Bayesian Uncertainty Quantification and Regularized Reconstruction for CT-Based Dimensional Metrology

Negin Khoeiniha[a,b,*], Patricio Guerrero[a,b], Tristan van Leeuwen[c,d], Wim Dewulf[a,b]

[a] *Department of Mechanical Engineering, KU Leuven, 3001 Leuven, Belgium*
[b] *M&A Corelab, Flanders Make @ KU Leuven, 3001 Leuven, Belgium*
[c] *Computational Imaging, Centrum Wiskunde Informatica, Amsterdam, The Netherlands*
[d] *Mathematical Institute, Utrecht University, Utrecht, The Netherlands*

Abstract

Statistical methods within the Bayesian framework have been widely used to address inverse imaging problems, such as computed tomography (CT) image reconstruction. These methods offer a probabilistic approach that is able to enhance the reconstruction quality by employing regularization methods while enabling uncertainty quantification of the result, providing valuable insights into the reliability of the reconstructed images. However, despite the flexibility and range of techniques within this framework, the computational intensity of this class of approaches is still impractical for large-scale datasets like those in CT. In this manuscript, we introduce a concept for determining the uncertainty in CT-based dimensional measurement using a rapid, regularized, Markov Chain Monte Carlo reconstruction technique. This method provides a volumetric model where each voxel is represented by a distribution, which is then transformed into a triplet of gray value models: one for the central value and one each for the upper and lower bound of the confidence interval. Bi-directional and uni-directional length measurements on results derived from each single-gray-value model, for real CT data, provide a task-specific measurement uncertainty. This method requires significantly less computation and storage capacity compared to classic Monte Carlo simulations by reducing the number of needed simulations for approximating a distribution while incorporating regularization techniques. The results are compared to the conventional non-regularized and regularized reconstruction methods such as Feldkamp–David–Kress (FDK) and gradient descent, followed by validating the determined uncertainty in real CT data.

Keywords: X-ray Computed Tomography, Uncertainty Quantification, Markov Chain Monte Carlo Simulations, Dimensional Metrology

1. Introduction

Computed tomography (CT) is a non-destructive imaging technique that has various industrial applications due to its ability to provide detailed insights into the internal parts of the object, allowing for the inspection of structures and dimensions of industrial components. CT scanning involves exposing an object to X-ray radiation from various angles and determining the reduction

*Corresponding author
Email addresses: negin.khoeiniha@kuleuven.be (Negin Khoeiniha), patricio.guerrero@kuleuven.be (Patricio Guerrero), T.van.Leeuwen@cwi.nl (Tristan van Leeuwen), wim.dewulf@kuleuven.be (Wim Dewulf)

in radiation intensity caused by absorption by the object. After capturing numerous projections, mathematical methods are used to reconstruct a three-dimensional (3D) representation of the object, showcasing its internal structure [1, 2].

Various methodologies have been developed for CT reconstruction [3], including analytical methods, e.g. Filtered Back Projection (FBP), and Feldkamp–David–Kress (FDK) [4], iterative methods like Simultaneous Iterative Reconstruction Technique (SIRT) [5], and statistical methods. For instance, Markov Chain Monte Carlo (MCMC) family of techniques [6]. Each methodology offers some advantages and limitations concerning factors such as accuracy, computational efficiency, and performance under conditions of data limitation and noise, for instance, regularized iterative methods are more suitable for handling limited data [7, 8]. Despite many advancements in each class of reconstruction method, there is still a need for a computationally efficient approach that can provide enough information to determine the uncertainty of the reconstruction result.

Multiple sources of error contribute to the final uncertainty of the reconstruction results [2, 9]. These sources, including the scanning system, numerical methods in reconstruction, and surface determination, are summarized in Table 1.

Table 1: Influence factors in CT

Source	**Uncertainty Contributors**
CT system	X-ray source
	Detector
	Kinematic system
	Geometrical misalignment
Data processing	3D reconstruction
	Data preprocessing and postprocessing.
	Segmentation and surface determination
Environment	Temperature
	Vibrations
	Humidity
Operator	Workpiece fixturing and orientation
	X-ray source settings
	Number of projections
	Measurement strategy

1.1. Related work

Considering the various sources of error in CT, uncertainty quantification (UQ) is an important task for industrial CT reconstruction and conducting dimensional measurements, which affects decision-making, and reliability of the measurement result [9].

UQ in dimensional measurements has long been focused on through standards such as the Guide to the Expression of Uncertainty in Measurement (GUM) [10], which provides a framework for evaluating and expressing measurement uncertainty in a standardized way. While GUM-based approaches are well-suited for analytical or model-based uncertainty evaluations, they often become less manageable for complex measurement systems like CT, where multiple sources of error are present. In such cases, simulation-based approaches such as Monte Carlo methods have gained popularity due to their flexibility in capturing these factors.

Various methods have been developed to address this challenge, such as statistical approaches based on Monte Carlo simulation [11, 12, 13]. In this family of techniques for determining the uncertainty, Monte Carlo simulations are used to approximate the distribution of the solution and assess the uncertainty of the result.

Although these methods have proven effective in the determination of the uncertainty, their application remains limited due to the computational load of the Monte Carlo simulations in dealing with high-resolution real data.

Recent research in the statistical imaging literature with an emphasis on the Bayesian framework [14, 15, 16] employs robust mathematical methods to ensure stability facing noisy data and support computationally feasible uncertainty quantification.

In this family of approaches, the underlying data observation process can be combined with prior knowledge, forming a realistic model of the scanning process while promoting the expected properties for the result, e.g., regularity or smoothness. In this methodology, the CT reconstruction problem can be solved via statistical inference techniques, such as maximum likelihood or Bayesian strategies [17]. Contrary to the analytical and iterative techniques, which provide a single gray value for each voxel, statistical methods yield a distribution for each voxel, which represents the distribution of the volume with respect to the observed data. In the Bayesian framework, this distribution is known as the posterior distribution, which combines the data observation process with the prior knowledge.

In [18], a joint posterior distribution is defined for both the reconstruction result and the geometrical parameters, and then suitable methods are used to analyze the distribution for determining the uncertainty of the result.

The approach presented in [19] introduces a model-discrepancy term into the CT model to account for uncertainties resulting from discrepancies between the assumed and actual geometric parameters.

The method outlined in [20] showed that by using the highest posterior density intervals, the uncertainty of specific features of the reconstruction result can be determined. For instance, quantifying the uncertainty in the estimated boundaries of inclusions within a CT image, which represents the most credible regions for the true boundary location, given the observed data and chosen prior distributions.

For addressing uncertainty quantification in X-ray CT reconstruction with uncertain view angles in [21], a comprehensive Bayesian model was defined that treats the unknown attenuation coefficients, the view angles, and the associated hyperparameters as random variables, determines the uncertainties in the reconstruction.

Despite the significant advancements in Bayesian methodologies for UQ in CT reconstruction, which address various sources of uncertainty in the results, there is still a need for further development in this methodology. A downside of these models is their computational demands, making them impractical for high-dimensional CT data without requiring data reduction. When the data size exceeds the computational capacity of these methods, we are forced either to switch to a less computationally demanding method or to subsample the data. This data reduction, however, introduces additional uncertainty into the analysis and complicates the analysis of results, potentially undermining the robustness of the uncertainty quantification process.

1.2. Contribution

In this manuscript, we focus on determining the uncertainty caused by noise in CT-based dimensional measurements using computationally efficient Bayesian techniques. We define two

posterior distribution models, each associated with a specific regularization method (Tikhonov and a smooth approximation of the total variation), that are used to estimate the gray-value uncertainty for each voxel via a computationally feasible sampling algorithm.

To this end, we explore two Bayesian inference approaches: the Maximum A Posteriori (MAP) estimate, obtained via convex optimization algorithms [22, 23], and the Markov Chain Monte Carlo (MCMC) [24, 25] sampling method, specifically the Metropolis-adjusted Langevin Algorithm with adaptive step size (LipMALA). This combination enables both efficient reconstruction and detailed posterior uncertainty analysis on real CT data.

We also demonstrate the propagation of this uncertainty to dimensional measurements, validating our approach through repeating the experiment on industrial parts.

The rest of the manuscript is organized as follows: in Section 2, we describe the problem formulation, focusing on the CT inverse problem and defining the Bayesian model for the reconstruction problem. In Section 3, the Bayesian framework will be presented in more detail, and a computational method to obtain a solution from the model will be discussed. In section 4, the details about the implementation of the discussed methods are presented, along with parameter selection and the data used in this manuscript, followed by the results obtained from applying Bayesian techniques on real CT data from two objects: an aluminum cylindrical step gauge and a 3D printed titanium part will be presented. This is followed by bi-directional length measurements on the reconstruction results, and then determining the related uncertainty in dimensional measurements performed from the reconstructed data. The paper is concluded in section 5, with suggestions for expanding this research.

2. Statistical formulation of the CT reconstruction problem

Image reconstruction is an inverse problem, a process of calculating the causal factors that produced a set of observations [17], which can be mathematically formulated as:

$$Ax = b, \tag{1}$$

where A is a linear operator, which in the context of CT is the forward projection model, x is the unknown image object, and b is the noisy observed data [17]. Therefore, by solving the CT inverse problem we aim to reconstruct the image object x, a 3D representation of the scanned object, from the projections b.

Inverse problems, especially CT image reconstruction, are generally ill-posed in the sense defined by Hadamard [26], meaning they are highly sensitive to input errors and are not guaranteed a unique solution. This makes them sensitive to errors and noise in the input data, mentioned in Table 1, posing a significant challenge for reliable image reconstruction.

Regularization techniques can be used to mitigate these issues by imposing prior knowledge on the solution, ensuring a meaningful and stable answer while dealing with noisy or limited data, which is critical in CT reconstruction.

In the Bayesian framework, an inverse problem can be formulated as a statistical inference problem, where the solution to the problem (1) is the distribution of the parameter x, given the observed data b. This distribution is called the posterior distribution. In Bayes' theorem, we can obtain the posterior probability density by combining the data and prior distributions:

$$\pi_{\text{post}}(x|b) \propto \pi_{\text{likelihood}}(b|x)\, \pi_{\text{prior}}(x), \tag{2}$$

where the likelihood function $\pi_{\text{likelihood}}(b|x)$ describes the conditional probability density for the data which was produced from scanning the object x, and $\pi_{\text{prior}}(x)$ encodes the prior information of the image object x before observing the data. Therefore, to form the Bayesian model for the CT reconstruction problem, a likelihood and prior distribution need to be defined.

2.1. Likelihood Function

The likelihood function $\pi_{\text{likelihood}}(b|x)$ models the relationship between the projections b and the image object x. The average behavior of the x-ray attenuation process is described by the Lambert-Beer law [27]

$$I = I_0 \exp^{-Ax},$$

where I_0 is the number of the x-ray photons that enter the object and I is the number of photons that reach the detector, hence $b = -\log(-I/I_0)$ The x-ray attenuation process in high-dose CT scans can be statistically modeled with a wide distribution, such as a Gaussian (normal) distribution, using the forward operator A, such that [1]:

$$b \sim \mathcal{N}(Ax, I), \tag{3}$$

where σ is the standard deviation of the data. In Section 4, we will discuss the estimation of σ from the observed data.

