampler that obtains approximate samples from the posterior distribution aided by modern optimization techniques. We stress that the proposed method is more expensive compared to many other geometry-correcting methods, and for high-quality low-noise data there is no need to use the expensive MCMC machinery since the resulting estimates and reconstructions are comparable in quality to cheaper simpler methods. However, for challenging high-noise and/or short-scan setups, the proposed method may yield superior results. We also remark that simple geometry correction methods can work excellently in tandem with the proposed MCMC method – we can obtain an excellent initial guess of the geometry using the simple methods, and then subsequently improve upon the estimate using the MCMC method. This may also significantly reduce the computational cost of the MCMC method since we can skip a large part of the initial transient phase in the MCMC chains.

The focus of this paper has been on 2D fan-beam CT with uncertainty in the center of rotation. However the proposed methodology is by nature highly flexible, and it can be extended in many ways: Firstly, since the methodology only relies on the ability to perform forward and back projections with a given projection geometry, we are not limited to a particular scan mode, and an interesting perspective would be to use the method for parallel-beam, 3D cone-beam, or other more custom scan modes. Secondly, it would be interesting to include multiple geometric parameters. For instance, in fan-beam setups, some geometric parameters of interest are the horizontal and vertical position of the radiation source and detector. Estimating multiple parameters simultaneously potentially allows us to correct more challenging misalignments in the projection geometry, which may result in higher-quality reconstructions. The only change needed in the proposed algorithm is including more geometric variables in the Metropolis-Hastings step. Finally, it would be interesting to extend the model by considering more informative reconstruction prior distributions instead of the Gaussian. This may improve the quality of the reconstructions, and it may also assist in obtaining even more accurate estimates of the geometric parameters.

## Disclosure statement

The authors of this manuscript have no conflicts of interest to disclose.

## Funding

This work was supported by The Novo Nordisk Foundation [Grant No. NNF20OC0061894] and by The Villum Foundation [Grant No. 25893].

## Data availability

The code and data to reproduce the results and generate the figures of the article are available from [31,32], respectively.

