Bayesian Inference with Projected Densities

Everink, Jasper Marijn; Dong, Yiqiu; Andersen, Martin Skovgaard

Published in:
SIAM-ASA Journal on Uncertainty Quantification

Link to article, DOI:
10.1137/22M150695X

Publication date:
2023

Document Version
Peer reviewed version

Bayesian Inference with Projected Densities*

Jasper M. Everink†, Yiqiu Dong†, and Martin S. Andersen†

Abstract. Constraints are a natural choice for prior information in Bayesian inference. In various applications, the parameters of interest lie on the boundary of the constraint set. In this paper, we use a method that implicitly defines a constrained prior such that the posterior assigns positive probability to the boundary of the constraint set. We show that by projecting posterior mass onto a polyhedral constraint set, we obtain a new posterior with a rich probabilistic structure on the boundary of that set. If the original posterior is a Gaussian, then such a projection can be done efficiently. We apply the method to Bayesian linear inverse problems, in which case samples can be obtained by repeatedly solving constrained least squares problems, similar to a MAP estimate, but with perturbations in the data. When combined into a Bayesian hierarchical model and the constraint set is a polyhedral cone, we can derive a Gibbs sampler to efficiently sample from the hierarchical model. To show the effect of projecting the posterior, we applied the method to deblurring and computed tomography examples.

Key words. Bayesian inference, constraints, inverse problems, uncertainty quantification

AMS subject classifications. 62F15, 65C05, 90C25

1. Introduction.

The goal of Bayesian inference is to analyze the probability distribution of a parameter obtained through Bayes' theorem [7]. Bayes' theorem states that an initial distribution $\pi(\boldsymbol{x})$ on a parameter $\boldsymbol{x} \in \mathbb{R}^n$ can be updated with the data $\boldsymbol{b}$ using the likelihood function $\pi(\boldsymbol{b}|\boldsymbol{x})$, more precisely, $\pi(\boldsymbol{x}|\boldsymbol{b}) \propto \pi(\boldsymbol{b}|\boldsymbol{x})\pi(\boldsymbol{x})$. In many applications, the variable of interest $\boldsymbol{x}$ must satisfy some constraints: the pixels in images are bounded, the attenuation coefficients in a CT scan are nonnegative and the mass of an object might be bounded from above. This kind of information is often incorporated in the prior distribution $\pi(\boldsymbol{x})$, with the goal of making the posterior more accurate and explainable.

Besides choosing common prior distributions that are naturally restricted to the constraint space, there are multiple other ways to add the constraints to the prior information. One of them is to truncate the prior distribution [8], where an unconstrained prior $\pi(\boldsymbol{x})$ is replaced by a prior proportional to $\pi(\boldsymbol{x})\mathbf{1}_{\mathbf{C}}(\boldsymbol{x})$, where $\mathbf{1}_{\mathbf{C}}(\boldsymbol{x})$ is 1 if $\boldsymbol{x} \in \mathbf{C}$ and 0 otherwise. This is equivalent to truncating the posterior, resulting in a posterior proportional to $\pi(\boldsymbol{x}|\boldsymbol{b})\mathbf{1}_{\mathbf{C}}(\boldsymbol{x})$. Samples from such a truncated posterior can be obtained using MCMC methods. Alternatively, the variable can be reparameterized to an unconstrained space. For example, for a positive variable x, we can write $x = e^z$ and define an unconstrained prior on z. If z is normally distributed, then x is referred to as the log-normal distribution [2]. These methods generally focus on the interior of the constraint set, while in many applications, the signals of interest lie on the boundary of the constraint set. For example, an image with at least

*Submitted to the editors July 2, 2022.

Funding: This work was supported by the Villum Foundation (grant no. 25893) and the Novo Nordisk Foundation (grant no. NNF20OC0061894).

†Department of Applied Mathematics and Computer Science, Technical University of Denmark. Richard Petersens Plads, Building 324, DK-2800 Kgs. Lyngby, Denmark. (jmev@dtu.dk, yido@dtu.dk, mskan@dtu.dk)

one zero-valued pixel already lies on the boundary of the nonnegative orthant and in many applications, a lot of pixels are expected to be zero. However, if the distribution is described by a density, then the probability of having at least one zero-valued pixel is zero.

Here, we consider a method that focuses on the boundary of the constraint set by projecting samples from an unconstrained posterior onto the constraint set. Such methods have been analyzed before in a Bayesian decision-theoretic framework [14]. Furthermore, in the case of Bayesian linear inverse problems with Gaussian likelihood and prior with nonnegativity constraints, it has been noted [3, 4] that samples from a projected posterior can be obtained by sampling from constrained versions of randomized linear least squares problems. More precisely, let $\boldsymbol{x}$ be an unknown signal which is observed through a linear forward operator $A : \mathbb{R}^n \to \mathbb{R}^m$. Examples of linear forward operators include convolutions [13] and CT-scans [1]. The observations are inaccurate measurements of the form $\boldsymbol{b} = A\boldsymbol{x} + \boldsymbol{e}$ with error $\boldsymbol{e} \in \mathbb{R}^m$. The components of the error $\boldsymbol{e}$ are often modeled as independent and identically distributed Gaussian random variables, which results in a likelihood of the form $\pi(\boldsymbol{b}|\boldsymbol{x}) \propto \exp\left(-\frac{\lambda}{2}\|A\boldsymbol{x} - \boldsymbol{b}\|_2^2\right)$. If we model our initial knowledge of $\boldsymbol{x}$ as the prior distribution $\pi(\boldsymbol{x}) \propto \exp\left(-\frac{\delta}{2}\|L\boldsymbol{x}\|_2^2\right)$, then the posterior distribution satisfies

$$\pi(\boldsymbol{x}|\boldsymbol{b}) \propto \pi(\boldsymbol{b}|\boldsymbol{x})\pi(\boldsymbol{x}) \propto \exp\left(-\frac{\lambda}{2}\|A\boldsymbol{x} - \boldsymbol{b}\|_2^2 - \frac{\delta}{2}\|L\boldsymbol{x}\|_2^2\right). \tag{1.1}$$

In [4], they observed that projecting (1.1) onto the nonnegative orthant with respect to the norm $\|\boldsymbol{x}\|_{\lambda A^TA+\delta L^TL} := \sqrt{\boldsymbol{x}^T(\lambda A^TA + \delta L^TL)\boldsymbol{x}}$, assuming $\lambda A^TA + \delta L^TL$ is positive definite, is equivalent to solving the randomized constrained least squares problem

$$\boldsymbol{x}^\star = \operatorname*{argmin}_{\boldsymbol{x}\in\mathbb{R}^n_{\geqslant 0}} \left\{ \frac{\lambda}{2}\|A\boldsymbol{x} - \hat{\boldsymbol{b}}\|_2^2 + \frac{\delta}{2}\|L\boldsymbol{x} - \hat{\boldsymbol{c}}\|_2^2 \right\},$$

where $\hat{\boldsymbol{b}} \sim \mathcal{N}(\boldsymbol{b}, \lambda^{-1}I)$ and $\hat{\boldsymbol{c}} \sim \mathcal{N}(\boldsymbol{0}, \delta^{-1}I)$. This approach results in additional computational and theoretical benefits compared the more general study of [14].

In this paper, we generalize the approach from [3, 4] to general constraints, with a focus on polyhedral sets, and analyze the projected Gaussian posterior obtained through solving randomized constrained least squares problems. We derive a characterization for the projected distribution and show how this theory can be applied to Bayesian linear inverse problems. We also discuss how the projected posterior relates to a constrained prior and derive a Gibbs sampler for when the constraints are polyhedral cones. Finally, we use numerical experiments to look into the effect of the projection to Bayesian linear inverse problems.

This paper is organized as follows. In Section 2, we discuss the theory behind the oblique projection of a multivariate Gaussian distribution onto a closed and convex set. In Section 3, we discuss how to apply this projection framework to uncertainty quantification for linear inverse problems. Finally, in Section 4, we discuss the projection framework applied to deblurring and computer tomography.

2. Projected multivariate Gaussian distribution.

In this section, we describe a framework of projecting Gaussian distributions onto closed convex sets, where the projection is with respect to the norm induced by the precision matrix of the Gaussian. We describe theory for

polyhedral constraint sets and give some explicit descriptions of the projected Gaussian for a few simple constraint sets.