For lower I_0, low-dose scans, or photon-limited settings leading to narrower distributions for I, a Poisson distribution might be more appropriate, though they introduce additional complexity in inference.

With this simplification, the likelihood distribution becomes [28]:

$$\pi_{\text{likelihood}}(b|x) \propto \exp\left[\frac{-1}{2\sigma^2}\|Ax - b\|_2^2\right]. \tag{4}$$

2.2. Prior distribution

When solving an inverse problem, the prior distribution reflects the prior assumptions about the solution, which helps to regularize the reconstruction process in our problem, incorporating additional information or constraints about the expected parameter values and penalizing heavy oscillations in pixel values to force smoothness in the image. In this manuscript, two prior distributions will be used: a Gaussian prior and a combination of Gaussian with Laplacian prior, which respectively correspond to two regularization techniques: Tikhonov regularization [29] and the smooth total variation (sTV) regularization [30], which will be discussed in detail.

2.2.1. Gaussian prior: Tikhonov regularization

Gaussian prior enforces global smoothness on the result, suppressing the large noise components by controlling the norm of u. This prior information corresponds to the Tikhonov regularization and is defined as:

$$\pi_{\text{prior}}(x) \propto \exp\left[-\lambda\|Lx\|_2^2\right],$$

where λ controls the influence of the prior distribution on the reconstruction compared to the likelihood, and L is a linear operator, e.g., a finite difference approximation of the gradient operator.

In the context of Tikhonov regularization, λ is called the regularization constant, which controls the trade-off between fitting the data and applying the regularization penalty. If the operator $L = I$, we obtain the standard Tikhonov regularization, and if L is an operator, the generalized Tikhonov regularization is obtained.

Therefore, the posterior distribution that utilizes the Tikhonov regularization can be formulated as:

$$\pi_{\text{post}}(x|b) \propto \exp(-J_{\text{tikh}}(x,b)), \quad J_{\text{tikh}}(x,b) = \frac{1}{2\sigma^2}\|Ax-b\|_2^2 + \lambda\|Lx\|_2^2, \tag{5}$$

2.2.2. Hybrid Gaussian-Laplacian prior: Smooth TV regularization

By considering a Laplacian prior on the gradient of the image, the total variation (TV) regularization is derived. This distribution promotes sparsity in the gradient, which encourages edge preservation by minimizing the TV of the changes in the image intensity [30]. TV measures the rate of change between neighboring pixels, penalizing large changes (noise) while allowing for sharp transitions (edges). Unlike methods that use a squared norm as a penalty term, such as Gaussian priors, the Laplacian prior can tolerate sharp transitions, making it an excellent choice for preserving edges in images that include discontinuities.

The standard form of TV is not smooth, and while algorithms for non-smooth optimization do exist, they generally suffer from slow convergence. Therefore, a combination of Laplacian and Gaussian distributions is used on the gradient of the image object, yielding a smooth approximation of TV regularization. This approach is commonly achieved through the Huber approximation of TV norm [31] defined as follows:

$$\pi_{\text{prior}}(x) \propto \exp\left[-\lambda H_\epsilon(Lx)\right],$$

with

$$H_\epsilon(Lx) = \sum_{i,j,k} h_\epsilon(\|(Lx)_{i,j,k}\|_2),$$

where ϵ is the smoothing constant, $(Lx)_{i,j,k}$ is a finite difference approximation to the gradient at pixel i,j,k and the Huber function $h_\epsilon(t)$ is defined as

$$h_\epsilon(t) = \begin{cases} \frac{t^2}{2\epsilon} & \text{if } t \le \epsilon, \\ |t| - \epsilon & \text{else.} \end{cases}$$

Therefore, the posterior distribution (2) for the smooth TV (sTV) will be of the form:

$$\pi_{\text{post}}(x|b) \propto \exp(-J_{\text{stv}}(x,b)), \quad J_{\text{stv}}(x,b) = \frac{1}{2\sigma^2}||Ax-b||_2^2 + \lambda H_\epsilon(Lx). \tag{6}$$

With this definition of $J(x,b)$ in (5) and (6), two log-concave posterior distributions $\pi_{\text{post}}(x|b)$ are formed, and suitable Bayesian methods used on this distribution are discussed in detail in the next section.

3. CT reconstruction and uncertainty: Algorithmic approaches

As mentioned in the introduction, there are two main approaches within the Bayesian methodology. While we first determine the MAP estimate using convex optimization techniques for an initial analysis of the reconstruction result, our primary focus in this work is on the MCMC methods, which provide UQ for the MAP estimate and are subsequently introduced.

3.1. MAP estimate

The Maximum A Posteriori, or MAP estimator, is a point x which maximizes the posterior probability distribution function π_{post}:

$$x_{\text{MAP}} = \text{argmax}_x \pi_{\text{post}}\left(x|b\right).$$

In section 2, two posterior distributions were introduced in (5) and (6). Here, we will discuss the optimization algorithms used to obtain x_{tikh} and x_{stv} in detail.

3.1.1. MAP with Tikhonov regularization: conjugate gradient descent

The MAP estimate x_{MAP} computed related to (5) will be denoted by x_{tikh}, and is calculated from solving the following minimization problem:

$$x_{\text{tikh}} = \min_x J_{\text{tikh}}(x, b). \tag{7}$$

The gradient of the objective function $J_{tikh}(x, b)$ is

$$\nabla J_{\text{tikh}}(x, b) = \frac{1}{\sigma^2} A^*(Ax - b) + 2\lambda L^* Lx, \tag{8}$$

where we denoted by M^* adjoint of a bounded linear operator M.

It can be simply verified that $\nabla J_{tikh}(x, b)$ is Lipschitz-continuous with parameter $lip_{\text{tikh}} = \|A\|_2/\sigma^2 + \lambda\|L\|^2$, where $\|L\|^2 \leq 12$ is the operator norm of L [32], and the norm of the operator A is the largest singular value of A which can be computed using the power method [33].

In solving the Tikhonov regularization problem, gradient-based methods, such as the Conjugate gradient (CG), could be used [22]. The first order optimality condition to solve (8) requires the gradient to be zero, then we have

$$(\frac{1}{\sigma^2} A^*A + \lambda L^*L)x = \frac{1}{\sigma^2} A^*b, \tag{9}$$

which can be rewritten as $Rx = f$ with $R = (\frac{1}{\sigma^2} A^*A + \lambda L^*L)$ and $f = \frac{1}{\sigma^2} A^*b$. For computing the regularized reconstruction result from (9), we are using the CG method, which is suitable for solving large-scale linear systems of the form $Rx = f$. CG iteratively updates the estimate of u by minimizing the quadratic objective function $J_{\text{tikh}}(x, b)$ from (5), at each iteration aiming to reduce the residual error $r_k = b - (A + \lambda L^*L)x_k$.

3.1.2. MAP with smooth TV regularization: accelerated gradient descent

Computing the MAP estimate form (6), denoted by x_{stv}, is similar to the Tikhonov problem in minimizing the corresponding objective function:

$$x_{\text{stv}} = \min_x J_{\text{stv}}(x, b) \tag{10}$$

and the gradient of the Huber approximation of the TV regularization is

$$\nabla J_{\text{stv}}(x, b) = \frac{1}{\sigma^2} A^*(Ax - b) + \lambda L^* \nabla H_\epsilon(Lx), \tag{11}$$

where

$$(\nabla H_\epsilon(Lx)) = \frac{\lambda(Lx)}{\max\{\epsilon, |(Lx)|_2\}}. \tag{12}$$

The function $J_{\text{stv}}(x, b)$ is continuously differentiable, and its gradient is Lipschitz-continuous with Lipschitz constant $lip_{\text{stv}} = \|A\|_2/\sigma^2 + (\lambda\|L\|^2)/\epsilon$ [34]. This property ensures that all gradient-based algorithms (e.g. optimization or sampling algorithms) are well defined [35]. The most commonly used methods for this model are Accelerated gradient descent (AGD), quasi-Newton method (l-BFGS) and gradient descent (GD). AGD is competitive with BFGS and simpler to implement in this problem [34], therefore the AGD method with fixed step size depending on the Lipschitz constant[36] is used to solve the problem (10). The outline of this method is presented in Algorithm 1.

Algorithm 1 Accelerated gradient descent (AGD) with fixed step.

Initialize: $x_0 = x_{-1} = y_0$ (initial guess), $\tau \leq 1/lip_{\text{stv}}$, and $t_0 = 1$
for $k = 1, ..., N$ **do**
  Compute $t_{k+1} = \frac{1+\sqrt{1+4t_k^2}}{2}$
  Compute $y_k = x_k + \frac{t_k - 1}{t_{k+1}}(x_k - x_{k-1})$
  Update $x_{k+1} = y_k - \tau \nabla J_{stv}(y_k, b)$
  $k = k + 1$
end for
Output: reconstruction result x

3.2. MCMC sampling

As was briefly mentioned in the previous section, the MCMC class of sampling methods within the Bayesian inference can be computationally demanding in cases like CT data, due to the large dimensions of the data, despite being highly effective in estimating complex distributions such as the posterior distributions defined in the previous section. Since we are interested in an in-depth investigation of the posterior distribution, a sampling method that rapidly converges to the solution is needed, which reduces the computation demand and storage need.

The Langevin dynamic family of sampling algorithms [37] offers numerous advantages over other methods, which makes them appealing for our application. Among these advantages, we can mention the incorporation of gradient information which prevents random walk behavior and guides the proposal steps toward regions of higher probability [38]. This property reduces the number of samples needed for approximating the distribution.

Proposing samples from the target distribution $\pi(x)$, which in our case is $\pi_{\text{post}}(x|b)$, with this family of methods is as follows:

$$x' = x_t - \tau\Sigma\nabla \log \pi_{\text{post}}(x_t \mid b) + \sqrt{2\tau}\Sigma^{1/2}\zeta_t, \tag{13}$$

where $\zeta_t \sim \mathcal{N}(0, I_{d\times d})$ is a vector of d-dimensional standard Gaussian random variables; Σ is a symmetric positive definite matrix, and τ is a step size which can be fixed or varied for all steps.

Adjusting the proposed samples in each step to filter them for our desired distribution with the Metropolis-Hastings acceptance/rejection step, results in the MALA sampling algorithm [39]. The accept–reject probability for the proposed sample x' can be calculated by:

$$\gamma(x'|x_t) = \min(1, \frac{\pi_{\text{post}}(x' \mid b)\mathcal{N}(x' \mid x_t - \tau\Sigma\nabla \log \pi_{\text{post}}(x_t \mid b), 2\tau\Sigma)}{\pi_{\text{post}}(x_t \mid b)\mathcal{N}(x_t|x' - \tau\Sigma\nabla \log \pi_{\text{post}}(x' \mid b), 2\tau\Sigma)}).$$

After calculation of $\gamma(x' \mid x_t)$, a random number $r \sim \mathcal{U}(0,1)$ is drawn and in case of $r < \gamma(x' \mid x_t)$, the proposed sample x' can be accepted to be stored, $x_{t+1} = x'$, and the next sample be proposed based on it. In case of rejection, the proposed sample will be discarded and another sample will be proposed based on the previous step.

This filtering of the samples reduces the storage need for storing all the proposed samples and will lead to a more precise approximation of the posterior distribution.

