

Fast Gibbs sampling for high-dimensional Bayesian inversion

This content has been downloaded from IOPscience. Please scroll down to see the full text.

2016 Inverse Problems 32 115019

(<http://iopscience.iop.org/0266-5611/32/11/115019>)

View [the table of contents for this issue](#), or go to the [journal homepage](#) for more

Download details:

IP Address: 128.41.35.12

This content was downloaded on 07/11/2016 at 15:41

Please note that [terms and conditions apply](#).

You may also be interested in:

[Fast Markov chain Monte Carlo sampling for sparse Bayesian inference in high-dimensional inverse problems using L1-type priors](#)

Felix Lucka

[Maximum a posteriori estimates in linear inverse problems with log-concave priors are proper Bayes estimators](#)

Martin Burger and Felix Lucka

[Inverse problems with Poisson data: statistical regularization theory, applications and algorithms](#)

Thorsten Hohage and Frank Werner

[Statistical inversion and Monte Carlo sampling methods in EIT](#)

Jari P Kaipio, Ville Kolehmainen, Erkki Somersalo et al.

[A TV-Gaussian prior for infinite-dimensional Bayesian inverse problems and its numerical implementations](#)

Zhewei Yao, Zixi Hu and Jinglai Li

[MAP estimators for piecewise continuous inversion](#)

M M Dunlop and A M Stuart

Fast Gibbs sampling for high-dimensional Bayesian inversion

Felix Lucka

Centre for Medical Image Computing, University College London, WC1E 6BT
London, UK

E-mail: f.lucka@ucl.ac.uk

Received 27 February 2016, revised 28 June 2016

Accepted for publication 7 September 2016

Published 7 October 2016



CrossMark

Abstract

Solving ill-posed inverse problems by Bayesian inference has recently attracted considerable attention. Compared to deterministic approaches, the probabilistic representation of the solution by the posterior distribution can be exploited to explore and quantify its uncertainties. In applications where the inverse solution is subject to further analysis procedures can be a significant advantage. Alongside theoretical progress, various new computational techniques allow us to sample very high dimensional posterior distributions: in (Lucka 2012 *Inverse Problems* 28 125012), and a Markov chain Monte Carlo posterior sampler was developed for linear inverse problems with ℓ_1 -type priors. In this article, we extend this single component (SC) Gibbs-type sampler to a wide range of priors used in Bayesian inversion, such as general ℓ_p^q priors with additional hard constraints. In addition, a fast computation of the conditional, SC densities in an explicit, parameterized form, a fast, robust and exact sampling from these one-dimensional densities is key to obtain an efficient algorithm. We demonstrate that a generalization of slice sampling can utilize their specific structure for this task and illustrate the performance of the resulting slice-within-Gibbs samplers by different computed examples. These new samplers allow us to perform sample-based Bayesian inference in high-dimensional scenarios with certain priors for the first time, including the inversion of computed tomography data with the popular isotropic total variation prior.

Keywords: Bayesian inversion, MCMC, Gibbs sampler, slice sampling, computed tomography, total variation prior

(Some figures may appear in colour only in the online journal)



Original content from this work may be used under the terms of the [Creative Commons Attribution 3.0 licence](https://creativecommons.org/licenses/by/3.0/). Any further distribution of this work must maintain attribution to the author(s) and the title of the work, journal citation and DOI.

1. Introduction

1.1. Bayesian inversion

We consider the task of inferring information about an unknown quantity from indirect, noisy measurements and assume that a reasonable mathematical model is given by a linear, ill-posed operator equation including additive noise terms. The following discrete forward model is used to carry out the computational inference:

$$f = A u + \varepsilon. \quad (1)$$

Here, $f \in \mathbb{R}^m$ represents the measured data, $u \in \mathbb{R}^n$ a suitable discretization of the unknown quantity we wish to reconstruct and $A \in \mathbb{R}^{m \times n}$ a corresponding discretization of the continuous forward operator. We assume that the statistics of the additive noise can be well-approximated by a Gaussian distribution $\mathcal{N}(\mu, \Sigma)$ and that f and A are already centered and decorrelated with respect to μ and Σ , i.e., $f = \Sigma^{-1/2}(\tilde{f} - \mu)$ and $A = \Sigma^{-1/2}\tilde{A}$, where \tilde{f} and \tilde{A} denote the original data and forward operator, respectively. This leads to $\varepsilon \sim \mathcal{N}(0, I_m)$ and the following *likelihood* distribution

$$p_{\text{like}}(f|u) \propto \exp\left(-\frac{1}{2} \|f - A u\|_2^2\right), \quad (2)$$

which is a probabilistic model of the measured data f given the unknown solution u . In typical inverse problems, solving (1) for u is ill-posed. As a consequence, the information that (2) contains about u is insufficient to carry out robust inference and we need to amend it by *a priori* information, encoded in a *prior* distribution $p_{\text{prior}}(u)$. Then, the total information on u we have *after* performing the measurement is encoded by the conditional distribution of u given f , the so-called *a posteriori* distribution. It can be computed by *Bayes' rule*:

$$p_{\text{post}}(u|f) = \frac{p_{\text{like}}(f|u)p_{\text{prior}}(u)}{p(f)}. \quad (3)$$

Figure 1 illustrates the inference process. Originating from statistical physics, *Gibbs distributions* are commonly used prior models:

$$p_{\text{prior}}(u) \propto \exp(-\lambda \mathcal{J}(u)). \quad (4)$$

The functional $\mathcal{J}(u)$ measures an *energy* of u . The use of Gibbs priors leads to

$$p_{\text{post}}(u|f) \propto \exp\left(-\frac{1}{2} \|f - A u\|_2^2 - \lambda \mathcal{J}(u)\right). \quad (5)$$

For a general introduction to *Bayesian inversion* we refer to [20, 28, 39] and references therein. The recent attention on this particular inversion approach is best reflected by the recent special issue of *Inverse Problems* [7], which also provides a good overview over current developments and trends.

1.2. Sample-based inference

While the posterior $p_{\text{post}}(u|f)$ represents our complete knowledge about u , *Bayesian estimation* tries to extract the information of interest from it. Classical examples thereof include the *maximum a posteriori estimate (MAP)* and the *conditional mean estimate (CM)*

$$\hat{u}_{\text{MAP}} := \operatorname{argmax}_{u \in \mathbb{R}^n} \{p_{\text{post}}(u|f)\}, \quad \hat{u}_{\text{CM}} := \mathbb{E}[u|f] = \int u p_{\text{post}}(u|f) du, \quad (6)$$

which both yield a single point estimate of u . The details on their properties and relationship can be found in [6]. More sophisticated estimators such as *conditional covariance (CCov)*, *conditional variance (CVar)* or *standard deviation (CStd)* estimates try to extract higher order

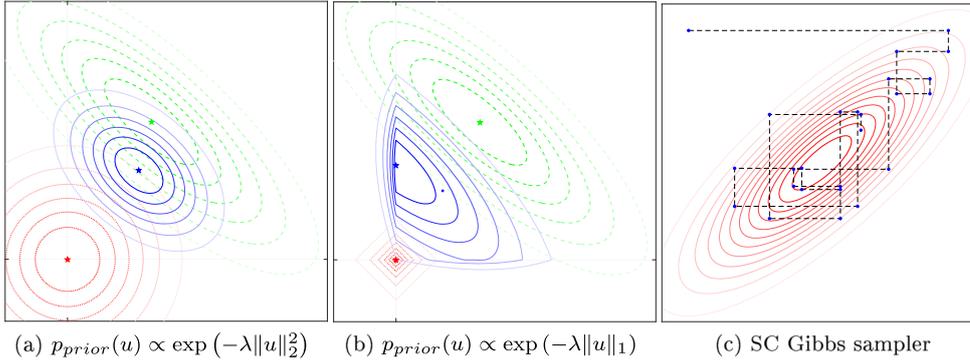


Figure 1. (a)–(b) Illustration of Bayesian inversion with different priors. Depicted are the level sets of likelihood (green, dashed), prior (red, dotted) and resulting posterior distribution (blue, solid). Maximal and expected values of the corresponding distributions are depicted by a star and a bullet, respectively. (c) Example of an MCMC chain generated by an SC Gibbs sampler (blue bullets connected by dashed lines) to sample a bivariate target density (level sets shown as red solid lines).

statistics of u , or try to quantify the uncertainties of u , for instance through *credible region/interval* and *extreme value probability* estimators.

Bayesian computation refers to the practical task of computing the above estimators. For most inversion scenarios and prior models, this involves solving high-dimensional optimization or integration tasks (see (6)), or even a mix of both. In this article, we are examining techniques that integrate $p_{\text{post}}(u|f)$ by *Monte Carlo integration*:

$$\int g(u) p_{\text{post}}(u|f) du \approx \frac{1}{K} \sum_i^K g(u^i), \quad (7)$$

where u^i are samples of $p_{\text{post}}(u|f)$ generated by a *sampling algorithm/sampler*. Due to the lack of efficient *direct samplers* that generate i.i.d. samples, *Markov chain Monte Carlo (MCMC)* samplers need to be employed in most situations. MCMC for high dimensional Bayesian inversion is a very active field of research, see, e.g., [1, 2, 4, 11, 12, 16, 22, 31–33, 36–38] for some examples of recent developments. In [27], an efficient MCMC sampler for Gibbs priors with ℓ_1 -norm-type energies (ℓ_1 -priors, see figure 1(b)) was presented:

$$p_{\text{prior}}(u) \propto \exp(-\lambda \|D^T u\|_1). \quad (8)$$

Such energies are commonly used to impose *sparsity constraints* on the solution of high dimensional image reconstruction problems, a direction of research closely related to the notion of *compressed sensing* [8, 14, 15]. A detailed discussion of sparsity as *a priori* information in Bayesian inversion can be found in [28]. The sampler developed in [27] belongs to the class of *single component (SC) Gibbs samplers*, which sample $p_{\text{post}}(u|f)$ by subsequently sampling along the conditional, SC densities $p_{\text{post}}(u_j|u_{-j}, f)$:

Algorithm 1. (SC-Gibbs sampling)

Define an initial state u^0 , a burn-in size K_0 and sample size K . For $i = 1, \dots, K_0 + K$ do:

A1.1 Choose a component j (deterministic or random).

A1.2 Draw $y \sim p_{\text{post}}(\cdot | u_{-j}, f)$.

A1.3 Set $u_j^{i+1} = y$, and $u_{-j}^{i+1} = u_{-j}^i$.