## ORCID

Frederik Hagsholm Pedersen  <http://orcid.org/0000-0003-4655-7763>

Jakob Sauer Jørgensen  <http://orcid.org/0000-0001-9114-754X>

Martin Skovgaard Andersen  <http://orcid.org/0000-0002-4654-3946>

## References

- [1] Hansen PC, Jørgensen JS, Lionheart WRB. Computed tomography: algorithms, insight, and just enough theory. Philadelphia: SIAM; 2021. (Fundamentals of algorithms).
- [2] Rubin GD. Computed tomography: revolutionizing the practice of medicine for 40 years. *Radiology*. 2014;273:S45–S74.
- [3] Cnudde V, Boone MN. High-resolution X-ray computed tomography in geosciences: a review of the current technology and applications. *Earth Sci Rev*. 2013;123:1–17.
- [4] Salvo L, Suery M, Marmottant A, et al. 3d imaging in material science: application of X-ray tomography. *C R Phys*. 2010;11:641–649.
- [5] Yang K, Kwan ALC, Miller DE, et al. A geometric calibration method for cone beam CT systems. *Med Phys*. 2006;33:1695–1706.
- [6] Azevedo SG, Schneberk DJ, Fitch JP, et al. Calculation of the rotational centers in computed tomography sinograms. *IEEE Trans Nucl Sci*. 1990;37:1525–1540.
- [7] Yang M, Pan J, Zhang J, et al. Center of rotation automatic measurement for fan-beam CT system based on sinogram image features. *Neurocomputing*. 2013;120:250–257.
- [8] Vo NT, Drakopoulos M, Atwood RC, et al. Reliable method for calculating the center of rotation in parallel-beam tomography. *Opt Express*. 2014;22:19078–19086.
- [9] Cheng CC, Ching YP, Ko YH. Correction of center of rotation and projection angle in synchrotron X-ray computed tomography. *Sci Rep*. 2018;8:9884.
- [10] Donath T, Beckmann F, Schreyer A. Automated determination of the center of rotation in tomography data. *J Opt Soc Am A*. 2006;23:1048–1057.
- [11] Houben L, Sadan M. Refinement procedure for the image alignment in high-resolution electron tomography. *Ultramicroscopy*. 2011;111:1512–1520.
- [12] Parkinson D, Knoechel C, Yang C, et al. Automatic alignment and reconstruction of images for soft X-ray tomography. *J Struct Biol*. 2012;177:259–266.
- [13] Ramos T, Jørgensen JS, Andreassen JW. Automated angular and translational tomographic alignment and application to phase-contrast imaging. *J Opt Soc Am A*. 2017;34:1830–1843.
- [14] Jørgensen JS, Coban SB, Lionheart WRB, et al. SparseBeads data: benchmarking sparsity-regularized computed tomography. *Meas Sci Technol*. 2017;28:124005.
- [15] Beck A, Teboulle M. A fast iterative Shrinkage–Thresholding algorithm for linear inverse problems. *SIAM J Imaging Sci*. 2009;2:183–202.
- [16] van Aarle W, Palenstijn WJ, Cant J, et al. Fast and flexible X-ray tomography using the ASTRA toolbox. *Opt Express*. 2016;24:25129–25147.
- [17] van Aarle W, Palenstijn WJ, Beenhouwer JD, et al. The ASTRA toolbox: a platform for advanced algorithm development in electron tomography. *Ultramicroscopy*. 2015;157:35–47.
- [18] Knapik BT, van der Vaart AW, van Zanten JH. Bayesian inverse problems with Gaussian priors. *Ann Stat*. 2011;39:2626–2657.
- [19] Markkanen M, Roininen L, Huttunen J, et al. Cauchy difference priors for edge-preserving Bayesian inversion. *J Inverse Ill-posed Probl*. 2019;27:225–240.
- [20] Pereyra M. Proximal Markov chain Monte Carlo algorithms. *Stat Comput*. 2016;26:745–760.
- [21] Bardsley JM. Computational uncertainty quantification for inverse problems. New York: SIAM; 2018. (Computational science and engineering).
- [22] Bardsley JM, Fox C. An MCMC method for uncertainty quantification in nonnegativity constrained inverse problems. *Inverse Probl*. 2012;20:477–498.
- [23] Bardsley JM, Hansen PC. MCMC algorithms for computational UQ of nonnegativity constrained linear inverse problems. *SIAM J Sci Comput*. 2020;42:A1269–A1288.
- [24] Féron O, Orioux F, Giovannelli J. Gradient scan Gibbs sampler: an efficient algorithm for high-dimensional Gaussian distributions. *IEEE J Sel Top Signal Process*. 2016;10:343–352.
- [25] Robert CP, Casella G. Monte Carlo statistical methods. New York: Springer; 2004. (Springer texts in statistics).
- [26] Roberts GO, Rosenthal JS. Optimal scaling for various Metropolis–Hastings algorithms. *Stat Sci*. 2001;16:351–367.

- [27] Müller P. A generic approach to posterior integration and Gibbs sampling Tech. Rep. 91-09. Department of Statistics, Purdue University; 1991.
- [28] Tierney L. Markov chains for exploring posterior distributions. *Ann Stat.* **1994**;22:1701–1762.
- [29] Gilks WR, Best NG, Tan KKC. Adaptive rejection Metropolis sampling within Gibbs sampling. *J R Stat Soc.* **1995**;44:455–472.
- [30] Andrieu C, Thomas J. A tutorial on adaptive MCMC. *Stat Comput.* **2008**;18:343–373.
- [31] Pedersen FH, Jørgensen JS, Andersen MS. Code accompanying article ‘A Bayesian approach to CT reconstruction with uncertain geometry’. 2022. DOI:[10.5281/zenodo.7411881](https://doi.org/10.5281/zenodo.7411881).
- [32] Pedersen FH, Andersen MS, Jørgensen JS. Data accompanying article ‘A Bayesian approach to CT reconstruction with uncertain geometry’. 2022. DOI:[10.5281/zenodo.7408468](https://doi.org/10.5281/zenodo.7408468).