Although the MALA sampling technique offers various properties that make it suitable for CT reconstruction, the performance of this approach relies on selecting the appropriate step size, τ, which is a challenge in dealing with high-dimensional data. The step size governs the extent to which the drift of the Langevin dynamics can change. Step size τ can be tuned such that the method achieves better performance; however, this is non-trivial and problem-dependent. To have a rapid convergence to the posterior distribution and the desired answer, we need to tune the step size based on the problem and the data.

Therefore, variants of the MALA method with adaptive step size are more appealing in CT reconstruction. We will be using the method proposed in [40], which offers an adaptive step size mechanism that adjusts the step size during the sampling process based on local Lipschitz constants L_J and the gradient of the target distribution:

$$\tau_t = min\{(1+\alpha_{t-1})\tau_{t-1}, L_J \frac{||x_t - x_{t-1}||}{||\Sigma\nabla \log \pi_{\text{post}}(x_t \mid b) - \Sigma\nabla \log \pi_{\text{post}}(x_{t-1} \mid b)||}\},$$

where α_t represents the ratio between two consecutive step sizes, and L_J is the inverse Lipschitz constant of the log-posterior distribution, corresponding to lip_{tikh} and lip_{stv}, depending on the chosen posterior model.

Lipschitz-MALA (Lip-MALA) effectively mitigates the need for manual tuning, by adapting the step size which enhances the robustness and convergence properties of the algorithm, making it a promising choice for reconstruction and uncertainty quantification in CT applications, and to the best of our knowledge, this method has not been tested on CT reconstruction applications.

Algorithm 2 describes the pseudocode for implementing the LipMALA method and shows the algorithmic steps involved in its execution. The proposal of new samples in the LipMALA method is based on the gradient of the log of the posterior distribution, similar to the classical MALA. Although this method has an additional step compared to MALA, it requires no additional computation by reusing the values that have been calculated in previous steps. Therefore, there is no additional complexity or computation cost to the algorithm.

Algorithm 2 Metropolis-adjusted Langevin algorithm with locally Lipschitz adaptive step size (Lip-MALA MCMC)

Require: $x_0, \tau_0 > 0, \alpha_0 = +\infty, L_C = 1/lip_J, \Sigma = P_{n\times n}$
Ensure: $x_0, ..., x_N$
  for $t = 1, ..., N$ **do**
    Draw vector $\zeta_t \sim \mathcal{N}(0, I_{n\times n})$
    Propose a sample $x' = x_t - \tau_t \Sigma \nabla \log \pi_{\text{post}}(x_t \mid b) + \sqrt{2\tau}\Sigma^{1/2}\zeta_t$
    Calculate the acceptance ratio $\gamma = \frac{\pi_{\text{post}}(x'|b)}{\pi_{\text{post}}(x'|b)} \cdot \frac{q(x_t|x')}{q(x'|x_t)}$,
    Generate a random number $u \sim \mathcal{U}(0, 1)$
    if $u < \gamma$ **then**
      Accept the candidate state: $x_{t+1} = x'$
      Update the step size: $\tau_t = min\{(1 + \alpha_{t-1})\tau_{t-1}, L_J \frac{||x_{t+1} - x_t||}{||\Sigma\nabla \log \pi_{\text{post}}(x_{t+1}|b) - \Sigma\nabla \log \pi_{\text{post}}(x_t|b)||}\}$
      Update $\alpha_t = \tau_t/\tau_{t-1}$
    else
      Reject the proposed sample: $x_{t+1} = x_t$
    end if
  end for

The MCMC CT reconstruction process based on the LipMALA method is shown in Figure 1, which takes the scanned projections as an input, and the output is a volumetric model x, where each voxel $x_{i,j,k}$ within this volume is assigned a distribution of gray values $\pi_{post}(x_{i,j,k} \mid b)$.

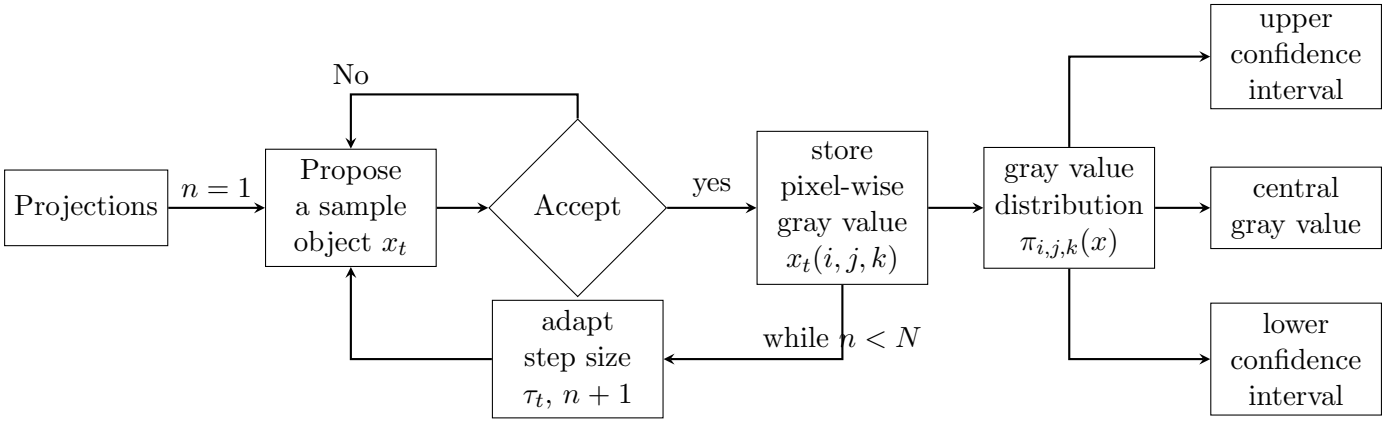

Figure 1: Illustration of LipMALA sampling method, incorporating adaptive step size adjustment upon acceptance of proposed samples. Upon completion of the sampling process, the mean and variance are calculated, with the mean representing the constructed volume and the voxel-wise uncertainty.

With the definitions of the two forms of the posterior distribution for Tikhonov and sTV regularization, and with both being Lipschitz continuous, we can use the MALA sampling method for accruing an approximation of the posterior distributions (5) and (6), which will give us adequate information to compute a lower and upper confidence interval for the MAP estimates x_{tikh} and x_{stv}.

In section 2, two methods were defined and in section 3 two optimization algorithms and one sampling technique were introduced. In the next section, the results of applying these methods on

real CT data will be presented and compared.

4. Numerical results

In this section, we outline the implementation details of the methods described in the previous sections, including the setup for regularization, optimization algorithms, and uncertainty quantification. We begin by a description of the data acquisition process and the reconstruction algorithms employed, followed by discussing the parameter selection for both the Tikhonov and sTV regularization methods. Additionally, we present the numerical results from applying the discussed methods to real CT data for reconstruction, followed by uncertainty determination around the result and dimensional measurements on the parts.

4.1. Data

The data used in this manuscript are real XCT projections. Scans are acquired using a Nikon XTH 225 ST computed tomography system. This system is equipped with a 2000 by 2000 X-ray detector with a pixel pitch of $0.2mm$.

We applied the methods on two objects: a cylindrical step gauge made out of aluminum [41] and a 3D printed Titanium manifold part, presented in Figure 2. The step gauge block is a reference object for dimensional measurements that is suitable for validation of the results and the manifold is a representation of a real-life application of the method.

The lab CT system employed a cone-beam setup with $N = 360$ equidistant projections for the step gauge block acquired over $360°$ and $N = 3600$ projections for the printed manifold, which was then subsampled to $N = 360$ projections to create a moderate few-view data. The source-to-center distance and source-to-detector distance were measured as 307.8mm and 1091.9mm, respectively, for the step gauge block and 129.4mm and 1091.9mm for the manifold.

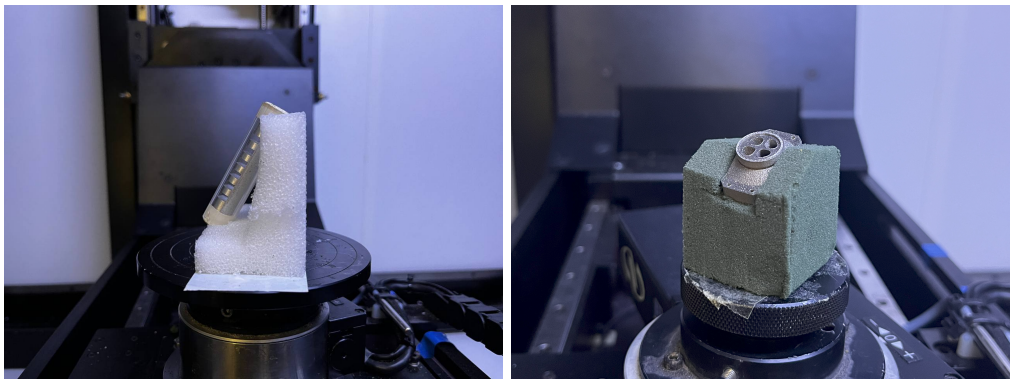

Figure 2: The aluminum cylindrical step gauge on the left, the Titanium 3D printed manifold on the right

A procedure described in [42] is used to determine the precise geometry of the XCT system. This method involves scanning a reference object containing high X-ray absorption spheres, extracting their projected centers in 2D detector images, and optimizing a ray-tracing model to fit the observed data. Since a highly precise estimation of the geometrical parameters is not required for our proposed method, alternative techniques, such as the approach presented in [43], can be employed for parameter estimation.

We use the ASTRA toolbox [44] and Tomosipo [45] to perform forward and backward projections, as well as the FDK reconstruction.

	Tikhonov		sTV
	standard	generalise	
Step gauge	0.19	0.12	0.009
Manifold	0.51	0.43	0.011

Table 2: regularization values for each method and object

4.1.1. Ground truth

To create a ground truth for a quantitative comparison between the results acquired from different methods, we are using a high-resolution scan, denoted as $\hat{b}$. The high-resolution scan for the step gauge block was acquired under identical settings to the data, with the sole difference being the number of projections. The scan was done with $N = 3600$ projections. The ground truth for the printed manifold is the hight resolution original scan.

4.2. Registration

To accurately compare the reconstructed results to each other and to the ground truth, it was essential to perform image registration for the results acquired from the step gauge block. Given that slight misalignment between the reconstructions can occur, subpixel image translation registration was employed to ensure precise alignment [46]. Translation offset and rotation differences between pairs of reconstruction results were calculated using the Python library scikit-image [47]. This method first estimates the cross-correlation peak using the Fast Fourier Transform, then refines the shift estimation by upsampling the Discrete Fourier Transform in a small neighborhood of the initial estimate through a matrix-multiplication approach.

This process was not necessary for the comparisons between the results from the printed manifold, as the ground truth and the data is from the same scan.

4.3. Parameters

4.3.1. Noise level

The standard deviation of the likelihood distribution (4), σ, is determined using the Python function restoration.estimate_sigma from the library Skimage applied on the data b. This function assumes that the noise follows a Gaussian distribution, and the estimation algorithm is based on the median absolute deviation of the wavelet detail coefficients, as described in Section 4.2 of [48].