Discard $\{u^i\}_{i=0}^{K_0}$ and use $\{u^i\}_{i=K_0+1}^{K_0+K}$ as a sample of $p_{\text{post}}(u|f)$.

We have used $x_{-j} := (x_1, \dots, x_{j-1}, x_{j+1}, \dots, x_n)^T$ to denote a vector with all components of x except for x_j . An illustration of Gibbs sampling is given in figure 1(c). In [27], this SC-Gibbs sampler was compared to the popular *Metropolis–Hastings (MH)* sampler: for the computational scenarios considered and the evaluation performed, it was demonstrated that in contrast to the MH, SC-Gibbs sampling gets more efficient when the level of sparsity or the dimension of the unknowns is increased. Thereby, it became possible to carry out sample-based inference with ℓ_1 priors in challenging inverse problems scenarios with $n > 10^6$:

- The theoretical predictions about the infinite dimensional limits of total variation (TV) priors posed in [10, 24] could be verified numerically (see [28]).
- *Computed tomography (CT)* inversion with Besov space priors (see [18, 21]) was examined for simulated and experimental data (see [6, 28]).
- The numerical results stimulated the development of new theoretical ideas about the relationship of MAP and CM estimates (see [6, 19]).

1.3. Previous limitations

As the sampler developed in [27] relies on a direct sampling of the SC densities, namely the *inverse cumulative distribution method (iCDF)*, we will call it the *direct ℓ_1 sampler* from now on. While the direct ℓ_1 sampler works well in the applications described above, it suffers from several limitations. To understand them, we recall that an efficient SC-Gibbs sampler needs to:

- (SC1) Compute the conditional, SC densities in an explicit, parameterized form in a fast way.
- (SC2) Employ a fast, robust and exact sampling scheme for the parameterized form of the SC densities.

In order to best fulfill (SC1) and (SC2), the direct ℓ_1 sampler was designed for a very particular setting: firstly, in addition to relying on a linear forward map (1) and a Gaussian noise model (2), it assumes that the operator $D \in \mathbb{R}^{n \times h}$ in (8) can be diagonalized (*synthesis priors*): there is a basis matrix V such that $D^T V$ is a diagonal matrix $W \in \mathbb{R}^{h \times n}$. The direct ℓ_1 sampler then samples over the coefficients of $u = V\xi$ in this basis:

$$p_{\text{post}}(\xi|f) \propto \exp\left(-\frac{1}{2} \|f - AV \xi\|_2^2 - \lambda \|W\xi\|_1\right). \quad (9)$$

This excludes the use of *frames* or *dictionaries* for D . Secondly, it only works for the ℓ_1 norm as a prior energy: a straight-forward extension of iCDF to examine more general ℓ_p^q -prior of the form $p_{\text{prior}}(u) \propto \exp(-\lambda \|D^T u\|_p^q)$ is not possible. This excludes the interesting cases of $q = p$, $p < 1$, which leads to a non-convex energy but also $p = 1$, $q > 1$, which was examined in [10]. Finally, a lot of interesting priors such as the popular isotropic TV prior in 2D/3D or related, *bloc/structured sparsity* priors have a more involved structure than (8) and cannot be treated with iCDF in an efficient and robust way as well. In all the above cases, including additional *hard constraints*, $u \in \mathcal{C}$, where \mathcal{C} is the *feasible set* of solutions is often advantageous:

$$p_{\text{prior}}(u) \propto \tilde{p}_{\text{prior}}(u) \cdot \mathbf{1}_{\mathcal{C}}(u) = \begin{cases} \tilde{p}_{\text{prior}}(u) & \text{if } u \in \mathcal{C}, \\ 0 & \text{otherwise.} \end{cases} \quad (10)$$

While such constraints have proven to be very useful as *a priori* knowledge [3, 40], their inclusion into the direct ℓ_1 sampler in a numerically stable way is cumbersome.

1.4. Contributions and structure

For most of the limitations discussed above, the main problem is not to fulfill (SC1), but to fulfill (SC2) by using a direct sampler such as iCDF for the parameterized SC densities in step A1.2. In this article, we sample from them by using a generalization of *slice sampling* that utilizes their specific structure instead and demonstrate the effectiveness of this replacement in different computed examples. This allows us to perform sample-based Bayesian inference in high-dimensional scenarios with the priors described above for the first time.

The paper is structured as follows: in section 2, we first derive the SC densities for the priors discussed above. Then, we introduce the basic and generalized slice sampler and discuss how to integrate it into the SC Gibbs sampler for Bayesian inversion. Section 3 contains computed examples and section 4 closes with a discussion. Several technical details are covered in section appendix A.

2. Sampling methods

For general and comprehensive introductions to MCMC sampling methods, we refer to [26, 34].

2.1. SC posterior densities

In this section, we briefly derive the SC posterior densities for the examined prior models in a simple, parameterized way, see (SC1). We first discuss the case where a basis $\{v_1, \dots, v_n\}$ helps to represent $u = \sum \xi_l v_l =: V\xi$ such that $p_{\text{post}}(\xi_j | \xi_{-j}, f)$ can be described using as few parameters as possible. Once such a basis is found, the part of $p_{\text{post}}(\xi_j | \xi_{-j}, f)$ coming from the likelihood is easy to derive: we define $\Psi := AV$ and $\varphi(j) := f - \Psi_{-j}\xi_{-j}$. Then, we find that

$$\begin{aligned} \frac{1}{2} \|f - Au\|_2^2 &= \frac{1}{2} \|f - AV\xi\|_2^2 = \frac{1}{2} \|f - \Psi\xi\|_2^2 = \frac{1}{2} \|f - (\Psi_{-j}\xi_{-j} + \Psi_j\xi_j)\|_2^2 \\ &= \frac{1}{2} \|\varphi(j) - \Psi_j\xi_j\|_2^2 \stackrel{\xi_j}{\propto} \frac{1}{2} \|\Psi_j\|_2^2 \xi_j^2 + \Psi_j^T \varphi(j) \xi_j =: ax^2 - bx, \end{aligned} \quad (11)$$

where we introduced $x := \xi_j$, $a := \frac{1}{2} \|\Psi_j\|_2^2$, and $b := \Psi_j^T \varphi(j) = \Psi_j^T f - (\Psi_j^T \Psi_{-j}) \xi_{-j}$ to ease the notation for the following sections. Note that while a and $\Psi_j^T f$ can be precomputed, $(\Psi_j^T \Psi_{-j}) \xi_{-j}$ relies on the current state of the ξ -chain and has to be computed in every step of the sampler. Especially for complicated forward operators in high-dimensional scenarios, this operation is the computational bottleneck of SC Gibbs samplers. Therefore, a careful, scenario-dependent implementation is important to obtain a fast sampler.

Now we proceed to determine V and the part of $p_{\text{post}}(\xi_j | \xi_{-j}, f)$ coming from the prior. The energies of the ℓ_p^q priors can be written as

$$\mathcal{J}(u) = \left(\sum_k^h |D_k^T u|^p \right)^{\frac{q}{p}} = \left(\sum_k^h \left| \sum_l (D_k^T v_l) \xi_l \right|^p \right)^{\frac{q}{p}}. \quad (12)$$

To obtain simple conditional densities for all ξ_j , we thus have to choose V such that

$$\max_l |D^T v_l|_0, \quad \text{where} \quad |u|_0 := \text{card}(\{k | u_k \neq 0\}), \quad (13)$$

is as small as possible. We first consider the special but important case of $D^T \in \mathbb{R}^{h \times n}$ having full rank and $h \leq n$. This includes the case where the columns D are elements of a basis, and thereby, the class of *Besov priors*, see [6, 13, 18, 21, 23] and the TV prior in 1D with

Neumann boundary conditions, which we will use in the computational studies. Due to the full rank, we can choose v_1, \dots, v_h such that $D^T v_l = e_l$ for $l = 1, \dots, h$, and v_{h+1}, \dots, v_n such that $D^T v_l = 0$ for $l = h+1, \dots, n$ (for D being a basis, we have $V = D$). With this transformation, (12) simplifies to:

$$\mathcal{J}(\xi) \propto \left(\sum_l^h |\xi_l|^p \right)^{\frac{q}{p}} = \left(|\xi_j|^p + \sum_{l \neq j}^h |\xi_l|^p \right)^{\frac{q}{p}}. \quad (14)$$

Defining $x := \xi_j$ as above, we can write the conditional SC posterior density as:

$$p(x) \propto \exp(-ax^2 + bx - c(|x|^p + d)^{q/p}), \quad c := \lambda \mathbb{1}_{\{j \leq h\}}, \quad d := \sum_{l \neq j}^h |\xi_l|^p, \quad (15)$$

which simplifies to

$$p(x) \propto \exp(-ax^2 + bx - c|x|^p), \quad c := \lambda \mathbb{1}_{\{j \leq h\}}, \quad (16)$$

for ℓ_p^p priors. In the case where D cannot be diagonalized, an explicit form is given by

$$p(x) \propto \exp \left(-ax^2 + bx - c \left(\sum_{k \in \text{supp}(D^T v_j)} |d_k x - e_k|^p \right)^{\frac{q}{p}} \right), \quad (17)$$

$$\text{where } c := \lambda \mathbb{1}_{\{j \leq h\}}, \quad d_k := (D^T v_j)_k, \quad e_k := (D^T V_{-j} \xi_{-j})_k. \quad (18)$$

Various generalizations of the standard ℓ_p^q priors with $\|D^T u\|_p^q$ -type energies first compute the ℓ_2 -norm of a *local feature* of u , e.g., of its gradient, and then measure the *global* ℓ_p^q energy of these local ℓ_2 norms. In this article, we will only discuss one prominent example thereof, which is given by the isotropic TV prior in 2D: if we assume that u represents an $N \times N$ discrete image, we can index the components of u as $u_{(k,l)}$ with $k = 1, \dots, N, l = 1, \dots, N, n = N^2$. We can then use forward differences in both spatial directions to define

$$\mathcal{J}_{\text{TV}}(u) = \sum_{(k,l)}^n \sqrt{(u_{(k+1,l)} - u_{(k,l)})^2 + (u_{(k,l+1)} - u_{(k,l)})^2}, \quad (19)$$

with appropriate additional boundary conditions. The local nature of the $\mathcal{J}_{\text{TV}}(u)$ allows us to derive a simple parameterization of the SC densities in the pixel basis $V = I_n$. Every $\xi_j = u_{(k,l)}$ only appears in three terms of the energy:

$$\begin{aligned} \mathcal{J}_{\text{TV}}(u_{(k,l)} | u_{-(k,l)}) &\stackrel{(k,l)}{\propto} \sqrt{(u_{(k+1,l)} - u_{(k,l)})^2 + (u_{(k,l+1)} - u_{(k,l)})^2} \\ &+ \sqrt{(u_{(k,l)} - u_{(k-1,l)})^2 + (u_{(k-1,l+1)} - u_{(k-1,l)})^2} \\ &+ \sqrt{(u_{(k+1,l-1)} - u_{(k,l-1)})^2 + (u_{(k,l)} - u_{(k,l-1)})^2}. \end{aligned} \quad (20)$$

Therefore, we can write the conditional SC posterior as

$$p(x) \propto \exp \left(-ax^2 + bx - c \sum_{k=1}^3 \sqrt{d_k (x - e_k)^2 + g_k} \right), \quad (21)$$

with appropriately computed parameters $d_k \in \{0, 1, 2\}$, $e_k, g_k \geq 0$.