2.1. Oblique projection of Gaussian distributions. Throughout this section, unless otherwise stated, we make the following assumption.

Assumption 2.1. *The random vector $\boldsymbol{x}^\star \in \mathbb{R}^n$ follows a Gaussian distribution with mean $\boldsymbol{\mu} \in \mathbb{R}^n$ and covariance matrix $\Sigma \in \mathbb{S}^n_{++}$, where $\mathbb{S}^n_{++}$ denotes the set of $n \times n$ symmetric positive definite matrices.*

For $\boldsymbol{x}^\star$ as defined by Assumption 2.1, we can equivalently write $\boldsymbol{x}^\star = \boldsymbol{\mu} + \Sigma\hat{\boldsymbol{w}}$, where $\hat{\boldsymbol{w}} \sim \mathcal{N}(\mathbf{0}, \Sigma^{-1})$. Define the quadratic function $g(\boldsymbol{x}; \boldsymbol{w}) := \frac{1}{2}\boldsymbol{x}^T\Sigma^{-1}\boldsymbol{x} - \boldsymbol{x}^T(\Sigma^{-1}\boldsymbol{\mu} + \boldsymbol{w})$, then $\boldsymbol{x}^\star$ satisfies $\nabla_{\boldsymbol{x}} g(\boldsymbol{x}^\star; \hat{\boldsymbol{w}}) = \mathbf{0}$. Therefore, $\boldsymbol{x}^\star$ is the solution to a randomized quadratic optimization problem of the form

$$\boldsymbol{x}^\star = \operatorname*{argmin}_{\boldsymbol{x} \in \mathbb{R}^n} \left\{ \frac{1}{2}\boldsymbol{x}^T\Sigma^{-1}\boldsymbol{x} - \boldsymbol{x}^T(\Sigma^{-1}\boldsymbol{\mu} + \hat{\boldsymbol{w}}) \right\}, \quad \hat{\boldsymbol{w}} \sim \mathcal{N}(\mathbf{0}, \Sigma^{-1}). \tag{2.1}$$

Thus, we can sample from $\boldsymbol{x}^\star$ by repeatedly solving optimization problem (2.1) for different samples of $\hat{\boldsymbol{w}}$.

The constrained optimization problem (2.1) restricted to a closed convex set $\mathbf{C} \subseteq \mathbb{R}^n$. The resulting optimization problem has the form

$$\boldsymbol{z}^\star = \operatorname*{argmin}_{\boldsymbol{z} \in \mathbf{C}} \left\{ \frac{1}{2}\boldsymbol{z}^T\Sigma^{-1}\boldsymbol{z} - \boldsymbol{z}^T(\Sigma^{-1}\boldsymbol{\mu} + \hat{\boldsymbol{w}}) \right\}, \quad \hat{\boldsymbol{w}} \sim \mathcal{N}(\mathbf{0}, \Sigma^{-1}),$$

or equivalently

$$\boldsymbol{z}^\star = \operatorname*{argmin}_{\boldsymbol{z} \in \mathbf{C}} \frac{1}{2}\|\boldsymbol{z} - (\boldsymbol{\mu} + \Sigma\hat{\boldsymbol{w}})\|^2_{\Sigma^{-1}} = \operatorname*{argmin}_{\boldsymbol{z} \in \mathbf{C}} \frac{1}{2}\|\boldsymbol{z} - \boldsymbol{x}^\star\|^2_{\Sigma^{-1}},$$

where $\|\boldsymbol{x}\|_{\Sigma^{-1}} := \sqrt{\boldsymbol{x}^T\Sigma^{-1}\boldsymbol{x}}$, i.e., $\boldsymbol{z}^\star$ is the oblique projection of $\boldsymbol{x}^\star \sim \mathcal{N}(\boldsymbol{\mu}, \Sigma)$ onto $\mathbf{C}$ with respect to the norm induced by the precision matrix Σ^{-1}. This obliquely projected Gaussian distribution is the main object of study in this work and is summarized in the following definition.

Definition 2.2. *Under Assumption 2.1, the oblique projection of $\mathcal{N}(\boldsymbol{\mu}, \Sigma)$ onto a closed, convex set $\mathbf{C} \subset \mathbb{R}^n$ is the oblique projection of $\boldsymbol{x}^\star$ onto the set $\mathbf{C}$ with respect to the norm induced by the precision matrix Σ^{-1}, that is,*

$$\boldsymbol{z}^\star = \Pi^{\Sigma^{-1}}_{\mathbf{C}}(\boldsymbol{x}^\star) := \operatorname*{argmin}_{\boldsymbol{z} \in \mathbf{C}} \frac{1}{2}\|\boldsymbol{x}^\star - \boldsymbol{z}\|^2_{\Sigma^{-1}}. \tag{2.2}$$

Because the precision matrix Σ^{-1} is positive definite, the oblique projection onto a closed, convex set $\Pi^{\Sigma^{-1}}_{\mathbf{C}}$ is well-defined and continuous, hence measurable. Therefore, the random vector (2.2) is well defined with distribution

$$\mathbb{P}(\Pi^{\Sigma^{-1}}_{\mathbf{C}}(\boldsymbol{x}^\star) \in \mathbf{E}) = \mathbb{P}\left(\boldsymbol{x}^\star \in \left[\Pi^{\Sigma^{-1}}_{\mathbf{C}}\right]^{-1}(\mathbf{E})\right), \tag{2.3}$$

for a measurable set $\mathbf{E}$ and where $\left[\Pi_{\mathbf{C}}^{\Sigma^{-1}}\right]^{-1}$ denotes the inverse image of $\Pi_{\mathbf{C}}^{\Sigma^{-1}}$.

Solving the optimization problem (2.2) can be computationally expensive, however, the inverse mapping can be analyzed more easily. The optimality condition of the oblique projection is given by

$$\mathbf{0} \in \Sigma^{-1}(\boldsymbol{z}^\star - \boldsymbol{x}^\star) + N_{\mathbf{C}}(\boldsymbol{z}^\star),$$

or equivalently

$$\boldsymbol{x}^\star \in \boldsymbol{z}^\star + \Sigma N_{\mathbf{C}}(\boldsymbol{z}^\star),$$

where $N_{\mathbf{C}}(\boldsymbol{z}) = \partial I_{\mathbf{C}}(\boldsymbol{z}) = \{\boldsymbol{v} \,|\, \boldsymbol{v}^T(\boldsymbol{y} - \boldsymbol{z}) \leqslant 0, \forall \boldsymbol{y} \in \mathbf{C}\}$ is the normal cone associated with $\mathbf{C}$.

From optimality condition (2.1) and equation (2.3) we get that the distribution of the projected Gaussian can be described by

$$\mathbb{P}\left(\Pi_{\mathbf{C}}^{\Sigma^{-1}}(\boldsymbol{x}^\star) \in \mathbf{E}\right) = \mathbb{P}\left(\boldsymbol{x}^\star \in \bigcup_{\boldsymbol{z}\in\mathbf{E}} [\boldsymbol{z} + \Sigma N_{\mathbf{C}}(\boldsymbol{z})]\right), \tag{2.4}$$

where $\mathbf{E}$ is a measurable set.

Because of the projection, a lot of the mass gets projected onto the boundary of the constraint set. If $\mathbf{C}$ is a polyhedral set, then the probability on the boundary can be described as in the following lemma.