4.3.2. Regularization parameter

As mentioned in section 2, the regularization parameter λ controls the balance between the regularization and the data fidelity term. The best parameter will reduce the effect of the noise in the data while keeping the valuable features in the image. In this paper, we will use the λ that minimizes the normalized mean squared error (NMSE) computed by $\|x_\lambda - \hat{x}\|_2^2 / \|x_\lambda\|_2^2$ between the regularized solution x_λ (which in this paper we note as x_{tikh}, x_{stv}) and the ground truth $\hat{x}$. The ground truth $\hat{x}$ is the FDK reconstruction of the high-resolution scan $\hat{b}$.

Using a grid search over a range of λ values, we selected those that minimized the NMSE between the reconstruction and the ground truth scan. Through this approach, the best-performing λ values for the step gauge block were found which are shown in Table 2.

4.3.3. Smoothing parameter

Given the problem specification, the Huber smoothing parameter that is used for our data is set at $\epsilon = 10^{-5}$ for the step gauge block and $\epsilon = 10^{-7}$ for the manifold. A smaller value results in less smoothing, which provides a closer approximation to the standard TV regularization problem ($\epsilon = 0$). However, reducing ϵ also slows the algorithm's convergence [23]. For a good approximation to the TV regularization, it is sufficient for ϵ to be small compared to the smallest difference between neighboring pixel values, which in this study was found to be around 10^{-4} to 10^{-3} for the step gauge and 10^{-6} to 10^{-7} for the manifold. Additionally, we evaluated the effect of ϵ by comparing reconstructions at $\epsilon \in \{10^{-2}, 10^{-3}, 10^{-4}, ..., 10^{-10}\}$. For $\epsilon \leq 10^{-5}$ and $\epsilon \leq 10^{-7}$, no substantial differences were observed for the step gauge block and the manifold, respectively.

4.4. Reconstruction results

As the first step, it is essential to establish a ground truth or benchmark against which to compare the reconstruction results. For this purpose, we present the FDK reconstruction of the moderate few-view data b, and also the ground truth $\hat{x}$. The reconstruction results for both parts are shown in Figure 3, and the high-resolution scan will serve as a reference point for evaluating the accuracy of all the other reconstruction methods.

The amount of projection data available for reconstruction indicates the difference between the reconstruction results of the high-resolution scan $\hat{b}$ and the moderate few-view scan b. The high-resolution scan provides much denser data, capturing more angular information and object detail, resulting in a more accurate, higher-quality reconstruction. The lower-resolution scan, on the other hand, has limited angular sampling, artifacts, and reduced resolution. This disparity in the projection data leads to a noticeable difference in the reconstruction quality.

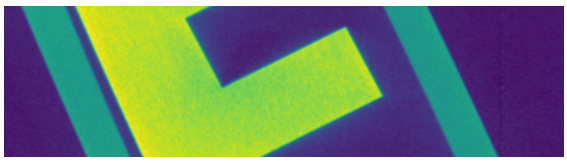

(a) Ground truth: step gauge block, high-resolution

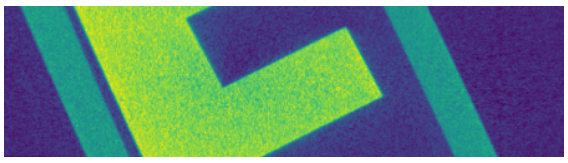

(b) FDK reconstruction: step gauge block, 360 projection

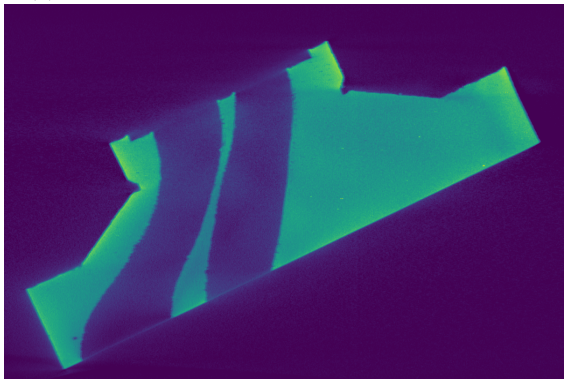

(c) Ground truth: manifold, high-resolution

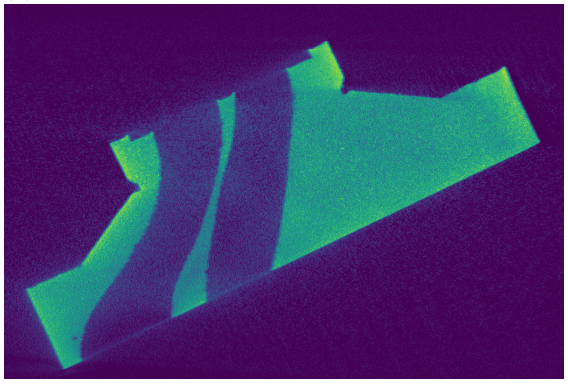

(d) FDK reconstruction: manifold, 360 projection

Figure 3: Reconstruction results with FDK for the two scanned objects, high-resolution and moderate few-view

It should be noted that each time the NMSE is used for comparing a reconstruction result with the ground truth, we perform image registration to account for any translation and rotation

differences between the two scans. By adding this step we ensure that the reconstruction and the ground truth are aligned properly, allowing for an accurate comparison. By fixing any misalignment, we avoid artificial errors in the NMSE calculation and ensure that the differences reflect only the reconstruction quality.

As was mentioned in section 2, the generalized Tikhonov regularization outperforms the standard Tikhonov regularization, as demonstrated in Figures 4a, 4b, 5a and 5b. Additionally, the NMSE comparison between the standard and generalized Tikhonov methods in Table 3 and Table 4, relative to the ground truth, further highlights the superior performance of the generalized form.

The performance of the sTV regularization method surpasses that of generalized Tikhonov in terms of reconstruction quality, as illustrated in Figure 4c and 5c and the comparison of the gray value profiles in Figure 4d, and 5d. This is further supported by the NMSE values for each reconstruction, calculated with respect to the ground truth and presented in Table 3 and Table 4. The superior performance of sTV in preserving fine details in the object is highlighted by preserving an artifact in the manifold, as shown in Figure 5e and 5f.

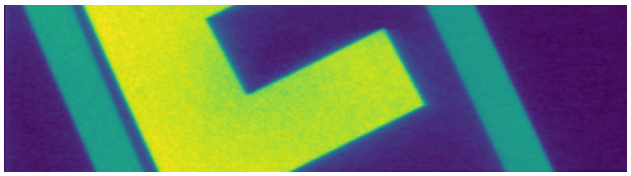

(a) Conjugate Gradient, standard Tihonov

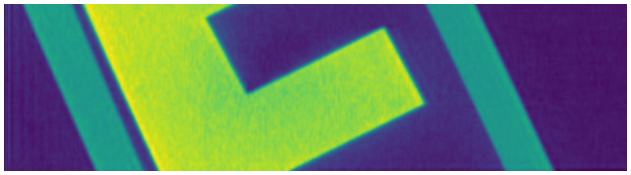

(b) Conjugate gradient, generalized Tikhonov

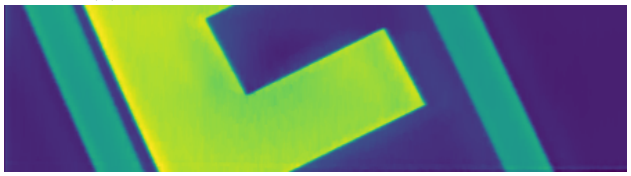

(c) Accelerated Gradient Descent, smooth TV

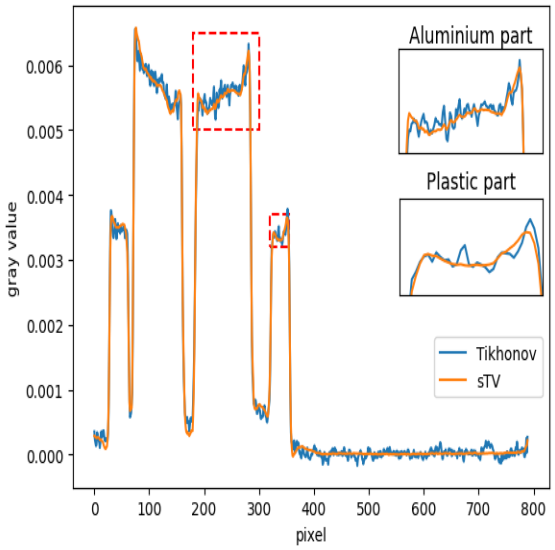

(d) profile comparison between Tikhonov and sTV

Figure 4: Reconstruction results from different regularization methods on the step gauge block

Table 3: NMSE between reconstruction results and the ground truth for the step gauge block

	FDK	CG	AGD	MALA	LipMALA
Unregularized	0.0219	-	-	-	-
Standard Tikhonov	-	0.0155	-	-	-
Generalized Tikhonov	-	0.0085	-	0.0096	0.0089
Smooth TV	-	-	**0.0050**	0.0077	**0.0054**

Although the sTV method achieves better reconstruction results, the convergence speed of the AGD algorithm for solving the sTV regularization is slower than the CG method used for solving

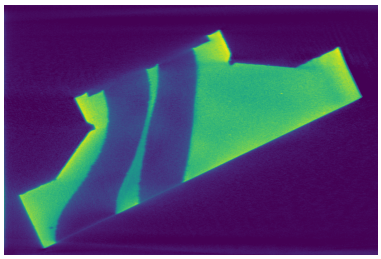

(a) CG, Standard Tihonov

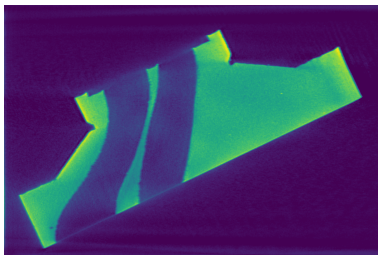

(b) CG, Generalized Tikhonov

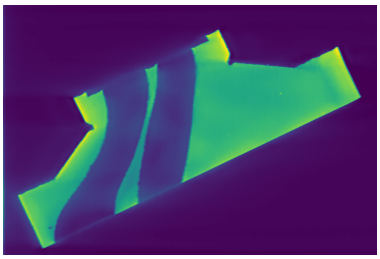

(c) AGD, Smooth TV

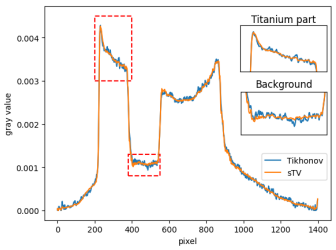

(d) profile comparison between Tikhonov and sTV

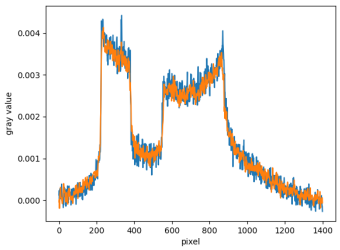

(e) profile comparison between Tikhonov and ground truth in preserving the artifact

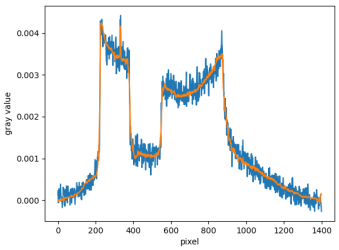

(f) profile comparison between sTV and ground truth in preserving artifact

Figure 5: Reconstruction results from different regularization methods on the 3D printed titanium manifold

Table 4: NMSE between reconstruction results and the ground truth for the manifold

	FDK	CG	AGD	MALA	LipMALA
Unregularized	0.0286	-	-	-	-
Standard Tikhonov	-	0.0128	-	-	-
Generalized Tikhonov	-	0.0101	-	0.0108	0.0104
Smooth TV	-	-	**0.0092**	0.0096	**0.0094**

the Tikhonov regularization, as illustrated in the residual logplot Figure 6, which compares the rate at which the residuals of the two methods decrease for each object. This difference is due to the fact that Tikhonov regularization is a quadratic form, while sTV is a smooth approximation of a non-differentiable term. As we can see in Figure 6, CG converges faster than AGD, with the residuals approaching zero more rapidly. This is consistent with the known behavior of CG for quadratic problems, where it typically converges in fewer iterations compared to gradient-based methods like AGD. However, despite the slower convergence of AGD, the regularized sTV method still yields superior reconstruction results due to its ability to preserve edges.