The difficulty of incorporating additional hard constraints (10) depends on the shape of the feasible set \mathcal{C} and the transformation V applied. In the following, we assume that they lead to a feasible (semi-)finite interval $[lb, ub]$ to which the continuous densities computed above

can be restricted to. In the case of \mathcal{C} being convex, such an interval always exists and there are computationally efficient ways to compute it.

2.2. MCMC-within-Gibbs sampling

The direct ℓ_1 sampler is sampling (16) with $p = 1$ by the iCDF method using an explicit form of the inverse CDF. For $p, q \neq 1$ or (21) this is not possible and one would need to integrate the CDF numerically to use the iCDF method as a SC density sampler. However, already for $p = 1$, a major technical difficulty was to develop a numerical implementation that worked for all possible combinations (a, b, c) : the first implementations broke down when the dimension n of the problem was increased and the ill-posedness became more severe. The reason was that combinations of (a, b, c) corresponding to extremely degenerate SC densities appeared more frequently for $n \rightarrow \infty$ and in general, the variability of SC densities grows. This trend will be an even more severe problem when one cannot find an explicit form of the inverse CDF and needs so resort to numerical integration. But also replacing the iCDF method by a univariate MCMC sampler (*MCMC-within-Gibbs sampling*) becomes challenging: The most commonly used *Metropolis-within-Gibbs* sampler, which utilizes an easy-to-implement MH sampler with a univariate Gaussian proposal $\mathcal{N}(x, \kappa)$ (where x is the current state) for the SC sampling step A1.2 will not work properly in such a situation: for an MH sampler to be efficient, finding a value of κ leading to an optimal *acceptance rate* is essential. However, the large variations in-between SC densities renders an automatic tuning of a single κ impossible. The alternative would be to tune and use a different κ_j for every component j , but the tuning procedure would require n times more samples than tuning one κ for all components. Thereby, the resulting algorithm would be more like an *adaptive* SC–MH sampler than a Gibbs sampler [17, 25].

2.3. Slice sampling

Slice sampling transfers the automatic adaptation of Gibbs sampling to univariate densities. While the basic version to sample arbitrary densities in a ‘back-box’ fashion was proposed in [30], we follow the presentation given in [34], which leads to a general version in which we can utilize several properties of our specific posterior densities to derive an efficient SC sampler. The starting point for slice sampling is the *Fundamental theorem of Simulation*, which states that sampling from a distribution $p(x)$ is equivalent to sampling uniformly from the area under the graph of $p(x)$: $\mathcal{G}_p := \{(x, z) | 0 \leq z \leq p(x)\}$. This simple observation is the basis of *accept–reject samplers*, a widely used class of samplers which draw uniform samples (x, z) from a region enclosing \mathcal{G}_p and only accept the sample if it fulfills $z \leq p(x)$. Figures 2(a) and (b) illustrate this principle. Slice sampling utilizes this principle in another way: it samples the auxiliary, bivariate density $\tilde{p}(x, z) \propto \mathbb{1}_{\mathcal{G}_p}(x, z)$ by a Gibbs sampler and only keeps the x samples, see figure 2(c):

Algorithm 2. (Basic slice sampling)

For a univariate density $p(x)$, define an initial state x^0 , a burn-in size K_0 and a sample size K . For $i = 1, \dots, K_0 + K$ do:

A2.1 Draw y uniform from $[0, p(x^i)]$ (vertical move).

A2.2 Draw x^{i+1} uniform from $S^y := \{z | p(z) \geq y\}$ (horizontal move).

Discard $\{x^i\}_{i=0}^{K_0}$ and use $\{x^i\}_{i=K_0+1}^{K_0+K}$ as a sample of $p(x)$.

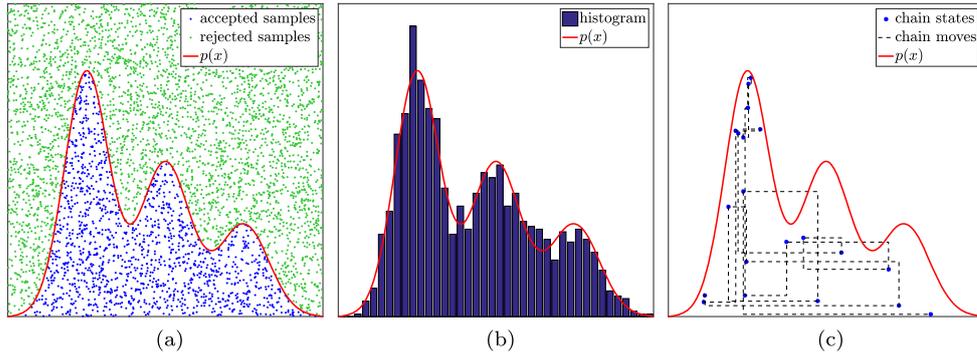


Figure 2. An illustration of accept–reject methods and slice sampling: (a) to sample from the density $p(x)$ (red line), uniform samples (x^i, y^i) (blue and green dots) are generated in a region enclosing its graph. All samples fulfilling $y^i \leq p(x^i)$ (blue dots) are accepted. (b) Histogram computed from the x values of all accepted samples. (c) Slice sampling. The x coordinates of the blue dots are samples of $p(x)$, while the dashed black line illustrates the path of the sampler on \mathcal{G}_p .

The difficulty of this basic slice sampling scheme as developed in [30] is determining S^y in Step A2.2. For the SC densities we want to sample from, determining S^y explicitly is not always feasible, and robust numerical approaches to compute it are difficult to design. For instance, using non-convex prior energies such as in ℓ_p^q priors with $q = p$, $p < 1$ leads to multi-modal SC densities and S_y may not be a single interval. Therefore, we will use a generalization of algorithm 2: slice sampling is a variant of *auxiliary variables algorithms* that introduce an additional variable y with a suitable density $p(y|x)$. Then, samples (x^i, y^i) from $p(x, y) = p(x)p(y|x)$ are obtained by a Gibbs sampler, which relies on $p(y|x)$ and $p(x|y)$, and only the x^i are kept. For the basic slice sampler, $p(y|x)$ is chosen as

$$p(y|x) = \frac{1}{p(x)} \mathbb{1}_{\{[0, p(x)]\}}(y), \quad (22)$$

i.e., as a uniform distribution on $[0, p(x^i)]$. We then have

$$p(x, y) = p(x) \frac{1}{p(x)} \mathbb{1}_{\{[0, p(x)]\}}(y) \quad (23)$$

$$p(x|y) \propto \mathbb{1}_{\{[0, p(x)]\}}(y) = \mathbb{1}_{\{x \mid p(x) \geq y\}}(x). \quad (24)$$

If $p(x)$ factorizes to $p(x) \propto p_1(x)p_2(x)$ we can define

$$p(y|x) = \frac{1}{p_2(x)} \mathbb{1}_{\{[0, p_2(x)]\}}(y), \quad (25)$$

which leads to

$$p(x, y) = p(x)p(y|x) = p(x) \frac{1}{p_2(x)} \mathbb{1}_{\{[0, p_2(x)]\}}(y) = p_1(x) \mathbb{1}_{\{[0, p_2(x)]\}}(y), \quad (26)$$

$$p(x|y) = p_1(x) \mathbb{1}_{S_2^y}(x), \quad \text{with } S_2^y := \{z \mid p_2(z) \geq y\}. \quad (27)$$

The corresponding sampler takes the form:

Algorithm 3. (Slice sampling)

For a univariate density $p(x) \propto p_1(x)p_2(x)$, define an initial state x^0 , a burn-in size K_0 and sample size K . For $i = 1, \dots, K_0 + K$ do:
 A3.1. Draw y uniform from $[0, p_2(x^i)]$ (vertical move).
 A3.2. Draw x^{i+1} from $p_1(x)\mathbb{1}_{S_2^y}(x)$ (weighted horizontal move).
 Discard $\{x^i\}_{i=0}^{K_0}$ and use $\{x^i\}_{i=K_0+1}^{K_0+K}$ as a sample of $p(x)$.

For all the methods presented in this section, $p(x)$ does not need to be normalized. Also note that for simplicity, we refer to algorithm 3 as the ‘slice sampler’, hopefully without causing confusion with the one presented in [30], which was included as the ‘basic slice sampler’ (algorithm 2) here for completeness of the presentation.

2.4. Slice-within-Gibbs sampling for Bayesian inversion

The implicit variable split introduced in algorithm 3 is appealing if $S_2^y = \{z \mid p_2(z) \geq y\}$ is a single interval and easy to determine and $p_1(x)$ constrained to an interval is easy to sample from. For the SC posterior densities we consider here, this holds if we split into likelihood plus hard constraints, i.e., $p_1(x) = \exp(-ax^2 + bx)\mathbb{1}_{[lb, ub]}(x)$, and prior parts $p_2(x)$. As the prior terms are unimodal and some even symmetric to zero, S_2^y is a single interval and can be determined easily: for (15), we have $p_2(x) \propto \exp(-c(|x|^p + d)^{q/p})$ and

$$\exp(-c(|x|^p + d)^{q/p}) \geq y \iff |x| \leq \left(\left(-\frac{\log(y)}{c} \right)^{p/q} - d \right)^{1/p}. \quad (28)$$

For the TV prior, (21), we need to compute S_2^y numerically. However, as the energy of $p_2(x)$ is convex, S_2^y is a single interval given by the solutions to $p_2(x) = y$. As the energy of $p_2(x)$ is also piecewise smooth and can be bounded from below, we can easily find starting points for fast, derivative-based root-finding-algorithms. The details are given in appendix A.2. A generalization to other convex, piecewise-smooth energies, such as (17) with suitable p, q , is straight-forward ($p = q = 1$ is a special case as $p_2(x) = y$ can be solved explicitly by a simple scheme). However, if $D^T V$ is dense the number terms in the prior energy is large and this step will become the computational bottleneck of the whole solver. Fortunately, many relevant operators D^T such as finite difference operators or dictionaries composed of local patches are sparse in the original basis, $V = I_n$.