Lemma 2.3. *Under Assumption 2.1, if $\mathbf{C}$ is a polyhedral set and $\mathbf{F}$ is a face of $\mathbf{C}$, then for any measurable set $\mathbf{E} \subseteq \operatorname{relint}(\mathbf{F})$, where $\operatorname{relint}(\mathbf{F})$ denotes the relative interior of the face $\mathbf{F}$, we have*

$$\mathbb{P}\left(\Pi_{\mathbf{C}}^{\Sigma^{-1}}(\boldsymbol{x}^\star) \in \mathbf{E}\right) = \int_{\mathbf{E}} \pi_{\mathbf{F}}(\boldsymbol{z}) \mathrm{d}\boldsymbol{z},$$

where $\pi_{\mathbf{F}}(\boldsymbol{z})$ is the $\dim(\mathbf{F})$-dimensional density

$$\pi_{\mathbf{F}}(\boldsymbol{z}) = \int_{\Sigma N_{\mathbf{C}}(\boldsymbol{z})} \pi_{\boldsymbol{x}^\star}(\boldsymbol{z} + \boldsymbol{v}) \mathrm{d}\boldsymbol{v}.$$

Proof. For any $\boldsymbol{z} \in \operatorname{relint}(\mathbf{F})$, the normal cone $N_{\mathbf{C}}(\boldsymbol{z})$ is independent of $\boldsymbol{z}$, therefore we can write (2.4) as

$$\begin{aligned}\mathbb{P}\left(\Pi_{\mathbf{C}}^{\Sigma^{-1}}(\boldsymbol{x}^\star) \in \mathbf{E}\right) &= \int_{\mathbf{E}+\Sigma N_{\mathbf{C}}} \pi_{\boldsymbol{x}^\star}(\boldsymbol{w}) \mathrm{d}\boldsymbol{w} \\ &= \int_{\mathbf{E}} \int_{\Sigma N_{\mathbf{C}}} \pi_{\boldsymbol{x}^\star}(\boldsymbol{z} + \boldsymbol{v}) \mathrm{d}\boldsymbol{v} \mathrm{d}\boldsymbol{z},\end{aligned}$$

where the final decomposition is valid due to Lemma A.1 in Appendix A.1. ∎

Lemma 2.3 above states that the projected Gaussian distribution onto a polyhedral set consists of a mixture of various densities of different dimensions on all of the faces of the polyhedral set.

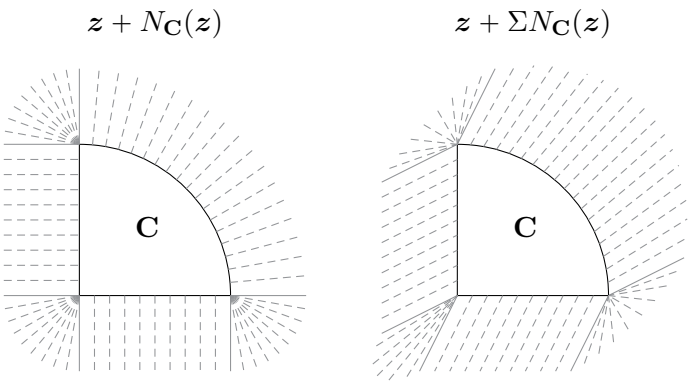

Figure 1: Visualization of obliquely projecting a density onto the different boundary components of a quarter disc.

Extending Lemma 2.3 to general closed, convex sets $\mathbf{C}$ is a more complicated procedure, but the idea of densities of different dimensions on different parts of the constraint set is the same. For example, consider the quarter disc as illustrated in Figure 1. The projection has no effect on the interior of the domain, therefore, there is a two-dimensional density on the interior of the set. The normal cones at the three corners are two dimensional, therefore, positive mass gets projected onto each of the corners. The normal cone at the straight and curved lines are one-dimensional, resulting in one-dimensional densities on these parts. Figure 2, shows these densities, except for the corner points. To further illustrate this, we will consider a few examples for which we can compute analytical expression for the densities.

2.2. Analytic examples. Computing analytical expressions for the densities on parts of the surface of the constraint set is in general intractable, but can be computed when the dimension of the normal cone $N_{\mathbf{C}}(\boldsymbol{x})$ is at most one-dimensional. In this subsection, we consider two such cases, half-space and disc constraints, and give analytical expressions for their densities.

2.2.1. Halfspace. Suppose that $\mathbf{C}$ is a halfspace defined by $\mathbf{C} = \{\boldsymbol{x} \in \mathbb{R}^n \,|\, \boldsymbol{a}^T\boldsymbol{x} \leqslant b\}$, where $\boldsymbol{a} \in \mathbb{R}^n$ is a nonzero normal vector and $b \in \mathbb{R}$. Denote by $F \in \mathbb{R}^{n\times(n-1)}$ a matrix whose columns form an orthonormal basis for $\mathrm{Null}(\boldsymbol{a}^T)$ and let $\boldsymbol{x}_0 \in \mathbb{R}^n$ satisfy $\boldsymbol{a}^T\boldsymbol{x}_0 = b$, then the boundary of the halfspace $\mathbf{C}$ can be parameterized by $\boldsymbol{x}_0 + F\boldsymbol{u}$ for $\boldsymbol{u} \in \mathbb{R}^{n-1}$. Under Assumption 2.1, the $n-1$ dimensional density of $\Pi_{\mathbf{C}}^{\Sigma^{-1}}(\boldsymbol{x}^\star)$ on the boundary of $\mathbf{C}$ is given by

$$\pi_{\mathrm{bd}}(\boldsymbol{u}) = \pi_{\boldsymbol{x}^\star}(\boldsymbol{x}_0 + F\boldsymbol{u})\sqrt{\frac{\boldsymbol{a}^T\Sigma\boldsymbol{a}}{\boldsymbol{a}^T\boldsymbol{a}}}\exp\left(\gamma^2\right)\sqrt{\frac{\pi}{2}}\,\mathrm{erfc}\left(\gamma\right), \tag{2.5}$$

where $\gamma = \frac{2(b-\boldsymbol{a}^T\boldsymbol{\mu})}{\sqrt{8\boldsymbol{a}^T\Sigma\boldsymbol{a}}}$ and erfc is the complementary error function.

2.2.2. Unit disc. Suppose that $\mathbf{C}$ is a unit disc defined by $\mathbf{C} = \{\boldsymbol{x} \in \mathbb{R}^2 \,|\, \|\boldsymbol{x}\|_2 \leqslant 1\}$ with boundary parameterization $\boldsymbol{n}(u) := (\cos(u), \sin(u))^T$ for $u \in [0, 2\pi)$. The density on the

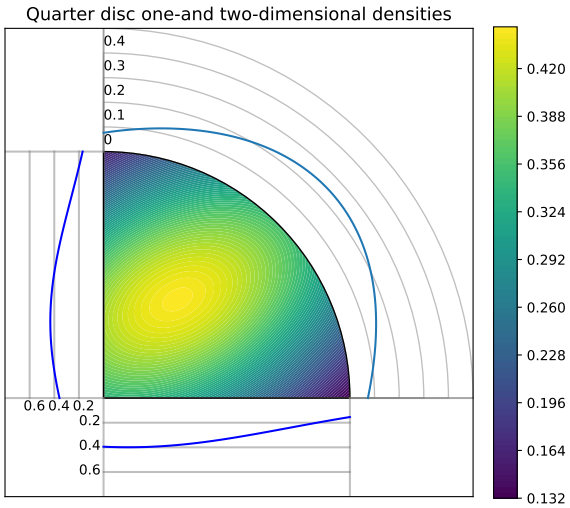

Figure 2: The one-and two-dimensional densities of the oblique projection of a Gaussian onto a quarter disc. The zero-dimensional densities in the corners are not shown.

boundary of the disc is given by

$$\pi_{\text{bd}}(u) = \pi_{\boldsymbol{x}^\star}(\boldsymbol{n}(u)) \left(\frac{\det(\Sigma)}{\alpha} + \frac{\sqrt{\pi}}{\sqrt{8}a^{3/2}} (2\alpha K(u) - \beta \det(\Sigma)) \exp\left(\frac{\beta^2}{8\alpha}\right) \operatorname{erfc}\left(\frac{\beta}{\sqrt{8\alpha}}\right) \right), \tag{2.6}$$

where $\alpha = \boldsymbol{n}(u)^T \Sigma \boldsymbol{n}(u)$, $\beta = 2\boldsymbol{n}(u)^T(\boldsymbol{n}(u) - \boldsymbol{\mu})$, $K(u) = \det[\Sigma \boldsymbol{n}(u)\ R\boldsymbol{n}(u)]$ and $R = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}$.

These computations are all based on being able to compute the integral of a Gaussian over a one-dimensional normal cone and the computations can be found in Appendix A.2. Furthermore, these distributions can be mixed. Figure 2 shows the one-and two-dimensional densities on the boundary of a quarter disc as described by the Equations (2.5) and (2.6).

2.3. Boundary properties. As seen in the analytical examples in Subsection 2.2, a lot of mass is projected onto the boundary of the constraint set $\mathbf{C}$. Nevertheless, there will always be a positive probability on the relative interior of the constraint set. The following lemma shows that both the boundary and relative interior of the constraint set will always have positive probability. Although the lemma is formulated in terms of the oblique projection of a Gaussian distribution with respect to its precision matrix, the proof extends to the projection with respect to any positive definite matrix and any continuous distribution whose support is $\mathbb{R}^n$.