Following this, we will delve deeper into the uncertainty determination for each result, examining the voxel-wise distributions generated from the LipMALA sampling method. This analysis will provide insights into the reliability of the reconstructed images by quantifying the uncertainty associated with both the Tikhonov and sTV reconstructions.

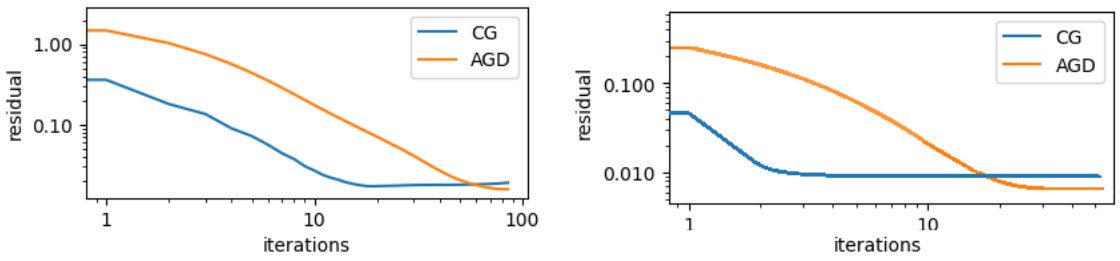

Figure 6: Residual convergence for the AGD and CG method, left: step gauge block, right: manifold

4.5. Bayesian Uncertainty Quantification

For the MALA and LipMALA methods, we performed $N = 1000$ and $N = 500$ sampling iterations respectively, for both posterior distributions (5) and (6) for both objects, discarding the first half as the burn-in. We observed that the MALA method required more samples for estimating the posterior distribution. Hence, the number of samples is higher than LipMALA. Each sampling process was initialized from the MAP estimate obtained from (7) and (10). The decision to warm-start the sampling from the MAP estimate was made to provide the algorithm with a favorable starting point, thus minimizing the time required for the samples to converge to the expected value, as well as reducing the storage space needed for keeping the samples. For the MALA method, the step size is fixed on the Lipschitz constant $\tau_{\text{tikh}} = lip_{\text{tikh}}$ and $\tau_{\text{stv}} = lip_{\text{stv}}$, for Tikhonov and sTV regularization, respectively.

Following the sampling process, the mean μ_c and standard deviation u of the samples were computed for each voxel, which we then determined a triplet of gray values: one for the central value μ_C and one each for the upper μ_{UCL}, and lower bound of the confidence interval μ_{LCL}.

In section 2, the posterior distributions (5) and (6) were defined which both form a Gaussian distribution. Therefore, the central value μ_C can be obtained by the mean gray value of the samples drawn with the LipMALA method. Lower and upper confidence limits are given by $\mu_{UCL} = \mu_C + k \cdot u$ and $\mu_{LCL} = \mu_C - k \cdot u$, where k is the coverage factor and u is the standard deviation of the samples. Different values for k can offer different confidence limits, for instance, $k = 2$ provides a 0.9545 confidence limit. The mean is an approximation of the MAP estimate x_{MAP}, which is the reconstruction result from the sampling method, as discussed previously.

The mean of the samples for both distributions (5) and (6) representing the Tikhonov and sTV regularization, respectively, are compared in Table 3 and Table 4, and the corresponding reconstruction results are shown in Figure 7. Lower NMSE values indicate that the reconstruction result is closer to the ground truth, reflecting a more accurate reconstruction from the LipMALA method compared to the MALA algorithm. Therefore, we will focus on the results from the LipMALA method in the rest of this section.

In Figure 8 and 9, we plot the μ_{LCL}, central value μ_C, and μ_{UCL} along two profile lines located at the edges of the parts. These profile lines highlight the boundaries, allowing us to assess the uncertainty in edge preservation and the accuracy of the reconstruction, particularly in regions where sharp transitions are expected.

4.6. Dimensional measurements

Dimensional measurements are performed on the mean of the LipMALA samples and the upper and lower confidence intervals for each object for both the Tikhonov and smooth TV methods,

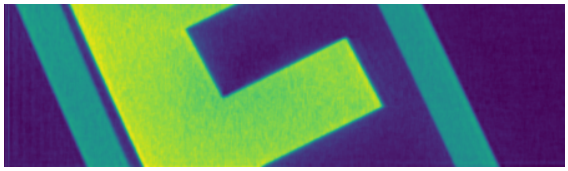

(a) μ_C, Tikhonov regularization, step gauge block

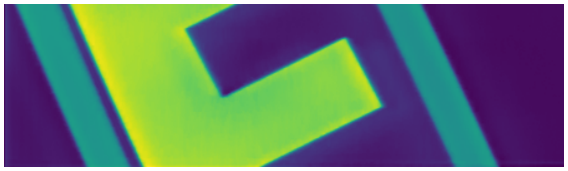

(b) μ_C, sTV regularization, step gauge block

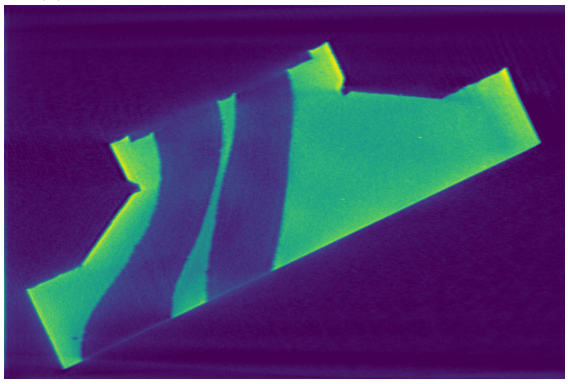

(c) μ_C, Tikhonov regularization, manifold

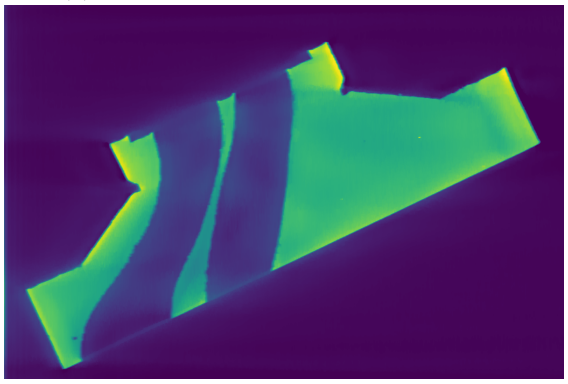

(d) μ_C, sTV regularization, manifold

Figure 7: Mean of the LipMALA samples for the posterior distribution representing the Tikhonov and sTV regularization

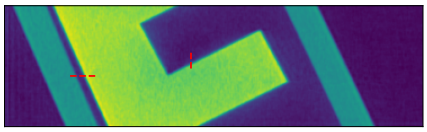

(a) mean LipMALA

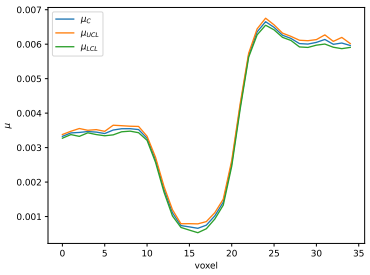

(b) Horizontal edge mark

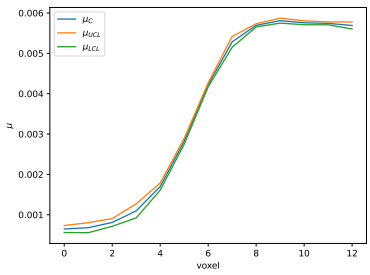

(c) Vertical edge mark

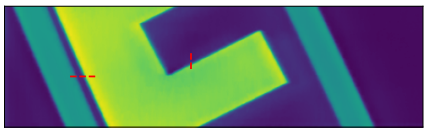

(d) Mean LipMALA

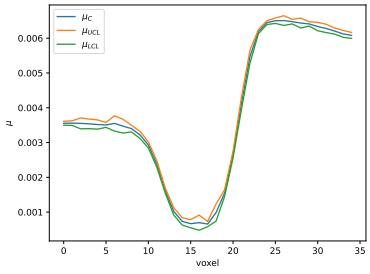

(e) Horizontal edge mark

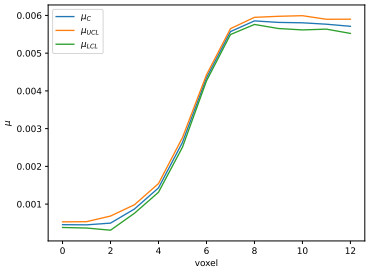

(f) Vertical edge mark

Figure 8: Gray value profile lines across edges, mean value, upper and lower confidence limits. top row: Tikhonov regularization, Bottom row: sTV regularization

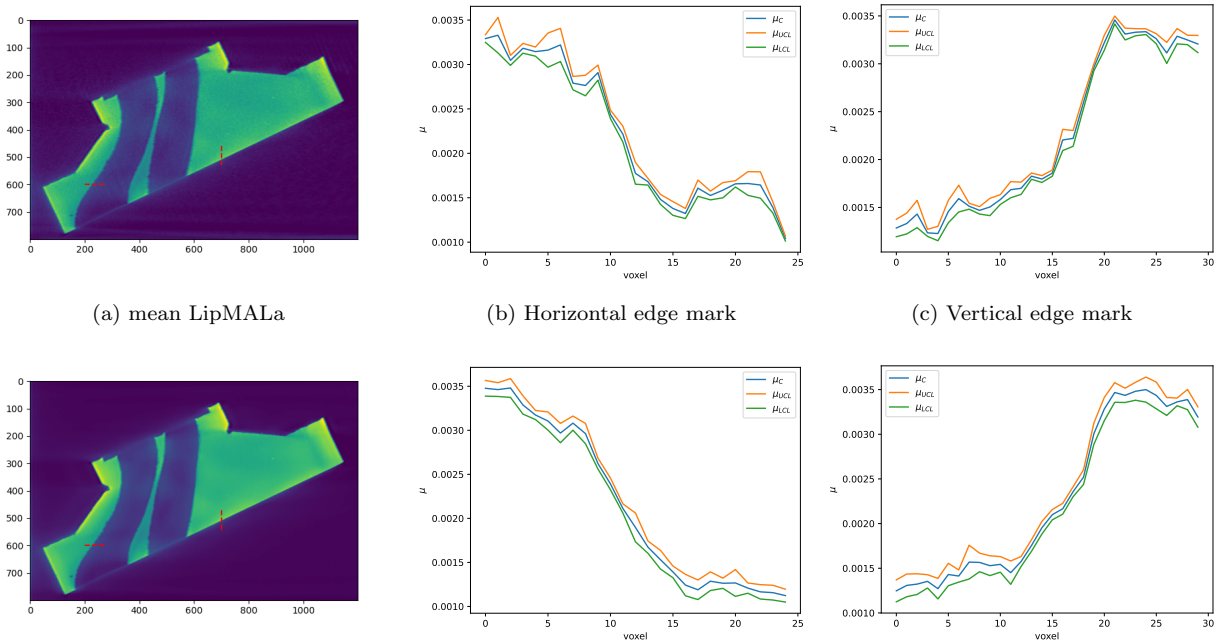

(a) mean LipMALa (b) Horizontal edge mark (c) Vertical edge mark

(d) Mean LipMALA (e) Horizontal edge mark (f) Vertical edge mark

Figure 9: Gray value profile lines across edges, mean value, upper and lower confidence limits. Top row: Tikhonov regularization, bottom row: sTV regularization