The likelihood part $p_1(x)$ is a Gaussian with $\mu_{SS} = b/(2a)$ and $\sigma_{SS}^2 = 1/(2a)$, truncated to the interval $I = S_2^y \cap [lb, ub]$. For sampling truncated Gaussians, various direct samplers were developed. Our implementation relies on a modified, more robust, version of [9]. Note that if the sampler is initialized in a feasible point $u^0 \in \mathcal{C}$, the probability of I being empty or a single point is zero in theory. In practice, finite precision can lead to $I = \{\tilde{x}\}$, in which case one has to set $x^{i+1} = \tilde{x}$.

Using the slice sampler presented above to sample from $p_{\text{post}}(\cdot \mid \xi_{-j}, f)$ in step A1.2 will be called *slice-within-Gibbs* sampling. In principle, it will generate a full Markov chain

$$\{\xi_j^{i_s}\}_{i_s=K_{s,0}+1}^{K_{s,0}+K_s} \sim p_{\text{post}}(\cdot \mid \xi_{-j}, f), \quad (29)$$

where we subscripted all variables belonging to the inner slice sampler with s . Practically, we only need one sample from $p_{\text{post}}(\cdot \mid \xi_{-j}, f)$. We will always initialize the slice sampler with the current value ξ_i of the component we want to update. Then, we only have to determine the

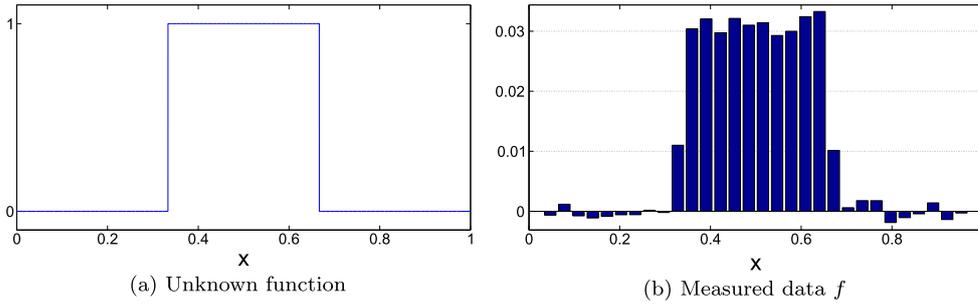


Figure 3. One-dimensional image deblurring scenario ('Boxcar').

length of the burn-in phase $K_{s,0}$ and choose the first sample of the real run as a sample of $p_{\text{post}}(\cdot | \xi_{-j}, f)$, i.e., $K_s = 1$.

The correctness and convergence of the slice-within-Gibbs sampler can be established by combining the properties of the slice sampler (algorithm 3) and the general Gibbs sampler (algorithm 1), which are discussed in [34].

3. Computed examples

3.1. Computational scenarios

'Boxcar'—image deblurring in 1D. For the initial evaluation studies, we use a simple image deblurring scenario in 1D that was adopted from [24] and also used in [27]. It is a simplification of the task to reconstruct a spatially distributed intensity image that is known to consist of piecewise homogeneous parts with sharp edges: the indicator function of $\left[\frac{1}{3}, \frac{2}{3}\right]$ is to be recovered from its integrals over $m = 30$ equidistant subintervals of $[0, 1]$, corrupted by noise with $\mu = 0$, $\Sigma = 10^{-3}I_m$ (see figure 3). The reconstruction is carried out on the grid $x_i^n = \frac{i}{256}$, $i = 1, \dots, n$, with $n = 255$ and the forward operator is discretized by the trapezoidal quadrature rule applied to that grid. Further details can be found in section 3.1.1 of [27].

The prior operator D^T will be given by the forward difference operator with Neumann boundary conditions:

$$D_i = e_{i+1} - e_i, \quad \implies D_i^T u = u_{i+1} - u_i, \quad i = 1, \dots, n-1. \quad (30)$$

D^T has full rank $h = n - 1$ and $V \in \mathbb{R}^{n \times n}$ given by

$$V_{(i,j)} = \begin{cases} 1 & \text{if } i \geq j, \\ 0 & \text{else} \end{cases} \quad (31)$$

is a basis matrix V such that $D^T V$ is a diagonal matrix. We will refer to priors based on this operator as *increment priors*. For the ℓ_1 increment prior, i.e., the conventional TV prior, we can also use the direct ℓ_1 sampler to sample from the posterior. By this, we can validate the approximation of the direct SC sampling via iCDF by the slice sampler proposed here. We will refer to this setting as the 'Boxcar' scenario in the following.

'Phantom-CT'—CT inversion in 2D. We consider an example of 2D *sparse angle* CT to demonstrate the potential of the proposed sampler for real-world applications. An approximate model of CT is given by the *Radon transform* \mathcal{R} : for a 2D function

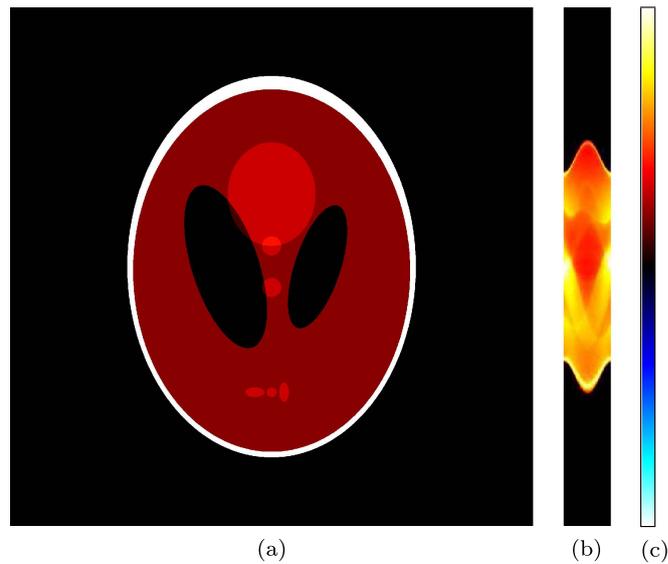


Figure 4. ‘Phantom-CT’ scenario. (a) Unknown function to recover. (b) Clean measurement data (‘sinogram’) for $m_s = 500$, $m_\theta = 45$. (c) Colormap used for all visualizations in this scenario, 0 corresponds to black.

$u^\infty \in L_2([-1, 1]^2)$, it computes integrals along straight lines which are parametrized by the angle θ of their normal vector and their (signed) distance s to the origin:

$$\begin{aligned} \mathcal{R}[u^\infty](\theta, s) &= \int_{l(\theta, s)} u^\infty(x(t), y(t)) dl(t) \\ &= \int_{-\infty}^{\infty} u^\infty(t \sin \theta + s \cos \theta, -t \cos \theta + s \sin \theta) dt. \end{aligned} \quad (32)$$

In sparse angle tomography, only a small number of such angular projections can be measured. In our study, we chose only $m_\theta = 45$ angles, evenly distributed in $[0, \pi)$. In addition, for a given angle θ_i , we practically only measure the integrals of $\mathcal{R}[u^\infty](\theta_i, s)$ over small s -intervals representing an array of $m_s = 500$ equal sized sensor pixels. In total, this leads to $m = m_s \cdot m_\theta = 18.000$ measurements. The forward operator A corresponds to the exact discretization of this measurement with respect to the pixel basis: all the operations involved in the measurement can be computed explicitly for indicator functions of rectangular sets. Further details of this step can be found in section 2.3 in [28].

The unknown function u^∞ to recover is a slightly scaled version of the *Shepp–Logan phantom* [35], a toy model of the human head defined by 10 ellipses. Figure 4(a) shows u^∞ and figure 4(b) the measurement data generated by discretizing u^∞ with a 768×768 pixel grid. We will refer to this scenario as ‘Phantom-CT’.

3.2. Accuracy assessment

To validate that the developed slice sampler accurately reproduces the distributions it is supposed to sample from, the convergence of the sample histograms to the underlying SC densities was checked for visually. Various (random) combinations of coefficients for the different SC densities were tested; see figure 5 for an example of such a comparison.

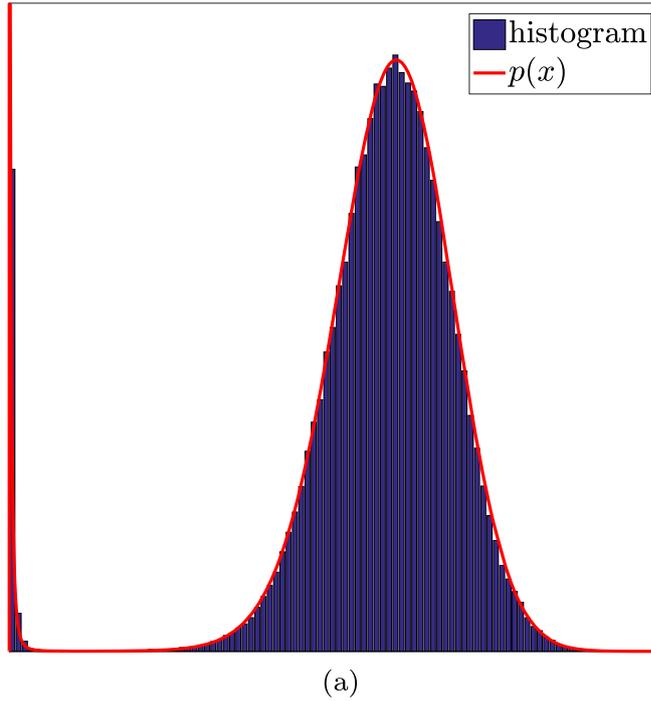


Figure 5. (a) Histogram (blue bars) of the slice sampler compared to targeted SC density (red line) given by (16) with $p = 0.8$. The parameters a and b were picked from a run of the direct ℓ_1 sampler applied to a 2D image deblurring scenario (described in [27]) and c was matched to the regularization parameter used therein.