Lemma 2.4. *Under Assumption 2.1, if $\mathbf{C} \subset \mathbb{R}^n$ is a non-empty, closed, convex set, then the probability of being on the boundary of $\mathbf{C}$ is positive and the probability of being in the*

relative interior of $\mathbf{C}$ *is positive, i.e.,*

$$\mathbb{P}(\Pi_{\mathbf{C}}^{\Sigma^{-1}}(\boldsymbol{x}^\star) \in \mathrm{bd}(\mathbf{C})) > 0 \quad \textit{and} \quad \mathbb{P}(\Pi_{\mathbf{C}}^{\Sigma^{-1}}(\boldsymbol{x}^\star) \in \mathrm{relint}(\mathbf{C})) > 0.$$

Proof. Because $\mathbf{C} \subset \mathbb{R}^n$ and closed, $\mathbb{R}^n \backslash \mathbf{C}$ is non-empty and open. Therefore

$$\mathbb{P}(\Pi_{\mathbf{C}}^{\Sigma^{-1}}(\boldsymbol{x}^\star) \in \mathrm{bd}(\mathbf{C})) \geqslant \mathbb{P}(\boldsymbol{x}^\star \in \mathbb{R}^n \backslash \mathbf{C}) > 0.$$

Let $\boldsymbol{z} \in \mathrm{relint}(\mathbf{C})$, then there exists $\epsilon > 0$ such that $B_\epsilon(\boldsymbol{z}) \cap \mathrm{aff}(\mathbf{C}) \subseteq \mathbf{C}$, where $B_\epsilon(\boldsymbol{z})$ is the closed ball of radius ϵ around $\boldsymbol{z}$ and $\mathrm{aff}(\mathbf{C})$ is the affine hull of $\mathbf{C}$. By Lemma A.1, $B_\epsilon(\boldsymbol{z}) \cap \mathrm{aff}(\mathbf{C}) + \Sigma\,\mathrm{aff}(\mathbf{C})^\perp$ is an n-dimensional convex set. Therefore we can conclude that ∎

$$\mathbb{P}(\Pi_{\mathbf{C}}^{\Sigma^{-1}}(\boldsymbol{x}^\star) \in \mathrm{relint}(\mathbf{C})) \geqslant \mathbb{P}(\boldsymbol{x}^\star \in B_\epsilon(\boldsymbol{z}) \cap \mathrm{aff}(\mathbf{C}) + \Sigma\,\mathrm{aff}(\mathbf{C})^\perp) > 0.$$

Typically, one would like to compute a single point estimate from the posterior and analyze the uncertainty of that estimate. A common choice for point estimate is the maximum a posteriori (MAP), which is the point of maximum posterior density. However, due to the projected posterior consisting of a mixture of different dimensional densities, the MAP estimate is undefined. Another common choice for point estimate is the mean of the posterior, however, as the following theorem shows, the mean of the projected posterior lies inside the relative interior of the constraint set. Similarly for the sample mean, a single sample in the relative interior already implies that the sample mean will lie in the relative interior. It is therefore not always a suitable point estimate if one is interested in the boundary of the constraint set. This motivates our choice of the componentwise median as point estimate in the numerical examples in Section 4.

Theorem 2.5. *Under Assumption 2.1, if* $\mathbf{C} \subset \mathbb{R}^n$ *is a non-empty, closed, convex set, then*

$$\mathbb{E}[\Pi_{\mathbf{C}}^{\Sigma^{-1}}(\boldsymbol{x}^\star)] \in \mathrm{relint}(\mathbf{C}).$$

Proof. By Lemma 2.4, $p := \mathbb{P}(\Pi_{\mathbf{C}}^{\Sigma^{-1}}(\boldsymbol{x}^\star) \in \mathrm{relint}(\mathbf{C})) > 0$. Let $\boldsymbol{z}^\star = \Pi_{\mathbf{C}}^{\Sigma^{-1}}(\boldsymbol{x}^\star)$, then consider the decomposition of the expectation on the relative interior, i.e.,

$$\mathbb{E}[\boldsymbol{z}^\star] = p\mathbb{E}[\boldsymbol{z}^\star | \boldsymbol{z}^\star \in \mathrm{relint}(\mathbf{C})] + (1-p)\mathbb{E}[\boldsymbol{z}^\star | \boldsymbol{z}^\star \in \mathrm{relbd}(\mathbf{C})].$$

We have that $\mathbb{E}[\boldsymbol{z}^\star | \boldsymbol{z}^\star \in \mathrm{relint}(\mathbf{C})] \in \mathrm{relint}(\mathbf{C})$ by the convexity of the relative interior and $\mathbb{E}[\boldsymbol{z}^\star | \boldsymbol{z}^\star \in \mathrm{relbd}(\mathbf{C})] \in \mathbf{C}$. Combined with the fact [11, Lemma 2.1.6] that for $p \in (0, 1]$, $p\,\mathrm{relint}(\mathbf{C}) + (1-p)\mathbf{C} \subseteq \mathrm{relint}(\mathbf{C})$, we can conclude that $\mathbb{E}[\boldsymbol{z}^\star] \in \mathrm{relint}(\mathbf{C})$. ∎

2.4. Gaussian decomposition. The two analytical examples in Subsection 2.2 illustrate that the densities on the boundary of the constraint set can become quite complicated, even for relatively simple examples. However, note that the density on the boundary of a half-space in (2.5) is proportional to the original Gaussian density, while on the boundary of a disc in (2.6), the additional factor depends on the surface coordinate. The property that the densities are proportional to the original Gaussian as observed in the half-space example can be generalized to polyhedral sets as stated in the following theorem, which is a key observation for the derivation of the Gibbs sampler in Subsections 3.2 and 3.3.

Theorem 2.6. *Under Assumption 2.1, if* $\mathbf{C} \subseteq \mathbb{R}^n$ *is a polyhedral set, then the density of the projected Gaussian* $\Pi_{\mathbf{C}}^{\Sigma^{-1}}(\boldsymbol{x}^\star)$ *on the relative interior of any face of* $\mathbf{C}$ *is proportional to the density of* $\boldsymbol{x}^\star$.

Proof. Consider a face $\mathbf{F}$ of the polyhedral set $\mathbf{C} \subseteq \mathbb{R}^n$ with linearly independent normal vectors $\boldsymbol{a}_1, \ldots, \boldsymbol{a}_k \in \mathbb{R}^n$, where k is the dimension of the face $\mathbf{F}$, and let $F \in \mathbb{R}^{n \times (n-k)}$ be a matrix whose columns form a basis for the space orthogonal to the span of the normal vectors $\boldsymbol{a}_1, \ldots, \boldsymbol{a}_k$, i.e., $F^T \boldsymbol{a}_i = \mathbf{0}$ for all $i = 1, \ldots, k$. If $\boldsymbol{x}_0$ is any point on $\mathbf{F}$, then each point $\boldsymbol{z} \in \mathbf{F}$ can be written as $\boldsymbol{x}_0 + F\boldsymbol{u}$ for some $\boldsymbol{u}$ and the corresponding normal cone can be parameterized as $\sum_{i=1}^k t_i \boldsymbol{a}_i$ for $t_i \geqslant 0$ for all $i = 1, \ldots, k$. The density $\pi_{\mathbf{F}}$ on the face then satisfies

$$\pi_{\mathbf{F}}(\boldsymbol{u}) \propto \int_0^\infty \cdots \int_0^\infty \pi_{\boldsymbol{x}^\star}\left(\boldsymbol{x}_0 + F\boldsymbol{u} + \sum_{i=1}^k t_i \Sigma \boldsymbol{a}_i\right) \mathrm{d}t_1 \cdots \mathrm{d}t_k.$$

Note that

$$\begin{aligned} &-2 \log \left(\frac{\pi_{\boldsymbol{x}^\star}\left(\boldsymbol{x}_0 + F\boldsymbol{u} + \sum_{i=1}^k t_i \Sigma \boldsymbol{a}_i\right)}{\pi_{\boldsymbol{x}^\star}\left(\boldsymbol{x}_0 + F\boldsymbol{u}\right)} \right) \\ &\quad = 2 \left(\sum_{i=1}^k t_i \boldsymbol{a}_i\right)^T \left(\boldsymbol{x}_0 + F\boldsymbol{u} - \boldsymbol{\mu}\right) + \left(\sum_{i=1}^k t_i \Sigma \boldsymbol{a}_i\right)^T \Sigma^{-1} \left(\sum_{i=1}^k t_i \Sigma \boldsymbol{a}_i\right). \end{aligned}$$