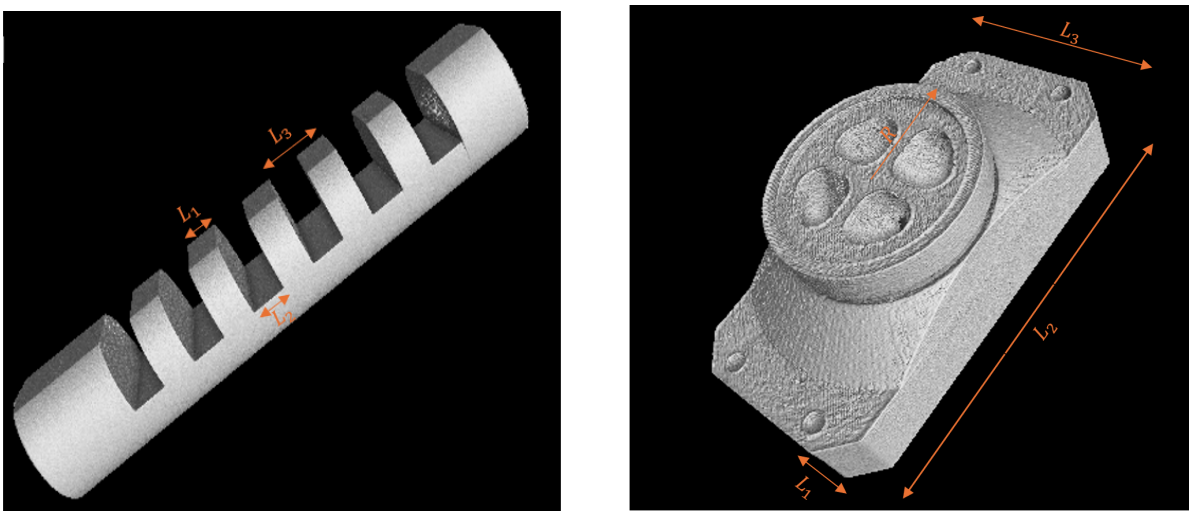

(a) 3D representation of the step gauge block

(b) 3D representation of the manifold

Figure 10: VGStuido max representation of the FDK reconstruction for both objects and the measured lengths

resulting in confidence limits for the dimensional measurements.

Each volumetric reconstruction model was imported into VGStudio MAX 2024.2.1 (Volume Graphics GmbH, Germany), and the segmentation was done on the volume with the advanced (classic) surface determination tool with automatic material definition. The measurement was repeated 3 times for each reconstruction and the average value of the three measurement results is reported. The 3D reconstruction result from the FDK method rendered in VGStudio is shown in Figure 10 for both objects, with indications of the lengths that are measured.

The dimensional measurements of the reconstructed parts are presented in Tables 5 and 6 for the mean, upper, and lower confidence intervals for the measurements on the results from the LipMALA sampling method for both the generalized Tikhonov and sTV regularization.

Table 5: Dimensional measurements on the mean, upper, and lower confidence limits of the step gauge block

	Ground Truth	Tikhonov			sTV		
		u_{LCL}	u_C	u_{UCL}	u_{LCL}	u_C	u_{UCL}
L_1(mm)	3.538	3.538	3.540	3.542	3.536	3.538	3.540
L_2(mm)	3.472	3.470	3.471	3.472	3.469	3.471	3.473
L_3(mm)	7.010	7.009	7.010	7.011	7.010	7.011	7.012

4.6.1. Validation of the UQ

In the previous section, we mentioned that using the coverage factor $k = 2$ provides a 95% confidence interval around the mean of the distribution. This means that if the experiment is repeated, 95% of the results are expected to fall within this interval.

We conducted an experiment by repeating the CT scan for both parts. The step gauge block was scanned 50 times under identical settings, the manifold was scanned twice with $N = 3600$ projections, and 20 subsamples with $N = 360$ projections were created. Each scan was then

Table 6: Dimensional measurements on the mean, upper, and lower confidence limits of the manifold

	Ground Truth	Tikhonov			sTV		
		u_{LCL}	u_C	u_{UCL}	u_{LCL}	u_C	u_{UCL}
L_1(mm)	4.038	3.907	3.990	4.074	3.899	3.984	4.069
L_2(mm)	26.492	26.487	26.525	26.563	26.484	26.524	26.564
L_3(mm)	13.106	12.997	13.083	13.169	12.961	13.050	13.139
R (mm)	6.527	6.501	6.553	6.605	6.515	6.573	6.631

reconstructed using both the AGD and CG methods, resulting in 50 reconstructions for each regularization method for the step gauge block and 20 reconstructions for each method for the manifold.

Using the standard deviation of the samples from LipMALA, we calculated the 95% confidence interval around the reconstruction. Therefore, we expect that 95% of the reconstructed results from the experiments lie within the confidence interval provided by the LipMALA method for both Tikhonov and sTV regularization of the corresponding object.

Since the UQ is voxel-wise, we assessed the percentage of the voxels from the 50 reconstructions inside the scanned object and excluded the background that fell within the 95% confidence interval generated by the LipMALA method for each method. The result showed that 86% and 87% of the voxels of the Tikhonv and sTV regularization for the step gauge and 84% and 85% for the manifold fell within this range, respectively, which is slightly lower than the expected 95% for both methods.

To further visualize the UQ results, we plotted the confidence intervals around the mean of the LipMALA samples μ_C for both the Tikhonov and sTV methods for both parts, focusing on a subset of voxels in the object, as well as the voxel values of the reconstruction of the scans for the same voxel.

For the step gauge block, we analyzed the aluminum part and the plastic compartment separately. Five voxels from each component were randomly chosen. Therefore, there are five confidence intervals for the aluminum part and five for the plastic compartment around it, and the intervals from the LipMALA method are plotted in red. The same voxels were chosen from the 50 scans as well, meaning there are 50 values for the same voxels plotted in green dots. These confidence interval plots are presented in Figure 11a, 11b for Tikhonov regularization and Figure 11c and 11d presenting the results from the sTV method, providing a clearer insight into the uncertainty across different regions of the object.

For the manifold, five voxels from the object were randomly chosen and, with the same principle as the step gauge block, are presented in Figure 13d and 13f for the generalized Tikhonov and sTV, respectively.

To achieve additional insight into the voxel-wise UQ attained by the LipMALA method, Figure 13 presents the chains of the samples after discarding the burn-in samples for selected voxels within the reconstructed volume for both posterior distributions (5) and (6). For the step gauge block, two voxels were chosen: one from the plastic compartment and one from the aluminum apart, while for the manifold, a single voxel within the object was selected. The plots show the sample trajectories over the iterations of the MCMC process. The dashed horizontal lines indicate the 95% credibility intervals based on the empirical standard deviation of the samples for each voxel.

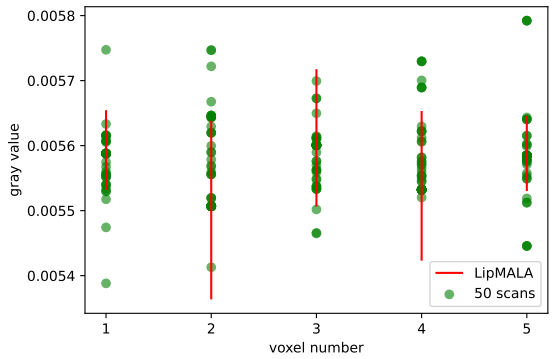

(a) Aluminium part

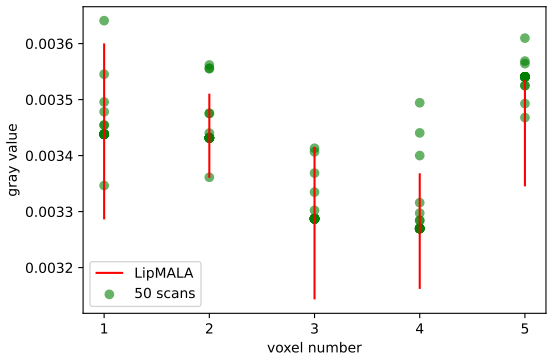

(b) plastic compartment

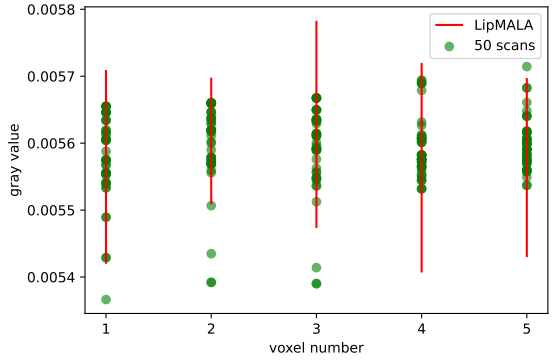

(c) Aluminium part

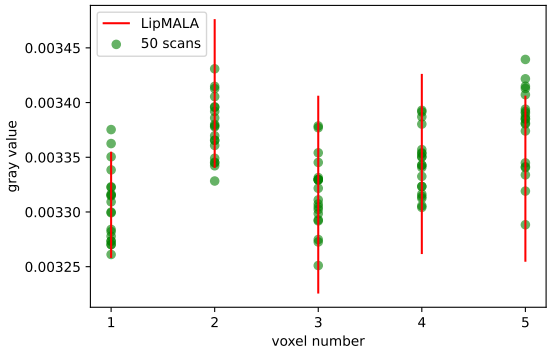

(d) plastic compartment

Figure 11: Confidence limits around μ_C of 5 random voxels in the object, in red, and the MAP estimate x_{MAP} of the 50 scans, green dots, of two different regions of the step gauge block. top row: Tikhonov regularization; bottom row: sTV regularization

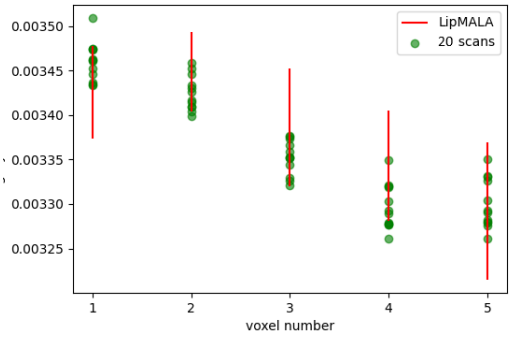

(a) Tikhonov regularization

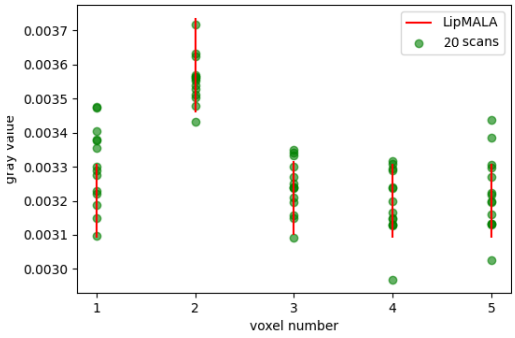

(b) sTV regularization

Figure 12: Confidence limits around μ_C of 5 random voxels in the object, in red, and the MAP estimate x_{MAP} of the 20 scans, green dots, of two different regions of the manifold.