3.3. Efficiency assessment

Once the accuracy of the slice sampler is established, the next crucial question is whether its use within a Gibbs sampler is *efficient*: in algorithm 1, we ideally want to replace the current values of the component j , u_i by a values that is both distributed following $p_{\text{post}}(\cdot | u_{-j}, f)$ and statistically independent of the current value u_i . While direct SC samplers, such as the iCDF, naturally fulfill these requirements, SC samplers relying on MCMC chains initialized with u_i fulfill them only asymptotically, in the limit $K_{s,0} \rightarrow \infty$. Using a fixed chain size $K_{s,0}$ will inevitably introduce additional correlation between subsequent samples and lower the *statistical efficiency* of slice-within-Gibbs samplers compared to Gibbs sampling relying on a direct sampler for the SC densities. In the following, we will asses this loss of statistical efficiency by autocorrelation analysis.

Autocorrelation analysis. Evaluating samplers in general rather than for a specific aim is a difficult task [26]. For the sake of a concise presentation, a detailed introduction and discussion is omitted here but can be found in section 4.1.6. of [28]. In this study, we will only examine the autocorrelation functions of the MCMC chains projected onto a test function $v \in \mathbb{R}^n$, i.e., of the chain

$$\{g^i\}_{i=1}^K = \{\langle v, u^i \rangle\}_{i=K_0+1}^{K_0+K}. \quad (33)$$

In the ‘Boxcar’ scenario, v is given as the largest eigenvector of the (pre-computed) posterior covariance matrix while in the ‘Phantom-CT’ scenario, it is the indicator function of the area

Table 1. Comparison of τ_{int} for direct and slice-within-Gibbs samplers using different burn-in lengths for the slice sampler. The ‘Boxcar’ scenario and a TV prior ($p = q = 1$) with $\lambda = 400$ is used and $K = 5 \times 10^6$, $\text{SSR} = n$.

$K_{s,0} = 10$	$K_{s,0} = 20$	$K_{s,0} = 40$	$K_{s,0} = 100$	$K_{s,0} = 200$	Direct
231.4 ± 8.6	149.2 ± 4.6	109.4 ± 2.9	102.0 ± 2.6	101.3 ± 2.6	97.8 ± 2.5

defined by $[-0.32, -0.12] \times [0.12, 0.32]$ (this area corresponds to the green box shown in figures 7(e)–(f)). To extract a quantitative measure from the autocorrelation functions, we will estimate their *integrated autocorrelation time* τ_{int} by the approach presented in [41]. In all computed examples, the component j to update in step A1.1 of algorithm 1 is drawn uniformly at random, (*random scan Gibbs sampler*) and a *sub-sampling rate* (SSR) of n is used, i.e., only every n th sample of the chains generated by algorithm 1 is actually stored and τ_{int} refers to the samples of this thinned chain. This means that, on average, we update all n components of $u \in \mathbb{R}^n$ between two steps of the chain (*full sweep*). In each scenario, the samplers were given a large burn-in time K_0 and K was chosen large enough to obtain sufficiently tight error bounds on τ_{int} [41].

Results. When using a conventional TV prior ($p = q = 1$) in the ‘Boxcar’ scenario, the direct ℓ_1 sampler using the iCDF method can be used as a reference to which the slice samplers can be compared to: the τ_{int} obtained by the direct sampler is a lower bound for all slice samplers. Table 1 lists the results. One can observe that already for small MCMC chain length $K_{s,0}$, the differences between direct and slice sampler in terms of statistical efficiency are negligible in practice. Similar examinations using ℓ_2 priors (where, again, a direct sampler can be used as a reference) showed that in this case, significant differences vanish for even smaller values of $K_{s,0}$ (results omitted here).

Tables 2(a), (b) and (c) show the results of similar examinations for an ℓ_p prior with $p = 1.2$, an ℓ_p^q prior with $p = 1$, $q = 10$ and the isotropic TV prior in the 2D ‘Phantom-CT’ scenario (using $n = 129 \times 129$), respectively. While we do not have a direct sampler as a reference here, one can clearly see that τ_{int} is converging to a limit for increasing $K_{s,0}$. In some cases, even using $K_{s,0} = 0$, i.e., only performing one step of the slice sampler, might be sufficient.

Computational complexity. In typical large scale inverse problems such as the one examined in section 3.5, the computational bottleneck is to compute the coefficients of the SC densities, not the process of sampling from them. Therefore, the computational complexity of the slice sampler is not a critical aspect of the whole algorithm. However, to give an indication of how increasing $K_{s,0}$ effects the total run time, table 3 compares the run time of the slice-within-Gibbs sampler to the direct ℓ_1 sampler and table 4 lists the run times of the slice-within-Gibbs sampler for the TV prior in 2D. While the implementation of the slice sampling part is more complicated in this situation (see appendix A.2), we see that even for a moderate sized scenario ($n = 256^2$) it does not significantly effect the run time. Therefore, one does not have to compromise statistical efficiency by choosing a small $K_{s,0}$ to obtain a better computational efficiency.

3.4. Application to the ℓ_p increment prior

Problems with the conventional TV prior (see [24] and the overview in section 5.2.2 in [28]) stimulated research into alternative, edge-preserving prior models. Here, we exemplify how the new slice-within-Gibbs sampler can be used to investigate such general questions in

Table 2. Comparison of τ_{int} for slice-within-RSG samplers using different burn-in lengths for the slice sampler. The ‘Boxcar’ scenario is used in (a) with an ℓ_p increment prior with $p = 1.2$, $\lambda = 400$ and in (b) with an ℓ_p^q increment prior with $p = 1$, $q = 10$, $\lambda = 0.02$. In both cases $K = 2 \times 10^6$ samples were drawn. In (c), the ‘Phantom-CT’ scenario with an isotropic TV prior with $\lambda = 500$ is used and $K = 2 \times 10^5$ samples were drawn.

(a)	$K_{s,0} = 1$	$K_{s,0} = 2$	$K_{s,0} = 4$	$K_{s,0} = 8$	$K_{s,0} = 16$	$K_{s,0} = 32$	$K_{s,0} = 64$
	41.9 ± 1.1	33.3 ± 0.8	23.4 ± 0.5	18.3 ± 0.3	15.8 ± 0.4	14.6 ± 0.3	14.8 ± 0.3
(b)	$K_{s,0} = 1$	$K_{s,0} = 2$	$K_{s,0} = 4$	$K_{s,0} = 8$	$K_{s,0} = 16$	$K_{s,0} = 32$	$K_{s,0} = 64$
	638 ± 46	425 ± 26	307 ± 16	198 ± 9	161 ± 6	155 ± 7	135 ± 6
(c)	$K_{s,0} = 0$	$K_{s,0} = 1$	$K_{s,0} = 2$	$K_{s,0} = 4$	$K_{s,0} = 8$	$K_{s,0} = 16$	$K_{s,0} = 32$
	6.0 ± 0.3	5.3 ± 0.3	5.3 ± 0.3	5.6 ± 0.3	5.2 ± 0.3	4.9 ± 0.3	5.2 ± 0.3

Bayesian inversion: We use it to compute both MAP and CM estimates for the ℓ_p increment prior with p decreasing from $p = 2$ (Gaussian prior) to $p = 1$ (TV prior) and even below $p < 1$ (non-logconcave prior). While the computation of the CM estimates is straight-forward, computing MAP estimates is done by using the sampler within a *simulated annealing* (SA) scheme, a stochastic meta-heuristic for global optimization. The details and an evaluation of using SA together with the proposed SC Gibbs samplers can be found in sections 4.2.4 and 5.1.5 in [28]. In both cases, $K = 10^5$ samples were drawn with $\text{SSR} = n$.

The results of computing MAP and CM estimates for different values of p are shown in figure 6. Here, λ was chosen such that all likelihood energies are equal and that $\lambda = 200$ for $p = 1$. The results suggest that using $p < 1$ leads to superior results for both MAP and CM estimates compared to $p = 1$. The MAP estimate is closer to the real solution as it is both sparser in the increment basis and the contrast loss is reduced. The CM estimate for $p = 0.8$ looks way more convincing compared to those for $p \geq 1$: it has clear pronounced edges that separate smooth, denoised parts. However, using the slice-within Gibbs samplers for $p < 1$ needs to be examined more carefully: while the results are visually convincing, we cannot be sure that the sampler explored the whole, possibly multimodal posterior and did not get stuck in a single mode.

3.5. Application to CT inversion with TV priors

Figure 7 shows MAP and CM estimates for the ‘Phantom-CT’ scenario using an isotropic TV prior with Neumann boundary conditions (19). Here, the MAP estimates were computed with the *alternating direction method of multipliers* [5]. In the Gibbs sampler, *oriented over-relaxation* (see [29] and section 4.3.1. in [28]) was used to accelerate convergence and $K = 2.5 \times 10^4$, 10^4 , 5×10^3 samples were drawn for $n = 64^2$, 128^2 , 256^2 , respectively ($\text{SSR} = n$).

Non-negativity constraints. As the slice-within-Gibbs sampler can easily incorporate additional hard constraints (10), it can be used to quantify their effect on the posterior $p_{\text{post}}(u|f)$. Figure 8 shows CM and CStd estimates computed with or without non-negativity constraints, $u \geq 0$. While the CM estimates look very similar, the CStd estimates reveal that

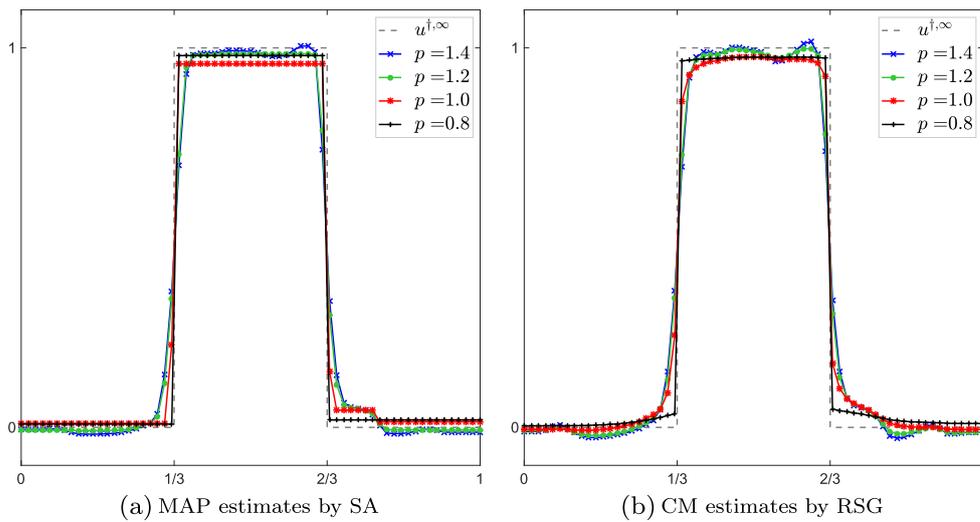


Figure 6. MAP and CM estimates for the 1D ‘Boxcar’ scenario using the ℓ_p increment prior and $n = 63$. The true function to recover (gray line plot) is denoted by $u^{\dagger, \infty}$.