The only term that depends on $\boldsymbol{u}$ vanishes because

$$\left(\sum_{i=1}^k t_i \boldsymbol{a}_i\right)^T F\boldsymbol{u} = \sum_{i=1}^k t_i \boldsymbol{a}_i^T F \boldsymbol{u} = 0,$$

and hence

$$\pi_{\boldsymbol{x}^\star}\left(\boldsymbol{x}_0 + F\boldsymbol{u} + \sum_{i=1}^k t_i \Sigma \boldsymbol{a}_i\right) \propto c(t_1, \ldots t_k) \pi_{\boldsymbol{x}^\star}\left(\boldsymbol{x}_0 + F\boldsymbol{u}\right),$$

where $c(t_1, \ldots t_k)$ only depends on the coefficients t_i. It follows that

$$\pi_{\mathbf{F}}(\boldsymbol{u}) \propto \int_0^\infty \ldots \int_0^\infty \pi_{\boldsymbol{x}^\star}\left(\boldsymbol{x}_0 + F\boldsymbol{u} + \sum_{i=1}^k t_i \Sigma \boldsymbol{a}_i\right) \mathrm{d}t_1 \ldots \mathrm{d}t_k \propto \pi_{\boldsymbol{x}^\star}\left(\boldsymbol{x}_0 + F\boldsymbol{u}\right),$$

i.e., the density of the projected Gaussian on a face of the polyhedral set is proportional to the unprojected Gaussian. ∎

3. Bayesian linear inverse problems with constraints.

In this section, we describe how to apply the theory for projected Gaussian distributions in Section 2 to linear inverse problems. We discuss the randomized optimization problems to be solved to obtain samples from the projected posteriors and how the projected posterior relates to a constrained prior. Finally, we define a Bayesian hierarchical model for the linear inverse problem and derive a Gibbs sampler for that model in the special case where the constraint set $\mathbf{C}$ is a polyhedral cone.

3.1. Bayesian model. Let us now consider the problem of recovering a signal $\boldsymbol{x} \in \mathbb{R}^n$ from noisy observations $\boldsymbol{b} = A\boldsymbol{x} + \boldsymbol{e}$ using a linear forward operator $A : \mathbb{R}^n \to \mathbb{R}^m$ and with noise $\boldsymbol{e} \sim \mathcal{N}(\mathbf{0}, \Sigma_{\boldsymbol{e}})$ for $\Sigma_{\boldsymbol{e}} \in \mathbb{S}^n_{++}$. This results in the likelihood function

$$\pi(\boldsymbol{b} \,|\, \boldsymbol{x}) \propto \exp\left(-\frac{1}{2}\|A\boldsymbol{x} - \boldsymbol{b}\|^2_{\Sigma_{\boldsymbol{e}}^{-1}}\right).$$

Furthermore, assume a priori that $\boldsymbol{x} \sim \mathcal{N}(\mathbf{0}, \Sigma_{\boldsymbol{x}})$ with $\Sigma_{\boldsymbol{x}} \in \mathbb{S}^n_{++}$, i.e.,

$$\pi(\boldsymbol{x}) \propto \exp\left(-\frac{1}{2}\|\boldsymbol{x}\|^2_{\Sigma_{\boldsymbol{x}}^{-1}}\right).$$

From Bayes' theorem we obtain the posterior distribution

$$\pi(\boldsymbol{x} \,|\, \boldsymbol{b}) \propto \pi(\boldsymbol{b} \,|\, \boldsymbol{x})\pi(\boldsymbol{x}) \propto \exp\left(-\frac{1}{2}\|A\boldsymbol{x} - \boldsymbol{b}\|^2_{\Sigma_{\boldsymbol{e}}^{-1}} - \frac{1}{2}\|\boldsymbol{x}\|^2_{\Sigma_{\boldsymbol{x}}^{-1}}\right), \tag{3.1}$$

which is a Gaussian distribution with covariance $\Sigma_{\boldsymbol{x}|\boldsymbol{b}} = (\Sigma_{\boldsymbol{x}}^{-1} + A^T\Sigma_{\boldsymbol{e}}^{-1}A)^{-1}$ and mean $\boldsymbol{\mu} = \Sigma_{\boldsymbol{x}|\boldsymbol{b}}A^T\Sigma_{\boldsymbol{e}}^{-1}\boldsymbol{b}$. Now consider a closed, convex set $\mathbf{C} \subset \mathbb{R}^n$. To sample from the projected posterior, instead of solving

$$\operatorname*{argmin}_{\boldsymbol{x}\in\mathbf{C}} \left\{\frac{1}{2}\boldsymbol{x}^T(\Sigma_{\boldsymbol{x}}^{-1} + A^T\Sigma_{\boldsymbol{e}}^{-1}A)\boldsymbol{x} - \boldsymbol{x}^T(A^T\Sigma_{\boldsymbol{e}}^{-1}\boldsymbol{b} + \hat{\boldsymbol{w}})\right\},$$

where $\hat{\boldsymbol{w}} \sim \mathcal{N}(\mathbf{0}, \Sigma_{\boldsymbol{x}|\boldsymbol{b}}^{-1})$, we can equivalently solve

$$\operatorname*{argmin}_{\boldsymbol{x}\in\mathbf{C}} \left\{\frac{1}{2}\|A\boldsymbol{x} - \hat{\boldsymbol{b}}\|^2_{\Sigma_{\boldsymbol{e}}^{-1}} + \frac{1}{2}\|\boldsymbol{x} - \hat{\boldsymbol{c}}\|^2_{\Sigma_{\boldsymbol{x}}^{-1}}\right\}, \tag{3.2}$$

where $\hat{\boldsymbol{b}} \sim \mathcal{N}(\boldsymbol{b}, \Sigma_{\boldsymbol{e}})$ and $\hat{\boldsymbol{c}} \sim \mathcal{N}(\mathbf{0}, \Sigma_{\boldsymbol{x}})$, i.e., solving randomized constrained linear least squares problems or randomized MAP estimates.

For a simple case where $\boldsymbol{e} \sim \mathcal{N}(\mathbf{0}, (\lambda I)^{-1})$ and $\boldsymbol{x} \sim \mathcal{N}(\mathbf{0}, (\delta L^T L)^{-1})$ with hyperparameters $\lambda, \delta > 0$ and L is a full-rank matrix, the posterior (3.1) simplifies to

$$\pi(\boldsymbol{x} \,|\, \boldsymbol{b}) \propto \pi(\boldsymbol{b} \,|\, \boldsymbol{x})\pi(\boldsymbol{x}) \propto \exp\left(-\frac{\lambda}{2}\|A\boldsymbol{x} - \boldsymbol{b}\|_2^2 - \frac{\delta}{2}\|L\boldsymbol{x}\|_2^2\right). \tag{3.3}$$

The optimization problem (3.2) then simplifies to

$$\operatorname*{argmin}_{\boldsymbol{x}\in\mathbf{C}} \left\{\frac{\lambda}{2}\|A\boldsymbol{x} - \hat{\boldsymbol{b}}\|_2^2 + \frac{\delta}{2}\|L\boldsymbol{x} - \hat{\boldsymbol{c}}\|_2^2\right\}, \tag{3.4}$$

where $\hat{\boldsymbol{b}} \sim \mathcal{N}(\boldsymbol{b}, (\lambda I)^{-1})$ and $\hat{\boldsymbol{c}} \sim \mathcal{N}(\mathbf{0}, (\delta I)^{-1})$.

Thus, by repeatedly solving the randomized constrained least squares problem (3.4), we obtain samples from the obliquely projected posterior onto the set $\mathbf{C}$.