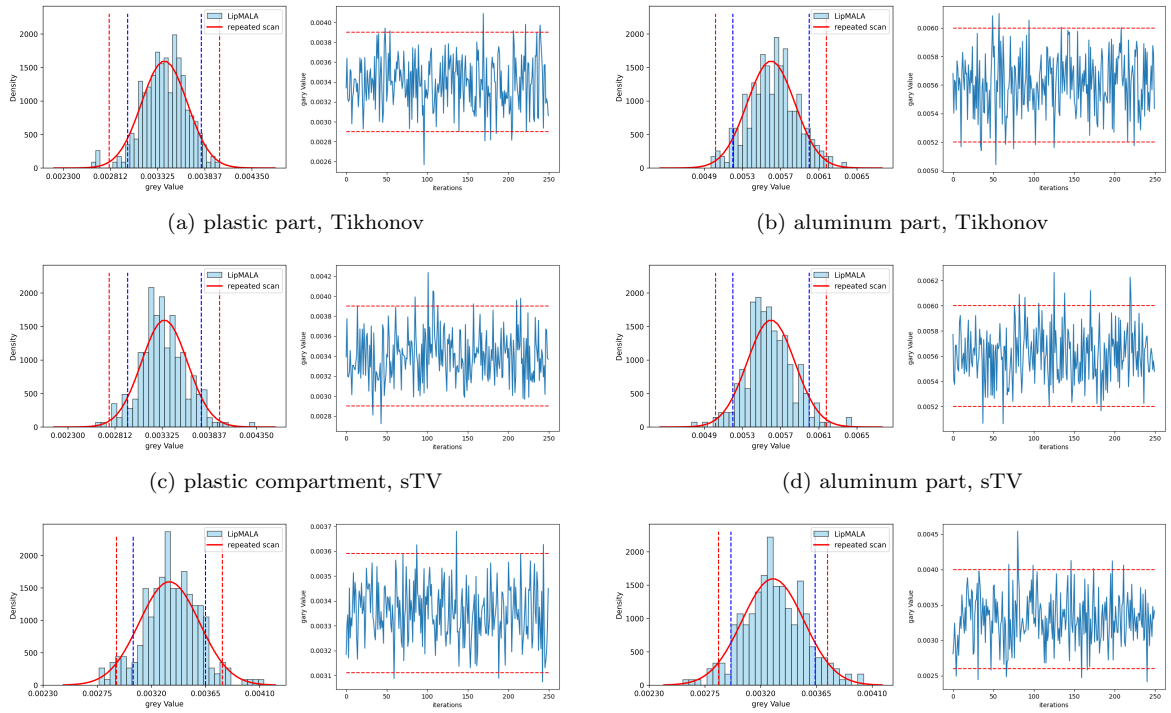

(a) plastic part, Tikhonov

(b) aluminum part, Tikhonov

(c) plastic compartment, sTV

(d) aluminum part, sTV

(e) manifold, Tikhonov

(f) manifold, sTV

Figure 13: Comparison between voxel-wise distributions obtained from repeated scans and LipMALA sampling for selected pixels inside the objects. Top and middle rows: step gauge block (plastic and aluminum regions); bottom row: manifold. Blue histograms: LipMALA sample distributions; blue curves: observed distributions from repeated scans. Red dashed lines indicate the 95% credibility intervals from LipMALA, and red dashed lines show the 95% confidence intervals from the repeated scans.

In addition to the chain traces, histograms of the corresponding sample distributions provide a visual approximation of the posterior distributions estimated by the LipMALA samples. Overlaid in red curve is the distribution of voxel values observed across repeated scans, representing a benchmark for evaluating the performance of the sampling method. As can be seen, the red distributions are consistently wider than the posterior approximations obtained from the LipMALA samples. This is more highlighted by the green dashed lines showing the 95% confidence limit from the repeated scans compared to the red dashed line representing the 95% confidence limit from the LipMALA method. This suggests that the observed voxel variability observed in repeated scans captures additional sources of uncertainty not fully accounted for by the posterior models used in this study, having a bigger confidence limit showing more sources of uncertainty.

The dimensional measurements on the reconstruction of both objects are validated by a similar principle, using the mean, upper, and lower confidence intervals of the measurements done on the repeated 50 scans reconstructed with both regularization methods. The measurement results are presented in Tables 7, and 8 for the step gauge block, and Tables 9 and 10 for the manifold, present measurements with their confidence limit for a clear comparison between the performance of the LipMALA sampling method and repeating the experiment for determining the measurement uncertainty. It is expected that 95% of the dimensional measurements on the 50 reconstruction results from each method should lie within the confidence interval calculated from the LipMALA

samples. However, when calculating the actual percentage, we observed that only 85% and 86% of the measurements on reconstructions of the step gauge block and 83% and 84% of the measurements on the reconstruction of the manifold, from Tikhonov and sTV regularization, fell within the corresponding intervals, respectively.

This discrepancy is reasonable for several reasons. First, only 50 scans may not be sufficient to approximate the true distribution of the reconstructions. Second, the UQ calculations primarily account for noise, but other factors, such as geometrical variations or limited focal spot size or spectrum, can also influence the reconstructions. Although we took measures to minimize the geometrical factors by aligning the reconstruction results through image registration and ensuring accurate geometric parameters, their effect was still present. Despite these challenges, achieving this accuracy in a real-data experiment demonstrates the capability of our UQ method to address practical uncertainties in CT reconstruction effectively.

Table 7: Dimensional measurements on the repeated scans, lower and upper confidence limits, step gauge block

	Tikhonov			sTV		
	u_{LCL}	u_C	u_{UCL}	u_{LCL}	u_C	u_{UCL}
L_1(mm)	3.537	3.539	3.541	3.536	3.538	3.540
L_2(mm)	3.469	3.471	3.473	3.469	3.471	3.473
L_3(mm)	7.009	7.010	7.011	7.010	7.011	7.012

Table 8: Comparison of the dimensional measurements and their confidence limits between the sampling results and the repeated scans for the step gauge block

	Tikhonov		sTV	
	samples	repeated scans	samples	repeated scans
L_1(mm)	3.540 ± 0.002	3.539 ± 0.002	3.538 ± 0.002	3.538 ± 0.002
L_2(mm)	3.471 ± 0.001	3.471 ± 0.002	3.471 ± 0.002	3.471 ± 0.002
L_3(mm)	7.010 ± 0.001	7.010 ± 0.001	7.011 ± 0.001	7.011 ± 0.001

Table 9: Dimensional measurements on the repeated scans, lower and upper confidence limits, manifold

	Tikhonov			sTV		
	u_{LCL}	u_C	u_{UCL}	u_{LCL}	u_C	u_{UCL}
L_1(mm)	3.919	4.014	4.109	3.904	4.001	4.098
L_2(mm)	26.466	26.524	26.582	26.451	26.512	26.573
L_3(mm)	12.978	13.074	13.171	12.940	13.037	13.134
R(mm)	6.473	6.539	6.605	6.498	6.576	6.652

Table 10: Comparison of the dimensional measurements and their confidence limits between the sampling results and the repeated scans for the manifold

	Tikhonov		sTV	
	samples	repeated scans	samples	repeated scans
L_1(mm)	3.990 ± 0.083	4.014 ± 0.095	3.984 ± 0.085	4.001 ± 0.097
L_2(mm)	26.525 ± 0.038	26.524 ± 0.058	26.524 ± 0.04	26.512 ± 0.061
L_3(mm)	13.083 ± 0.086	13.074 ± 0.096	13.050 ± 0.089	13.037 ± 0.097
R(mm)	6.553 ± 0.052	6.539 ± 0.066	6.573 ± 0.058	6.576 ± 0.076

4.7. Discussion

In this study, the primary aim was to determine the measurement uncertainty caused by noise with computationally feasible techniques that are practical in facing real industrial CT data. We focused on two common prior distributions, the Gaussian and Laplacian, which lead to the Tikhonov and sTV regularization methods, which are differentiable and effective for achieving reliable results while maintaining edge features. These methods were chosen for their simplicity and robustness, as well as their compatibility with the Bayesian framework employed in this work. Advanced regularization techniques designed for better preservation of fine details could be considered in future studies. While such techniques may change the form of the posterior distribution, the flexibility of MCMC methods makes them adaptable to these changes. In case the posterior is differentiable, the LipMALA method is a suitable and efficient option, as shown in this study. In cases where the posterior distribution is non-differentiable, alternative MCMC sampling techniques are available.

The Bayesian framework offers considerable flexibility to tailor the posterior distribution to specific needs or the characteristics of the object being scanned. For instance, a model can be adapted to preserve intricate details in more complex objects. In this work, we selected a well-manufactured object for validation of the results, designed for dimensional measurements, minimizing the effects of defects, artifacts, and pores on the final uncertainty quantification. The manifold is selected as a real-life example of the applications of the method.

In choosing the MCMC method, we focused on the Langevin family of techniques adapted by the Metropolis acceptance/rejection step. The choice of the Langevin family is due to employing the gradient information of the posterior distribution, which guides the samples to the parts with a higher probability, resulting in a more efficient sampling process. Additionally, each sample is checked by the acceptance probability from the Metropolis method, which again reduces the number of samples needed to be stored for further analysis. With a fixed step size for this sampling method, we have the MALA method. Although the MALA method offers various benefits, the manual tuning of the step size for the CT reconstruction problem requires effort. The LipMALA method offers an adaptive step size property, which not only eliminates manual step size tuning but also reduces the number of samples needed for a reliable estimation of the posterior distribution.

The choice of the prior distribution also affects the uncertainty introduced by the numerical algorithms used to calculate the MAP estimate. The CG method, which finds the MAP from the Gaussian distribution, and the AGD method, which calculates the MAP from the Laplacian distribution, differ in their convergence properties. This leads to variations in the confidence intervals around dimensional measurements. This highlights the interplay between regularization, optimization algorithms, and the resulting uncertainty.