Table 3. Total run time of the slice-within-Gibbs sampler using different burn-in lengths divided by the run time of the direct ℓ_1 sampler. The ‘Boxcar’ scenario and an ℓ_p increment prior with $p = 1.2$, $\lambda = 400$ is used.

$K_{s,0} = 10$	$K_{s,0} = 20$	$K_{s,0} = 40$	$K_{s,0} = 100$	$K_{s,0} = 200$
1.30	1.38	1.48	1.75	2.20

Table 4. Total run time of the slice-within-Gibbs sampler using different burn-in lengths $K_{s,0}$ divided by the run time for $K_{s,0} = 0$. The ‘Phantom-CT’ scenario ($n = 256 \times 256$) and a TV prior ($p = 1$) with $\lambda = 500$ is used.

$K_{s,0} = 0$	$K_{s,0} = 1$	$K_{s,0} = 2$	$K_{s,0} = 4$	$K_{s,0} = 8$	$K_{s,0} = 16$	$K_{s,0} = 32$
1.00	0.98	1.04	1.03	1.05	1.07	1.11

the non-negativity constraints lead to a significant reduction of the posterior variance in some regions.

Gradient estimates. we further present one example of how the samples of u generated by the sampler can be used to compute statistics and uncertainties of a feature of u (see section 1.2): in figure 9(a), we computed the CM estimate of the gradient of the image u

$$\mathbb{E}[\|\nabla u\|_2 | f] = \int \|\nabla u\|_2 p_{\text{post}}(u | f) \, du \quad (34)$$

and figure 9(b) shows the corresponding CStd estimate.

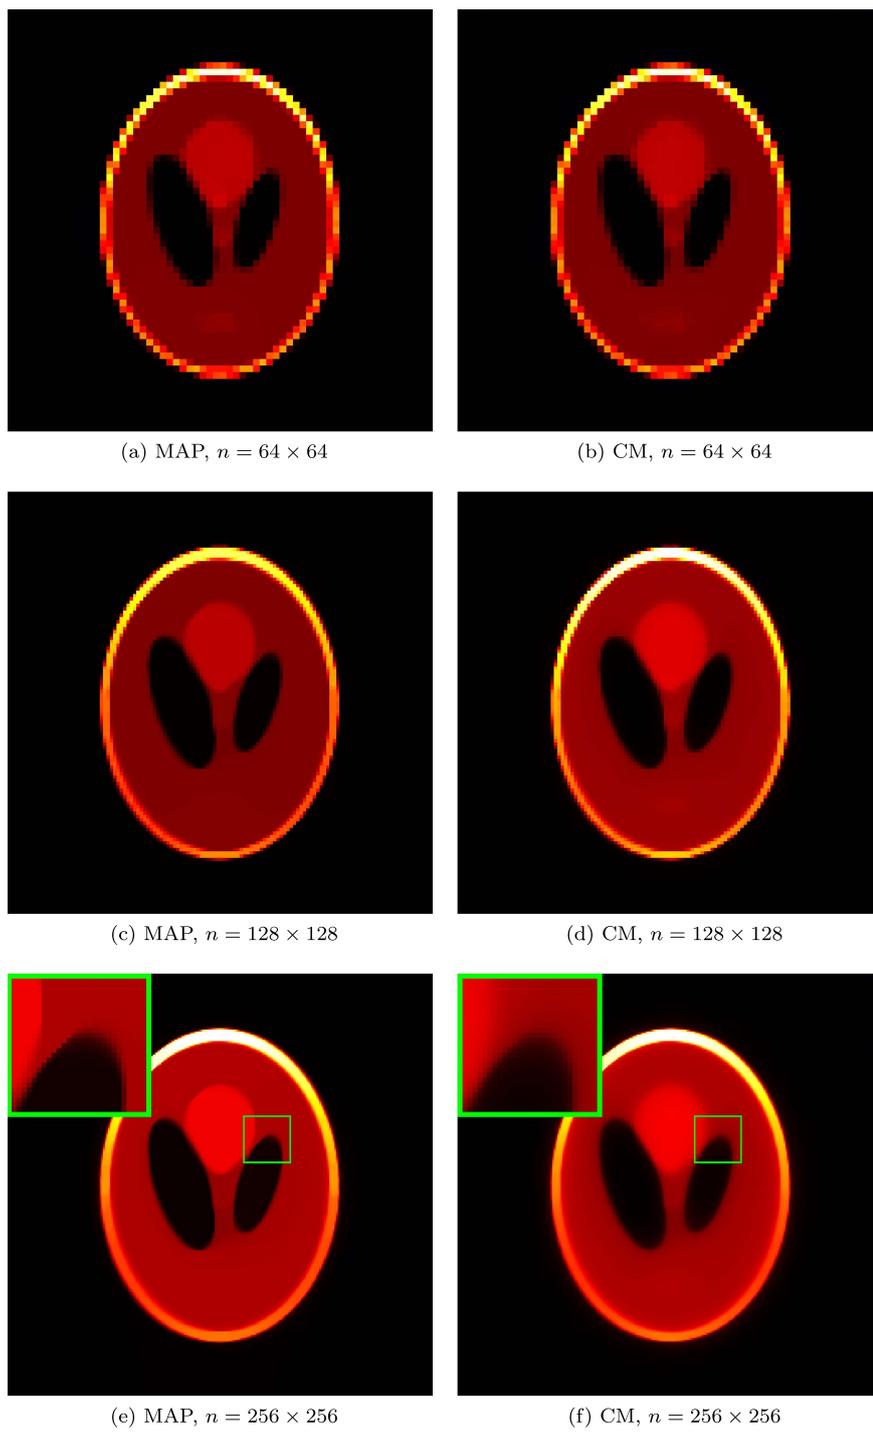


Figure 7. MAP and CM estimates for the 'Phantom-CT' scenario using an isotropic TV prior with $\lambda = 500$, computed with increasing spatial resolution. In the highest resolution, a zoom inset is added.

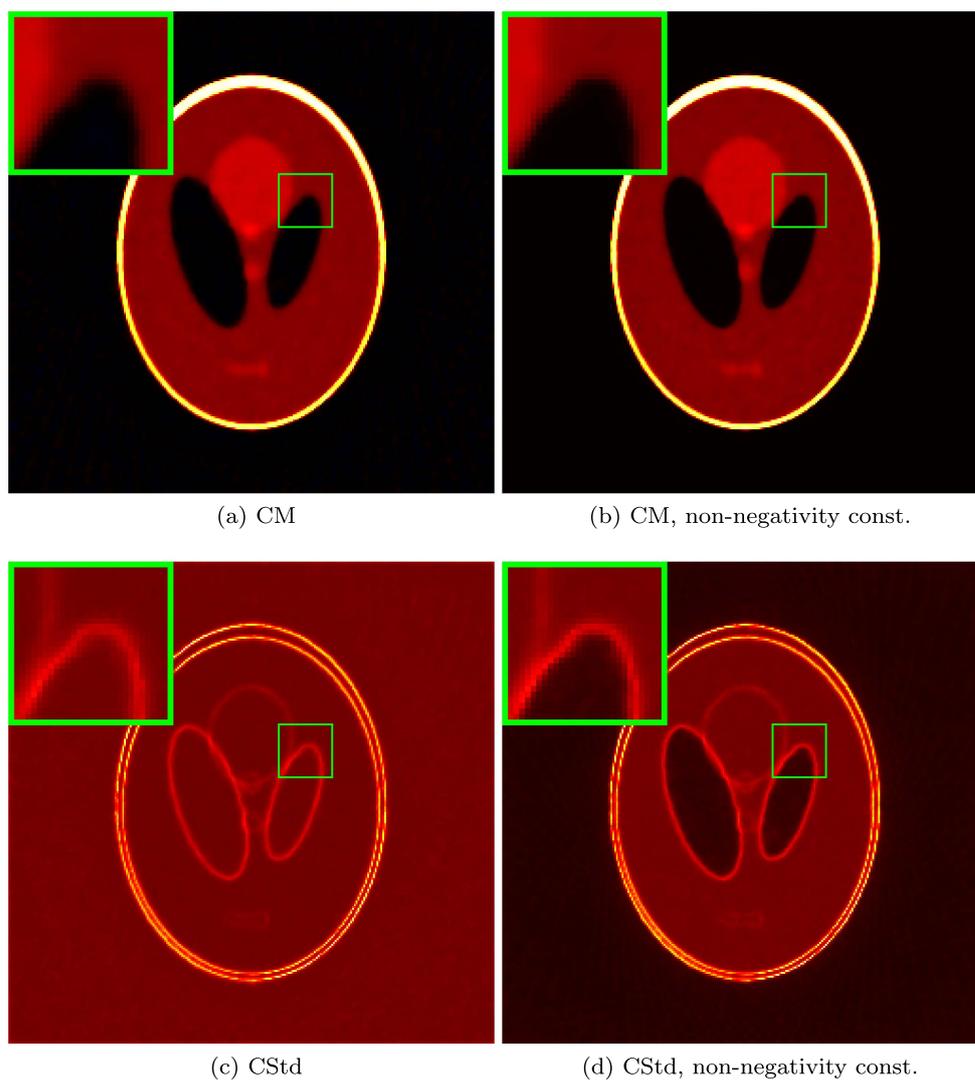


Figure 8. Influence of including non-negativity constraints on CM and CStd estimates in the ‘Phantom-CT’ scenario using an isotropic TV prior with $\lambda = 50$ and a spatial resolution of $n = 256 \times 256$. A zoom inset is added and both CM and both CStd estimates share the same color scale.

4. Discussion and conclusions

In this article, we presented and evaluated a new MCMC sampler that allows us to carry out sample-based Bayesian inversion for a wide range of scenarios and prior models. It is based on the extension of the SC Gibbs-type sampler developed in [27] by a problem-specific adaptation and implementation of generalized slice sampling and enables efficient posterior sampling in high-dimensional scenarios with certain priors for the first time.