3.2. Constrained prior. Although we defined the posterior implicitly through modification of the unconstrained posterior, we can also interpret this modification as having implicitly defined a constrained prior. Let $\mathbf{C} \subseteq \mathbb{R}^n$ be a polyhedral set and define the projected posterior onto $\mathbf{C}$ as in (2.2). By Theorem 2.6, the projected posterior density on any face $\mathbf{F}$ of $\mathbf{C}$ is proportional to the unprojected posterior. Because constraining the signal is based on prior information, we will assume that the likelihood $\pi(\boldsymbol{b}|\boldsymbol{x})$ is not affected by the projection. Hence, by Bayes' formula, we obtain the prior density on the face,

$$\pi_{\boldsymbol{x},\mathbf{F}}(\boldsymbol{u}) \propto \frac{\pi_{\boldsymbol{x}|\boldsymbol{b},\mathbf{F}}(\boldsymbol{u})}{\pi_{\boldsymbol{b}|\boldsymbol{x}}(\boldsymbol{x}_0 + F\boldsymbol{u})} \propto \frac{\exp\left(-\frac{1}{2}\|A\boldsymbol{x} - \boldsymbol{b}\|^2_{\Sigma_{\boldsymbol{e}}^{-1}} - \frac{1}{2}\|\boldsymbol{x}\|^2_{\Sigma_{\boldsymbol{x}}^{-1}}\right)}{\exp\left(-\frac{1}{2}\|A\boldsymbol{x} - \boldsymbol{b}\|^2_{\Sigma_{\boldsymbol{e}}^{-1}}\right)} \tag{3.5}$$

$$\propto \exp\left(-\frac{1}{2}\|\boldsymbol{x}\|^2_{\Sigma_{\boldsymbol{x}}^{-1}}\right) \propto \pi_{\boldsymbol{x}}(\boldsymbol{x}_0 + F\boldsymbol{u}),$$

where we denote by $\boldsymbol{x} = \boldsymbol{x}_0 + F\boldsymbol{u}$ the point of the face $\mathbf{F}$ parameterized by $\boldsymbol{u}$.

This shows that the corresponding constrained prior is proportional to the unconstrained prior on any face of $\mathbf{C}$ and therefore has a similar structure on $\mathbf{C}$ as the posterior. However, this constrained prior is generally not the same as obliquely projecting the prior onto the constraint set.

3.3. Bayesian hierarchical model and Gibbs sampler. Let us now consider adding priors to the hyperparameters λ and δ. To exploit that Gaussian and Gamma distributions are conjugate, let the hyperpriors be $\lambda \sim \Gamma(\alpha_\lambda, \beta_\lambda)$ and $\delta \sim \Gamma(\alpha_\delta, \beta_\delta)$, i.e.,

$$\pi(\lambda) \propto \lambda^{\alpha_\lambda - 1} \exp(-\beta_\lambda \lambda), \qquad \text{for } \lambda > 0 \quad \text{and}$$
$$\pi(\delta) \propto \delta^{\alpha_\delta - 1} \exp(-\beta_\delta \delta), \qquad \text{for } \delta > 0.$$

Combining these hyperpriors with the likelihood and prior of the previous section results in a Bayesian hierarchical model. A common method for sampling for the signal $\boldsymbol{x}$ and hyperparameters λ and δ is a hierarchical Gibbs sampler [2, Algorithm 5.1]. For simplicity of presentation, we consider $\boldsymbol{e} \sim \mathcal{N}(\mathbf{0}, (\lambda I)^{-1})$ and $\boldsymbol{x} \sim \mathcal{N}(\mathbf{0}, (\delta L^T L)^{-1})$ as in the later parts of Subsection 3.1. However, similar derivations can be done for more complicated covariance matrices. For the unconstrained setting of (3.3), it can be derived that

$$\pi_{x,\lambda,\delta|b}(x, \lambda, \delta) \propto \lambda^{m/2+\alpha_\lambda-1}\delta^{n/2+\alpha_\delta-1} \times \exp\left(-\frac{\lambda}{2}\|A\boldsymbol{x} - \boldsymbol{b}\|_2^2 - \frac{\delta}{2}\|L\boldsymbol{x}\|_2^2 - \beta_\lambda\lambda - \beta_\delta\delta\right),$$

from which it follows that

$$\lambda \,|\, \boldsymbol{x}, \boldsymbol{b} \sim \Gamma(m/2 + \alpha_\lambda, \frac{1}{2}\|A\boldsymbol{x} - \boldsymbol{b}\|_2^2 + \beta_\lambda),$$
$$\delta \,|\, \boldsymbol{x}, \boldsymbol{b} \sim \Gamma(n/2 + \alpha_\delta, \frac{1}{2}\|L\boldsymbol{x}\|_2^2 + \beta_\delta), \tag{3.6}$$

and $\boldsymbol{x} \,|\, \lambda, \delta, \boldsymbol{b}$ is a Gaussian described by (3.3). A Gibbs sampler then alternates among sampling from these conditional distributions.

Now for the constrained setting, let $\mathbf{C}$ be a polyhedral cone, i.e., $\mathbf{C}$ is the conic hull of finitely many vectors. Combined with the constrained prior in (3.5), we can compute the normalization constant K of the density of the prior conditioned on the face $\mathbf{F}$ of $\mathbf{C}$ as follows,

$$1 = \int_{\mathbf{F}} \pi_{\boldsymbol{x}|\mathbf{F}}(\boldsymbol{u})\mathrm{d}\boldsymbol{u} = K \int_{\mathbf{F}} \exp\left(-\frac{\delta}{2}\|LF\boldsymbol{u}\|_2^2\right) \mathrm{d}\boldsymbol{u} = K\delta^{-\dim(\mathbf{F})/2} \int_{\mathbf{F}} \exp\left(-\frac{1}{2}\|LF\boldsymbol{v}\|_2^2\right) \mathrm{d}\boldsymbol{v},$$

where we used that the face of any polyhedral cone is again a polyhedral cone, hence $c\mathbf{F} = \mathbf{F}$ for any $c > 0$ and we can take $\boldsymbol{x}_0 = \mathbf{0}$. The normalization constant K is therefore

$$K = \frac{\delta^{\dim(\mathbf{F})/2}}{\int_{\mathbf{F}} \exp\left(-\frac{1}{2}\|LF\boldsymbol{v}\|_2^2\right) \mathrm{d}\boldsymbol{v}},$$

hence the distribution of the prior conditioned on the face satisfies

$$\pi_{\boldsymbol{x}|\mathbf{F}}(\boldsymbol{u}) \propto \delta^{\dim(\mathbf{F})/2} \exp\left(-\frac{\delta}{2}\|LF\boldsymbol{u}\|_2^2\right),$$

where the proportionality does not depend on δ anymore.

Now we can obtain the distribution of the (hyper)parameters using Bayes' formula,

$$\begin{aligned}
\pi_{\boldsymbol{x},\lambda,\delta|\boldsymbol{b},\mathbf{F}}(\boldsymbol{u},\lambda,\delta) &\propto \pi_{\boldsymbol{b}|\boldsymbol{x},\lambda,\delta}(F\boldsymbol{u})\pi_{\boldsymbol{x}|\delta,\mathbf{F}}(\boldsymbol{u})\pi_\lambda(\lambda)\pi_\delta(\delta) \\
&\propto \lambda^{m/2+\alpha_\lambda-1}\delta^{\dim(\mathbf{F}(F\boldsymbol{u}))/2+\alpha_\delta-1} \\
&\quad \times \exp\left(-\frac{\lambda}{2}\|AF\boldsymbol{u}-\boldsymbol{b}\|_2^2 - \frac{\delta}{2}\|LF\boldsymbol{u}\|_2^2 - \beta_\lambda\lambda - \beta_\delta\delta\right),
\end{aligned}$$

where $\mathbf{F}(\boldsymbol{x})$ is the smallest face of $\mathbf{C}$ that contains $\boldsymbol{x}$. Therefore, we can conclude the following conditional distributions of the hyperparameters,

$$\begin{aligned}
\pi(\lambda \,|\, \boldsymbol{x},\boldsymbol{b},\mathbf{F}) &\propto \lambda^{m/2+\alpha_\lambda-1}\exp(-\beta_\lambda\lambda - \frac{\lambda}{2}\|A\boldsymbol{x}-\boldsymbol{b}\|_2^2), \text{ for } \lambda > 0 \quad \text{and} \\
\pi(\delta \,|\, \boldsymbol{x},\boldsymbol{b},\mathbf{F}) &\propto \delta^{\dim(\mathbf{F}(\boldsymbol{x}))/2+\alpha_\delta-1}\exp(-\beta_\delta\delta - \frac{\delta}{2}\|L\boldsymbol{x}\|_2^2), \text{ for } \delta > 0,
\end{aligned}$$

or equivalently

$$\begin{aligned}
\lambda \,|\, \boldsymbol{x},\boldsymbol{b},\mathbf{F} &\sim \Gamma(m/2+\alpha_\lambda, \frac{1}{2}\|A\boldsymbol{x}-\boldsymbol{b}\|_2^2 + \beta_\lambda), \\
\delta \,|\, \boldsymbol{x},\boldsymbol{b},\mathbf{F} &\sim \Gamma(\dim(\mathbf{F}(\boldsymbol{x}))/2+\alpha_\delta, \frac{1}{2}\|L\boldsymbol{x}\|_2^2 + \beta_\delta).
\end{aligned} \tag{3.7}$$

From these conditional distributions, we obtain the Polyhedral Cone Hierarchical Gibbs Sampler 3.1.