Overall, the Bayesian framework demonstrated in this work is adaptable and versatile, capable of addressing multiple error sources. By focusing on UQ caused by noise, we showcased how the LipMALA method can efficiently handle real CT data, providing a strong foundation for extending this approach to include systematic errors or more sophisticated model assumptions in the future.

5. Conclusion

In this paper, we demonstrated the superior performance of sTV regularization over Tikhonov regularization in the context of CT image reconstruction. sTV regularization proved to be more effective in preserving edges and structural details, making it suitable for applications requiring high accuracy in details. In contrast, Tikhonov regularization produced smoother but less detailed reconstructions.

For the determination of the uncertainty around the result, we employed the Bayesian approach with Langevin sampling techniques, specifically the Metropolis-adjusted Langevin methods. By automating the step size tuning process with respect to the Lipschitz constant of the distribution, the LipMALA method overcame a key limitation of standard MALA by eliminating the manual adjustment of the step size, facilitating the sampling process while improving efficiency. This adaptive approach significantly reduced the number of samples needed to estimate the posterior distribution, enabling computationally feasible and robust UQ for real CT data.

Furthermore, the efficiency of LipMALA, compared to traditional methods like Monte Carlo simulations, allowed us to handle real-world data and provide detailed UQ in the form of upper and lower confidence intervals derived from the posterior distribution. These confidence intervals proved valuable for conducting dimensional measurements on the reconstructed images, ensuring reliable and practical results.

Our method combines regularization techniques with an optimized sampling approach to offer a comprehensive UQ framework for industrial CT applications. By addressing the complexities of real data while maintaining computational efficiency, this framework demonstrates its potential for accurate and practical UQ in a variety of settings.

Future work could extend this approach to incorporate additional sources of uncertainty, such as geometrical parameters, into the Bayesian framework. This would further enhance the method's versatility and its ability to handle a broader range of applications.

Acknowledgments

This research work was undertaken in the context of Horizon 2020 European Innovative Training Networks (ITN) project - xCTing ("Enabling X-ray CT based industry 4.0 process chains by training next generation research experts", https://xcting-itn.eu/) with Project ID: 956172.

References

[1] P. C. Hansen, J. Jørgensen, W. R. Lionheart, Computed tomography: algorithms, insight, and just enough theory, SIAM, 2021.
[2] S. Carmignato, W. Dewulf, R. Leach, Industrial X-ray computed tomography, volume 10, Springer, 2018.
[3] F. Natterer, F. Wübbeling, Mathematical methods in image reconstruction, SIAM, 2001.
[4] L. A. Feldkamp, L. C. Davis, J. W. Kress, Practical cone-beam algorithm, Josa a 1 (1984) 612–619.
[5] J. Trampert, J.-J. Leveque, Simultaneous iterative reconstruction technique: Physical interpretation based on the generalized least squares solution, Journal of Geophysical Research: Solid Earth 95 (1990) 12553–12559.

[6] M. Pereyra, P. Schniter, E. Chouzenoux, J.-C. Pesquet, J.-Y. Tourneret, A. O. Hero, S. McLaughlin, A survey of stochastic simulation and optimization methods in signal processing, IEEE Journal of Selected Topics in Signal Processing 10 (2015) 224–241.
[7] J. Bian, J. H. Siewerdsen, X. Han, E. Y. Sidky, J. L. Prince, C. A. Pelizzari, X. Pan, Evaluation of sparse-view reconstruction from flat-panel-detector cone-beam ct, Physics in Medicine & Biology 55 (2010) 6575.
[8] G.-H. Chen, J. Tang, S. Leng, Prior image constrained compressed sensing (piccs): a method to accurately reconstruct dynamic ct images from highly undersampled projection data sets, Medical physics 35 (2008) 660–663.
[9] W. Dewulf, H. Bosse, S. Carmignato, R. Leach, Advances in the metrological traceability and performance of x-ray computed tomography, CIRP Annals 71 (2022) 693–716.
[10] Joint Committee for Guides in Metrology (JCGM), Guide to the expression of uncertainty in measurement — part 1: Introduction, https://www.bipm.org/documents/20126/2071204/JCGM_GUM-1.pdf, 2023. JCGM GUM-1:2023.
[11] E. Ametova, M. Ferrucci, S. Chilingaryan, W. Dewulf, A computationally inexpensive model for estimating dimensional measurement uncertainty due to x-ray computed tomography instrument misalignments, Measurement Science and Technology 29 (2018) 065007.
[12] J. Bredemann, R. H. Schmitt, Task-specific uncertainty estimation for medical ct measurements, Journal of Sensors and Sensor Systems 7 (2018) 627–635.
[13] M. Ferrucci, E. Ametova, W. Dewulf, Monte carlo reconstruction: a concept for propagating uncertainty in computed tomography, Measurement Science and Technology 32 (2021) 115006.
[14] P. J. Green, K. Łatuszyński, M. Pereyra, C. P. Robert, Bayesian computation: a summary of the current state, and samples backwards and forwards, Statistics and Computing 25 (2015) 835–862.
[15] J. P. Kaipio, E. Somersalo, Statistical inversion theory, Statistical and computational inverse problems (2005) 49–114.
[16] A. Durmus, E. Moulines, M. Pereyra, Efficient bayesian computation by proximal markov chain monte carlo: when langevin meets moreau, SIAM Journal on Imaging Sciences 11 (2018) 473–506.
[17] M. Dashti, A. M. Stuart, The bayesian approach to inverse problems, in: Handbook of uncertainty quantification, Springer, 2017, pp. 311–428.
[18] J. S. J. Frederik Hagsholm Pedersen, M. S. Andersen, A bayesian approach to ct reconstruction with uncertain geometry, Applied Mathematics in Science and Engineering 31 (2023) 2166041. doi:10.1080/27690911.2023.2166041.
[19] N. A. B. Riis, Y. Dong, P. C. Hansen, Computed tomography with view angle estimation using uncertainty quantification, Inverse Problems 37 (2021) 065007. URL: https://dx.doi.org/10.1088/1361-6420/abf5ba. doi:10.1088/1361-6420/abf5ba.
[20] B. M. Afkham, Y. Dong, P. C. Hansen, Uncertainty quantification of inclusion boundaries in the context of x-ray tomography, SIAM/ASA Journal on Uncertainty Quantification 11 (2023) 31–61. doi:10.1137/21M1433782.
[21] F. Uribe, J. M. Bardsley, Y. Dong, P. C. Hansen, N. A. B. Riis, A hybrid gibbs sampler for edge-preserving tomographic reconstruction with uncertain view angles, SIAM/ASA Journal on Uncertainty Quantification 10 (2022) 1293–1320. doi:10.1137/21M1412268.
[22] S. Ruder, An overview of gradient descent optimization algorithms, arXiv preprint arXiv:1609.04747 (2016).
[23] T. L. Jensen, J. H. Jørgensen, P. C. Hansen, S. H. Jensen, Implementation of an optimal first-order method for strongly convex total variation regularization, BIT Numerical Mathematics 52 (2012) 329–356.
[24] W. R. Gilks, S. Richardson, D. Spiegelhalter, Markov chain Monte Carlo in practice, CRC press, 1995.
[25] L. Tierney, Markov chains for exploring posterior distributions, the Annals of Statistics (1994) 1701–1728.
[26] J. Hadamard, Sur les problèmes aux dérivées partielles et leur signification physique, Princeton university bulletin (1902) 49–52.
[27] A. C. Kak, M. Slaney, Principles of computerized tomographic imaging, SIAM, 2001.
[28] B. Jorgensen, Statistical properties of the generalized inverse Gaussian distribution, volume 9, Springer Science & Business Media, 2012.
[29] A. N. Tikhonov, Solutions of ill-posed problems, VH Winston and Sons (1977).
[30] L. I. Rudin, S. Osher, E. Fatemi, Nonlinear total variation based noise removal algorithms, Physica D: nonlinear phenomena 60 (1992) 259–268.
[31] P. J. Huber, E. M. Ronchetti, Robust statistics, John Wiley & Sons, 2011.
[32] A. Chambolle, An approximation result for special functions with bounded deformation, Journal de mathématiques pures et appliquées 83 (2004) 929–954.
[33] E. Y. Sidky, J. H. Jørgensen, X. Pan, Convex optimization problem prototyping for image reconstruction in

computed tomography with the chambolle–pock algorithm, Physics in Medicine & Biology 57 (2012) 3065.
[34] A. Chambolle, T. Pock, An introduction to continuous optimization for imaging, Acta Numerica 25 (2016) 161–319.
[35] J. Nocedal, S. J. Wright, Numerical optimization, Springer, 1999.
[36] Y. Nesterov, A method for solving the convex programming problem with convergence rate o (1/k2), in: Dokl akad nauk Sssr, volume 269, 1983, p. 543.
[37] N. Brosse, A. Durmus, E. Moulines, The promises and pitfalls of stochastic gradient langevin dynamics, Advances in Neural Information Processing Systems 31 (2018).
[38] P. Langevin, et al., On the theory of brownian motion, CR Acad. Sci. Paris 146 (1908) 530.
[39] S. Brooks, A. Gelman, G. Jones, X.-L. Meng, Handbook of Markov chain monte carlo, CRC press, 2011.
[40] M. Izzatullah, T. van Leeuwen, D. Peter, Langevin dynamics markov chain monte carlo solution for seismic inversion, in: 82nd EAGE Annual Conference & Exhibition, volume 2021, EAGE Publications BV, 2021, pp. 1–5.
[41] A. Stolfi, L. De Chiffre, Selection of items for "interaqct comparison on assemblies", in: 6th Conference on Industrial Computed Tomography, 2016.
[42] M. Ferrucci, P. Heřmánek, E. Ametova, S. Carmignato, W. Dewulf, Measurement of the x-ray computed tomography instrument geometry by minimization of reprojection errors—implementation on simulated data, Precision Engineering 54 (2018) 7–20.
[43] P. Guerrero, S. Bellens, R. Santander, W. Dewulf, Automatic and computationally efficient alignment in fan-and cone-beam tomography, IEEE Transactions on Computational Imaging (2024).
[44] W. Van Aarle, W. J. Palenstijn, J. Cant, E. Janssens, F. Bleichrodt, A. Dabravolski, J. De Beenhouwer, K. Joost Batenburg, J. Sijbers, Fast and flexible x-ray tomography using the astra toolbox, Optics express 24 (2016) 25129–25147.
[45] A. Hendriksen, D. Schut, W. J. Palenstijn, N. Viganò, J. Kim, D. Pelt, T. van Leeuwen, K. J. Batenburg, Tomosipo: Fast, flexible, and convenient 3D tomography for complex scanning geometries in Python, Optics Express (2021). URL: https://doi.org/10.1364/oe.439909. doi:10.1364/oe.439909.
[46] M. Guizar-Sicairos, S. T. Thurman, J. R. Fienup, Efficient subpixel image registration algorithms, Optics letters 33 (2008) 156–158.
[47] S. Van der Walt, J. L. Schönberger, J. Nunez-Iglesias, F. Boulogne, J. D. Warner, N. Yager, E. Gouillart, T. Yu, scikit-image: image processing in python, PeerJ 2 (2014) e453.
[48] D. L. Donoho, I. M. Johnstone, Ideal spatial adaptation by wavelet shrinkage, biometrika 81 (1994) 425–455.