The results in sections 3.2 and 3.3 show that using generalized slice sampling to sample from the one-dimensional conditional, SC densities can lead to a fast, robust and accurate

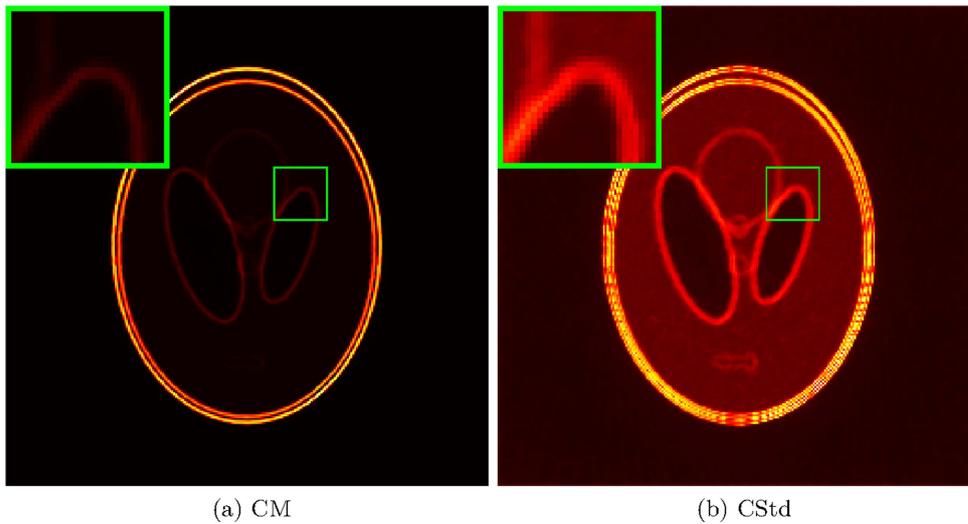


Figure 9. CM and CStd estimates of $\|\nabla u\|_2$ for the ‘Phantom-CT’ scenario using an isotropic TV prior with $\lambda = 50$, non-negativity constraints and a spatial resolution of $n = 256 \times 256$. A zoom inset is added.

posterior sampler for the inverse problems scenario (1) and is therefore an attractive option whenever a fast direct sampler such as iCDF is not available. The computed results in section 3.4 exemplified the use of the new slice-within-Gibbs sampler to examine recent topics in Bayesian inversion and section 3.5 demonstrated how it can lead to interesting results for Bayesian estimation in challenging, high-dimensional inverse problems scenarios. In particular, we examined that TV prior model in 2D: the theoretical analysis of the TV prior carried out, e.g. in [23, 24], is restricted to 1D, only, and, to the best of our knowledge, no theoretical results are available for higher dimensions, yet. The development of the slice-within-Gibbs sampler now enabled us to examine the use the TV prior for the important inverse problem scenario of CT inversion in 2D, for the first time. The results show that, contrary to the 1D case, the CM estimates seem to get smoother for a constant value of λ as the resolution increases. This observation could be the starting point for a new theoretical analysis and has to be examined in higher spatial dimensions by computational studies.

More generally, while our results and those of others (see section 1.2) have demonstrated that sampling high-dimensional posterior distributions is feasible for many important inverse problems scenarios nowadays, an important future challenge lies in extracting the information of interest from the samples generated: while we demonstrated, e.g., how to compute CStd estimates to examine how the posterior variance is influenced by non-negativity constraints (see figure 8) or estimates of a feature $g(u)$ of u (see figure 9), we did not discuss how to interpret the corresponding results. This requires a concrete application and objective and will be topic of future investigations based on the methods presented here.

Related to the last point, we only used simulated data scenarios in this study to focus on the sampling algorithm. The application to experimental CT data will be the subject of a forthcoming publication covering more general aspects of Bayesian inversion in practical applications (see section 5.3. in [28]). Furthermore, only prior models based on ℓ_p^q -norms were considered here, while the sampler can, in principle, be implemented for more general prior models. A more fundamental limitation and future challenge is the current restriction to

linear forward maps (1) and Gaussian noise models (2). Both nonlinear forward maps and non-Gaussian noise models typically conflict with condition (SC1), i.e., they make it very difficult to find an explicit parameterization of the SC densities. In addition, problems related to using SC-Gibbs sampling for multimodal posteriors (see section 3.4) may occur as well.

Code to reproduce all the computed examples will be provided as part of the release of a Matlab-based toolbox for Bayesian inversion.

Acknowledgments

This work was conducted as part of the PhD thesis [28] at the University of Münster. The author would like to thank the two supervisors, Martin Burger and Carsten H Wolters.

Appendix A. Details of the implementation

A.1. Computation of the likelihood coefficients

To implement the SC Gibbs sampler for the ‘Boxcar’ scenario, we compute $\Psi = AV$ and pre-compute $a := \frac{1}{2} \|\Psi_i\|_2^2$. For computing $b = \Psi_i^T \varphi(i) = \Psi_i^T f - (\Psi_i^T \Psi_{-i}) \xi_{-i}$, we pre-compute $\Psi_i^T f$ and $\|\psi_i\|_2^2$ for all i and build the $n \times n$ matrix $\Phi := \Psi^T \Psi$. Then, computing $(\Psi_i^T \Psi_{-i}) \xi_{-i}$ can be performed by using

$$(\Psi_i^T \Psi_{[-i]}) \xi_{[-i]} = \xi^T \Phi_{(:,i)} - \xi_i \|\psi_i\|_2^2, \quad (\text{A.1})$$

which involves a scalar product of dimension n as the most expensive operation.

For the ‘Phantom-CT’ scenario, there are two possible implementations: For the image sizes considered here (n up to 256×256) we can still compute the matrices $\Psi = A$ and $\Phi := \Psi^T \Psi$ explicitly and use the same implementation as in the ‘Boxcar’ scenario ($\Psi = A$ as we stay in the pixel basis, i.e., $V = I_n$). For larger n or 3D applications, we might not be able to store Φ or Ψ . An alternative implementation that does not require storing any matrices uses

$$b = \Psi_i^T f - \Psi_i^T (\Psi \xi) + \xi_i \|\Psi_i\|_2^2 \quad (\text{A.2})$$

to compute b in the following way:

- We again pre-compute $\Psi_i^T f$ and $\|\Psi_i\|_2^2$ for all i . Then, we store the measurement that the current state ξ would cause as f_ξ and initialize it by Au^0 . In principle, f_ξ is given as $\Psi \xi$, and can be directly computed at any time but this computation is too expensive to be performed at every SC update.
- For a given pixel i that is to be updated, we construct Ψ_i and compute the scalar product $\Psi_i^T f_\xi$ to update b by the above formula (note that $\Psi_i^T (f - \Psi \xi)$ is just a projection of Ψ_i onto the current residual of $f_\xi = \Psi \xi$). With the constructed Ψ_i and the change, δ_i , in ξ_i caused by the sampling step, we can then update $f_\xi = f_\xi + \delta_i \Psi_i$.
- While this iterative updating of f_ξ is fast, inaccuracies can accumulate over time, leading to a misfit between f_ξ and $\Psi \xi$. Therefore, we compute $\Psi \xi$ explicitly every n steps and reset f_ξ to this exact value.

The computational bottleneck of this procedure is to compute Ψ_i , i.e., the Radon transform of a pixel (or voxel in 3D). For the parallel beam geometry used here, explicit formulas relying on basic operations that can be parallelized over the angles θ can be derived. For more complicated beam geometries, e.g., the cone beam geometry for 3D reconstruction,

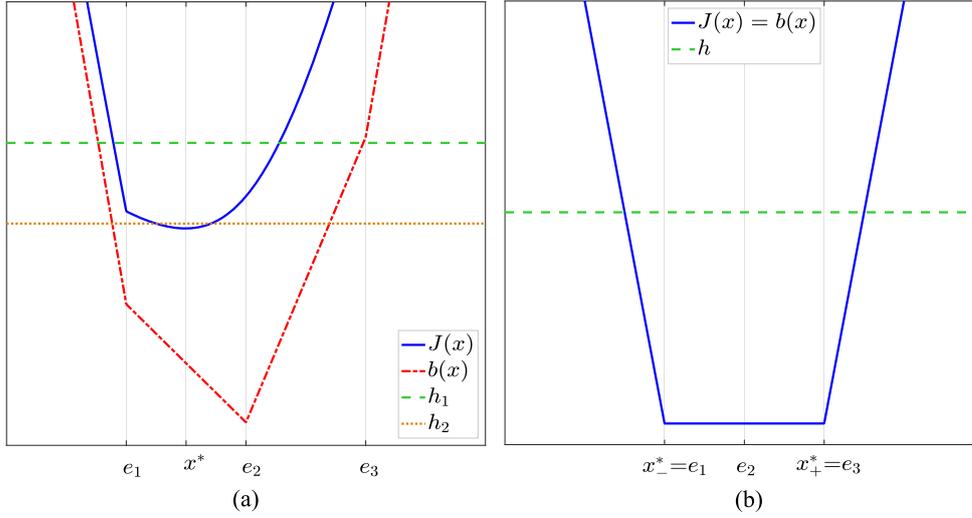


Figure A1. Illustration of two SC density energies for the slice sampler implementation of the TV prior in 2D: $J(x)$ (blue line), $b(x)$ (red line) and h (green and yellow lines) for $(e_1, e_2, e_3) = (-1, 0, 1)$ and (a) $(d_1, d_2, d_3) = (2, 1, 1)$, $(g_1, g_2, g_3) = (0, 0.5, 1)$, (b) $(d_1, d_2, d_3) = (1, 0, 1)$, $(g_1, g_2, g_3) = (0, 0, 0)$.

approximations relying on basic operations from computer graphics can be derived and implemented very efficiently and parallelized on GPUs.

A.2. Slice sampling with TV priors

From (21), we have

$$p_2(x) = \exp\left(-c \sum_{j=1}^3 \sqrt{d_j(x - e_j)^2 + g_j}\right), \quad d_j \in \{0, 1, 2\}, \quad g_j \geq 0 \quad (\text{A.3})$$

and want to determine $S_2^y = \{z \mid p_2(z) \geq y\}$ by solving

$$y = p_2(x) \iff -\frac{\log(y)}{c} = \sum_{j=1}^3 \sqrt{d_j(x - e_j)^2 + g_j}, \quad (\text{A.4})$$

where $y \in (0, p_2(x))$ with probability 1 and $p_2(x) \leq 1$. Assume that $\{e_1, e_2, e_3\}$ are sorted and define $J_j(x) := \sqrt{d_j(x - e_j)^2 + g_j}$ and $h := -\log(y)/c$. Then, $J(x) := \sum_j J_j(x)$ is convex and smooth in $I_1 := (-\infty, e_1)$, $I_2 := (e_1, e_2)$, $I_3 := (e_2, e_3)$ and $I_4 := (e_3, \infty)$. It is monotonic in I_1 and I_4 and is bounded from below by $b(x) := \sum_j \sqrt{d_j}|x - e_j|$. Define $[x_-^*, x_+^*] = \operatorname{argmin} J(x)$ as the interval of minimizers and x_-, x_+ as the solutions to $y = p_2(x)$. We have $x_- < x_-^*$, $x_+ > x_+^*$, $x_- \in I_1 \cup I_2 \cup I_3$ and $x_+ \in I_2 \cup I_3 \cup I_4$ with probability 1 and $[x_-^*, x_+^*] \subset [e_1, e_3]$. See figure A1 for two illustrations.