The main difference with the ordinary Hierarchical Gibbs sampler is δ, see the difference between (3.6) and (3.7). It can be quite expensive to compute $\dim(\mathbf{F}(\boldsymbol{x}^k))$, but in a few cases there are simpler expressions. In the simplest setting where $\mathbf{C} = \mathbb{R}^n$, i.e., the unconstrained setting, then the only face is the whole space $\mathbb{R}^n$, hence $\dim(\mathbf{F}(\boldsymbol{x}^k)) = n$ and the Polyhedral

Input: $\boldsymbol{x}^0, \alpha_\lambda, \beta_\lambda, \alpha_\delta, \beta_\delta, k_{\max}$
for $k = 1$ **to** $k_{\max}$ **do**
  Compute $(\lambda_k, \delta_k) \sim \pi_{\lambda,\delta|\boldsymbol{x},\boldsymbol{b}}$ as follows:
    $\lambda_k \sim \Gamma\left(m/2 + \alpha_\lambda, \frac{1}{2}\|A\boldsymbol{x}^{k-1} - \boldsymbol{b}\|_2^2 + \beta_\lambda\right)$,
    $\delta_k \sim \Gamma\left(\dim(\mathbf{F}(\boldsymbol{x}^{k-1}))/2 + \alpha_\delta, \frac{1}{2}\|L\boldsymbol{x}^{k-1}\|_2^2 + \beta_\delta\right)$.
  Compute $\boldsymbol{x}^k \sim \pi_{\boldsymbol{x}|\boldsymbol{b},\lambda^k,\delta^k}$ using (3.4)
end for
return $\{(\boldsymbol{x}^k, \lambda_k, \delta_k)\}_{k=1,\ldots,k_{\max}}$

Sampler 3.1: Polyhedral Cone Hierarchical Gibbs Sampler

Cone Hierarchical Gibbs Sampler reduces to an ordinary Hierarchical Gibbs Sampler. In the more complicated setting where $\mathbf{C} = \mathbb{R}^n_{\geqslant 0}$, i.e., nonnegativity constraints, then the faces are characterized by the zero values of the vector and $\dim(\mathbf{F}(\boldsymbol{x}))$ simplifies to the number of non-zero values elements of $\boldsymbol{x}$. The Polyhedral Cone Hierarchical Gibbs Sampler above then simplifies to and provides a rigorous justification for the Nonnegative Hierarchical Gibbs Sampler of [4].

4. Numerical examples. In this section, we present two numerical examples. First, we consider a one-dimensional deblurring problem and investigate the effect of constraints on the posterior distribution. Second, we consider a Bayesian hierarchical model for a two-dimensional computed tomography (CT) problem. For this problem, we use the Gibbs sampler as described in the Subsection 3.3 and take particular interest in efficiently and approximately solving the constrained linear least squares problems required for sampling.

4.1. One-dimensional deblurring with different constraints. Let us first consider a one-dimensional Gaussian deblurring problem defined by

$$\boldsymbol{b} = A\boldsymbol{x} + \boldsymbol{e}, \tag{4.1}$$

for a true signal $\boldsymbol{x} \in [0,1]^n$, noise $\boldsymbol{e} \sim \mathcal{N}(\mathbf{0}, \lambda^{-1}I)$ with hyperparameter $\lambda > 0$ and forward operator A defined by the Toeplitz matrix

$$A_{ij} = \frac{h}{\gamma\sqrt{2\pi}} \exp\left(-\frac{1}{2}\left(\frac{h(i-j)}{\gamma}\right)^2\right), \qquad \text{for } i,j = 1,\ldots,n$$

where $n = 128$, $h = 1/n$ and $\gamma = 0.02$. Assume a priori that $\boldsymbol{x} \sim \mathcal{N}(\mathbf{0}, (\delta L^T L)^{-1})$, where $\delta > 0$ is a hyperparameter and L is a periodic first-order finite difference matrix. Let the hyperparameters be fixed to $\lambda = 1000$ and $\delta = 150$.

Figure 3 shows the true signal and the noisy measurements obtained through (4.1) with $\lambda = 1000$. For the specific instance of $\boldsymbol{e}$ we have $\|\boldsymbol{e}\|/\|A\boldsymbol{x}\| \approx 6\%$. The true signal is divided into multiple components in order to illustrate the impact of constraints on different signal behavior. These components include instantaneous changes between the extreme values 0 and 1, a small deviation from an extreme value and a smooth transition between extreme values. A lot of the components of the true signal have extreme values, hence we will be a priori

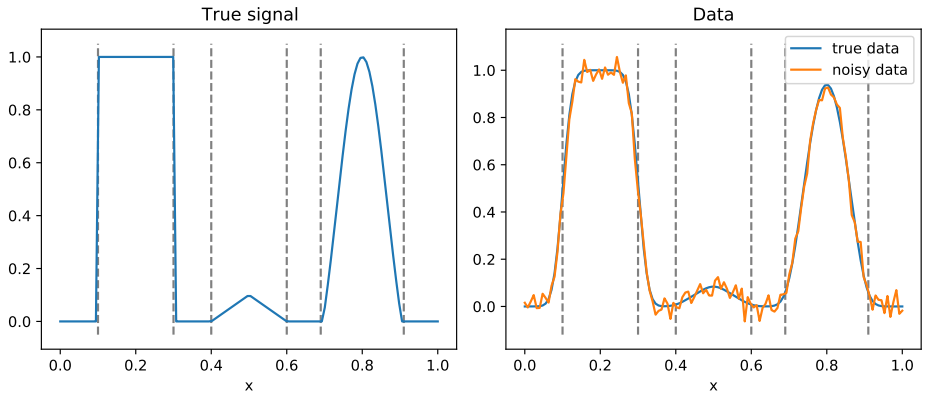

Figure 3: True signal, true data and noisy data for a Gaussian deblurring problem.

interested in signals on the boundary of $[0,1]^n$, i.e., the set of signals for which at least one element is 0 or 1.

Figure 4a shows the component-wise median and 95% credible intervals for the posterior from 10000 samples obtained by repeatedly solving optimization problem (3.4) in the unconstrained ($\mathbb{R}^n$), nonnegative constrained ($\mathbb{R}^n_{\geqslant 0}$) and box constrained ($[0,1]^n$) settings. Note that the fluctuating features near the extreme values 0 and 1 get flattened when applying the constraints. From the viewpoint of [14], the oblique projection is just one possible projection. Alternatively, the Euclidean projector provides an efficient method when applied to unconstrained samples. However, when applying the Euclidean projector, as illustrated in Figure 4b, these fluctuations around the extreme values are stronger.

As we use the median as central point estimate, we measured the variation of the samples using the width of component-wise credible intervals. Figure 5 shows the width of 95% component-wise credible intervals for the examples in Figure 4. Figure 5 shows that the width is reduced most where the signal values lie close the extreme values, but the width is also reduced for values close to extreme. Furthermore, using the oblique projection generally reduces the width more than the Euclidean projector. A possible explanation is that the solutions to constrained linear least squares problems have better stability properties than their unconstrained counterparts, see [5].