We will compute x_- by a Newton's method:

$$x_-^i = x_-^{i-1} - \frac{J(x_-^{i-1}) - h}{J'(x_-^{i-1})}, \quad (\text{A.5})$$

initialized in a point x_-^0 such that $x_-^0 \leq x_-$ and $J(x)$ is smooth on $[x_-^0, x_-]$. In each step, the Newton's method approximates $J(x)$ by a tangent in x_-^{i-1} . Due to the convexity of $J(x)$ and $x_-^0 \leq x_- < x_-^*$, the iterates never overshoot: $x_-^0 \leq x_-^i \leq x_-$ for all i . Thereby, they stay in $[x_-^0, x_-]$ and the derivative exists. Finding such an initialization x_-^0 requires some simple considerations:

The subdifferential $\partial J(x)$ is given as the sum of the subdifferentials of $J_i(x)$ (in the set-valued sense of addition):

$$\partial J_j(x) = \begin{cases} \left\{ \frac{d_j(x - e_j)}{\sqrt{d_j(x - e_j)^2 + g_j}} \right\}, & \text{if } x \neq e_j \text{ or } g_j > 0, \\ [-\sqrt{d_j}, \sqrt{d_j}], & \text{if } x = e_j \text{ and } g_j = 0. \end{cases} \quad (\text{A.6})$$

Now, let $J_e^* := \min_j J(e_j)$. We can distinguish two cases:

$h > J_e^*$: in this case, we check the following conditions in sequence:

- If $h > J(e_1)$, x_- is in I_1 . We use the lower bound $b(x)$ to determine x_-^0 such that $b(x_-^0) = h$:

$$x_-^0 = e_1 + \frac{J(e_1) - h}{\sum_j \sqrt{d_j}} \quad (\text{A.7})$$

As $b(x) \leq J(x)$, and both are monotonic in I_1 , we have that $x_-^0 < x_-$.

- Else if $h > J(e_2)$, x_- is in I_2 . We perform one Newton step from e_1 using the maximal subgradient in e_1 :

$$x_-^0 = e_1 - \frac{J(e_1) - h}{\max(\partial J(e_1))} \quad (\text{A.8})$$

This way, $x_-^0 \leq x_-$ and $[x_-^0, x_-] \subset I_2$, i.e., $J(x)$ is differentiable for all iterates.

- Else, $h > J(e_3)$ and x_- is in I_3 . With a similar reasoning, we set

$$x_-^0 = e_2 - \frac{J(e_2) - h}{\max(\partial J(e_2))}. \quad (\text{A.9})$$

For finding x_+ , a similar reasoning can be applied. In the locations of non-differentiability, the minimal subgradient has to be used.

$h < J_e^*$: in this case, $J(x)$ is not piecewise linear (see the yellow line in figure 9(a)) and the unique minimizer x^* is not in $\{e_1, e_2, e_3\}$. The convexity ensures that $x_- < x^* < x_+$ are all either in I_2 or I_3 . If $\max(\partial J(e_1)) < 0$ and $\min(\partial J(e_2)) > 0$ we have that x^* (and thereby x_- and x_+) are in I_2 . Otherwise, they are in I_3 . As above, initial points x_-^0 and x_+^0 fulfilling the conditions can be found by performing one Newton step from the corners of the interval using the maximal/minimal subgradient.

The case $h = J_e^*$ has probability zero.

References

- [1] Agapiou S, Bardsley J M, Papaspiliopoulos O and Stuart A M 2014 Analysis of the Gibbs sampler for hierarchical inverse problems *SIAM/ASA J. Uncertain. Quantification* **2** 511–44
- [2] Bardsley J M 2012 MCMC-Based Image reconstruction with uncertainty quantification *SIAM J. Sci. Comput.* **34** A1316–32

- [3] Bardsley J M and Fox C 2012 An MCMC method for uncertainty quantification in nonnegativity constrained inverse problems *Inverse Problems Sci. Eng.* **20** 477–98
- [4] Bardsley J M, Solonen A, Haario H and Laine M 2014 Randomize-then-optimize: a method for sampling from posterior distributions in nonlinear inverse problems *SIAM J. Sci. Comput.* **36** A1895–910
- [5] Boyd S, Parikh N, Chu E, Peleato B and Eckstein J 2011 Distributed optimization and statistical learning via the alternating direction method of multipliers *Found. Trends Mach. Learn.* **3** 1–22
- [6] Burger M and Lucka F 2014 Maximum *a posteriori* estimates in linear inverse problems with log-concave priors are proper Bayes estimators *Inverse Problems* **30** 114004
- [7] Calvetti D, Kaipio J P and Somersalo E 2014 Inverse problems in the Bayesian framework *Inverse Problems* **30** 110301
- [8] Candes E, Romberg J and Tao T 2006 Robust uncertainty principles: exact signal reconstruction from highly incomplete frequency information *IEEE Trans. Inf. Theory* **52** 489–509
- [9] Chopin N 2011 Fast simulation of truncated Gaussian distributions *Stat. Comput.* **21** 275–88
- [10] Comelli S 2011 A novel class of priors for edge-preserving methods in Bayesian inversion *Master's Thesis* University of Milan, Italy
- [11] Cui T, Fox C and O'Sullivan M J 2011 Bayesian calibration of a large-scale geothermal reservoir model by a new adaptive delayed acceptance Metropolis Hastings algorithm *Water Resour. Res.* **47** W10521
- [12] Cui T, Law K J and Marzouk Y M 2016 Dimension-independent likelihood-informed MCMC *J. Comput. Phys.* **304** 109–37
- [13] Dashti M, Harris S and Stuart A 2012 Besov priors for Bayesian inverse problems *Inverse Problems Imag.* **6** 183–200
- [14] Donoho D L 2006 Compressed sensing *IEEE Trans. Inf. Theory* **52** 1289–306
- [15] Foucart S and Rauhut H 2013 *A Mathematical Introduction to Compressive Sensing* (Basel: Birkhäuser)
- [16] Haario H, Laine M, Mira A and Saksman E 2006 DRAM: efficient adaptive MCMC *Stat. Comput.* **16** 339–54
- [17] Haario H, Saksman E and Tamminen J 2005 Componentwise adaptation for high dimensional MCMC *Comput. Stat.* **20** 265–73
- [18] Hämäläinen K, Kallonen A, Kolehmainen V, Lassas M, Niinimäki K and Siltanen S 2013 Sparse tomography *SIAM J. Sci. Comput.* **35** B644–65
- [19] Helin T and Burger M 2015 Maximum *a posteriori* probability estimates in infinite-dimensional Bayesian inverse problems *Inverse Problems* **31** 085009
- [20] Kaipio J P and Somersalo E 2005 *Statistical and Computational Inverse Problems (Applied Mathematical Sciences vol 160)* (New York: Springer)
- [21] Kolehmainen V, Lassas M, Niinimäki K and Siltanen S 2012 Sparsity-promoting Bayesian inversion *Inverse Problems* **28** 025005
- [22] Lanfer B 2014 Automatic generation of volume conductor models of the human head for EEG source analysis *PhD Thesis* University of Muenster
- [23] Lassas M, Saksman E and Siltanen S 2009 Discretization invariant Bayesian inversion and Besov space priors *Inverse Problems Imag.* **3** 87–122
- [24] Lassas M and Siltanen S 2004 Can one use total variation prior for edge-preserving Bayesian inversion? *Inverse Problems* **20** 1537–63
- [25] Latuszynski K, Roberts G O and Rosenthal J S 2013 Adaptive Gibbs samplers and related MCMC methods *Ann. Appl. Probab.* **23** 66–98
- [26] Liu J 2008 *Monte Carlo Strategies in Scientific Computing (Springer Series in Statistics)* (New York: Springer)
- [27] Lucka F 2012 Fast Markov chain Monte Carlo sampling for sparse Bayesian inference in high-dimensional inverse problems using L1-type priors *Inverse Problems* **28** 125012
- [28] Lucka F 2014 Bayesian inversion in biomedical imaging *PhD Thesis* University of Münster
- [29] Neal R M 1995 Suppressing random walks in Markov chain Monte Carlo using ordered overrelaxation *Technical Report* 9508 Learning in Graphical Models
- [30] Neal R M 2003 Slice sampling *Ann. Stat.* **31** 705–67
- [31] Palafox A, Capistran M A and Christen J A 2014 Effective parameter dimension via Bayesian model selection in the inverse acoustic scattering problem *Math. Problems Eng.* **2014** 12
- [32] Pereyra M 2015 Proximal Markov chain Monte Carlo algorithms *Stat. Comput.* **26** 745–60

- [33] Pereyra M, Schniter P, Chouzenoux E, Pesquet J-C, Tourneret J-Y, Hero A and McLaughlin S 2016 A Survey of stochastic simulation and optimization methods in signal processing *IEEE J. Sel. Top. Signal Process.* **10** 224–41
- [34] Robert C P and Casella G 2005 *Monte Carlo Statistical Methods (Springer Texts in Statistics)* (New York: Springer)
- [35] Shepp L and Logan B 1974 The fourier reconstruction of a head section *IEEE Trans. Nucl. Sci.* **21** 21–43
- [36] Cotter S L, Roberts G O, Stuart A and White D 2013 MCMC methods for functions: modifying old algorithms to make them faster *Stat. Sci.* **28** 424–46
- [37] Sommariva S and Sorrentino A 2014 Sequential Monte Carlo samplers for semi-linear inverse problems and application to magnetoencephalography *Inverse Problems* **30** 114020
- [38] Sorrentino A, Luria G and Aramini R 2014 Bayesian multi-dipole modelling of a single topography in MEG by adaptive sequential Monte Carlo samplers *Inverse Problems* **30** 045010
- [39] Stuart A M 2010 Inverse problems: a Bayesian perspective *Acta Numer.* **19** 451–559
- [40] Vogel C R 2002 *Computational Methods for Inverse Problems* (Philadelphia, PA: Society for Industrial and Applied Mathematics)
- [41] Wolff U 2004 Monte Carlo errors with less errors *Comput. Phys. Commun.* **156** 143–53