4.2. Gibbs sampler for CT reconstruction with nonnegativity constraints. Let us now consider a CT problem [9] of the form

$$\boldsymbol{b} = A\boldsymbol{x} + \boldsymbol{e},$$

where the true signal $\boldsymbol{x} \in \mathbb{R}^{100\times 100}$ is the Shepp-Logan phantom shown in Figure 6, the noise satisfies $\boldsymbol{e} \sim \mathcal{N}(\boldsymbol{0}, \lambda^{-1} I)$ with hyperparameter $\lambda > 0$, and forward operator A obtained from AIR Tools II [10] is a discretized Radon transform at 180 angles with 140 rays and a parallel-beam geometry. Assume a priori that $\boldsymbol{x} \sim \mathcal{N}(\boldsymbol{0}, (\delta L^T L)^{-1})$ where $\delta > 0$ is a hyperparameter and L is a first-order finite difference matrix. We generated noisy data using $\lambda = 5$. For the

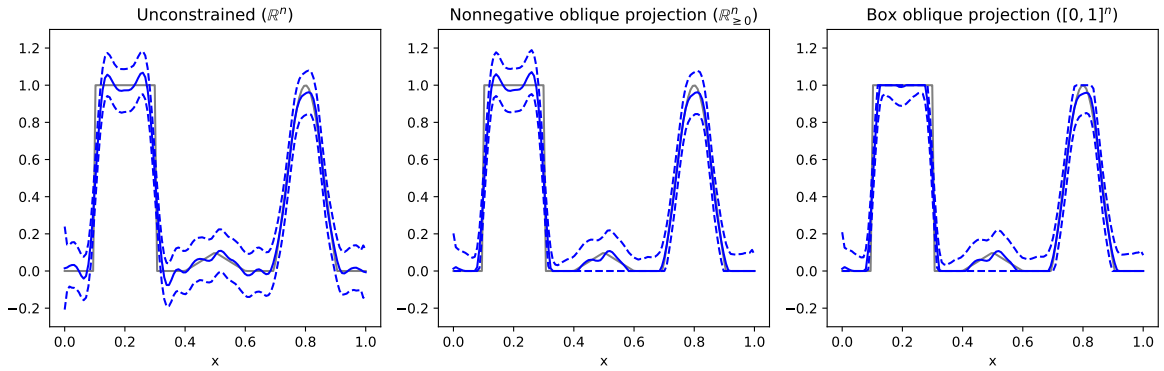

(a) Different constraints with oblique projection.

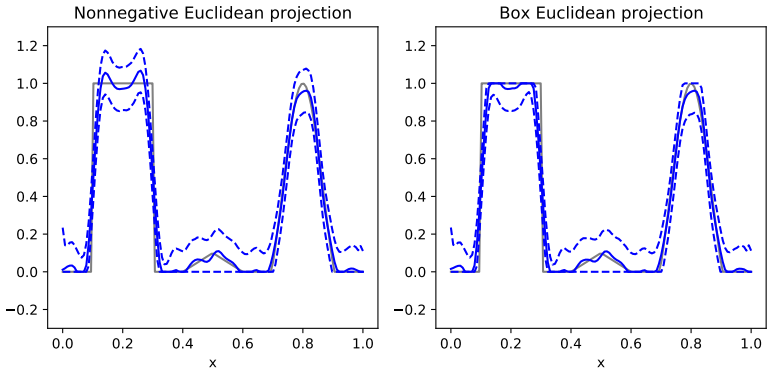

(b) Different constraints with standard Euclidean projection.

Figure 4: Component-wise median and 95% credible intervals for 10000 samples of the Gaussian deblurring model.

specific instance of $\boldsymbol{e}$ we have $\|\boldsymbol{e}\|/\|A\boldsymbol{x}\| \approx 4\%$.

In CT, the signal $\boldsymbol{x}$ represents attenuation coefficients that are bounded from below by the corresponding background value, generally air. For simplicity, this constraint has been modeled as nonnegatity. In this experiment, we used the Polyhedral Cone Hierarchical Gibbs Sampler to sample from the posterior both with and without nonnegativity constraints. The parameters for the hyperprior were chosen to be $\alpha_\lambda = \alpha_\delta = 1$ and $\beta_\lambda = \beta_\delta = 10^{-4}$, similar to [4].

Repeatedly solving optimization problem (3.4) to high accuracy is a very costly procedure. Therefore, instead of solving the optimization problem from scratch at each iteration of the Gibbs sampler, we solve the optimization problem each time using a few iterations of FISTA [6] with the previous sample as warm-start. Using a small number of iterations, greatly reduces the computation time, but the samples become more correlated and do not have to be samples from the target distribution anymore. This is similar to using a few steps of CG in a gradient

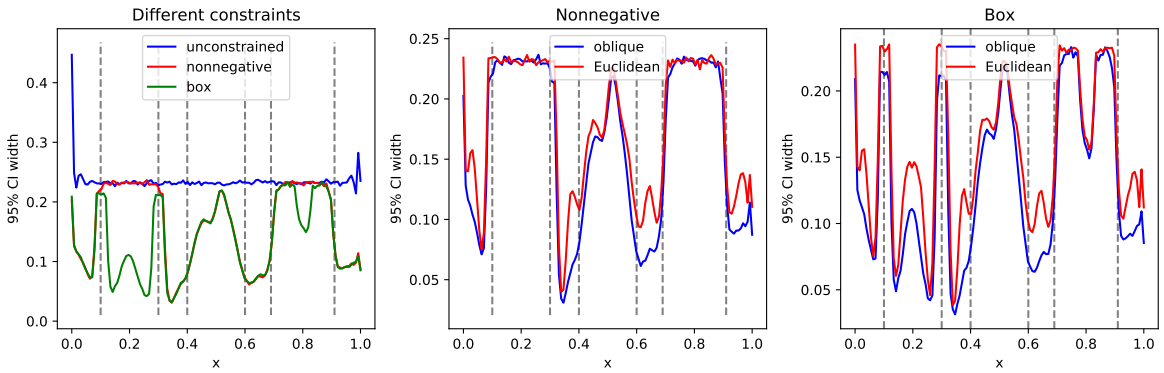

Figure 5: Component-wise width of 95% credible intervals for different constraints and projectors (oblique and Euclidean) obtained from 10000 samples of the Gaussian deblurring model.

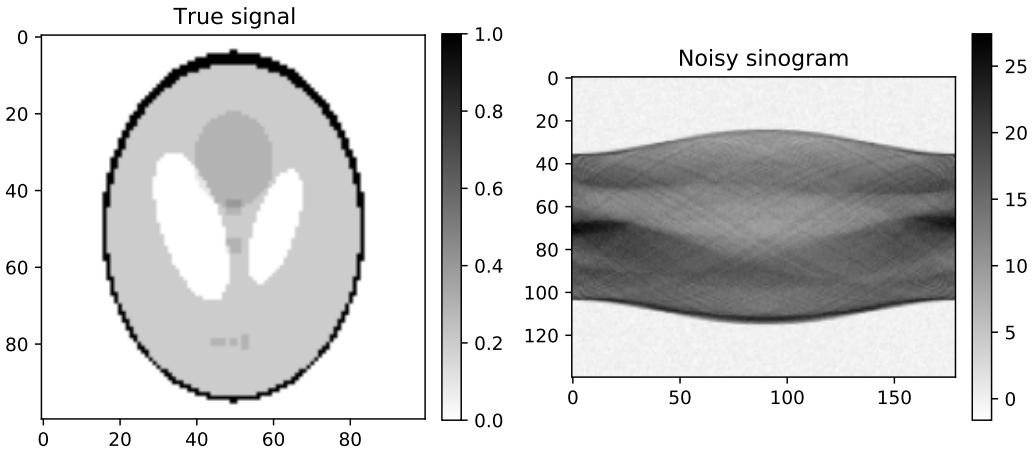

Figure 6: True Shepp-Logan phantom (left) and noisy sinogram (right).

scan Gibbs sampler [2, Algorithm 5.4], although this guarantees convergence to the target distribution.

We chose to run 100 iterations of FISTA for each sample with the dynamic stepsize $0.99(\lambda\|A^TA\|_2 + \delta\|L^TL\|_2)^{-1}$. This stepsize is a positive lower-bound on the inverse of the Lipschitz constant of the objective function and therefore guarantees convergence. We ran the algorithm for 15000 samples and removed the first 1000 samples as burn-in.

Figure 7 shows the autocorrelation function (ACF) and the distribution of the hyperparameters λ and δ. First, note that the autocorrelation function for λ decays slightly faster in the nonnegative setting than the unconstrained setting. This shows that the nonnegative samples are slightly less correlated. Second, although the actual noise level is the same in

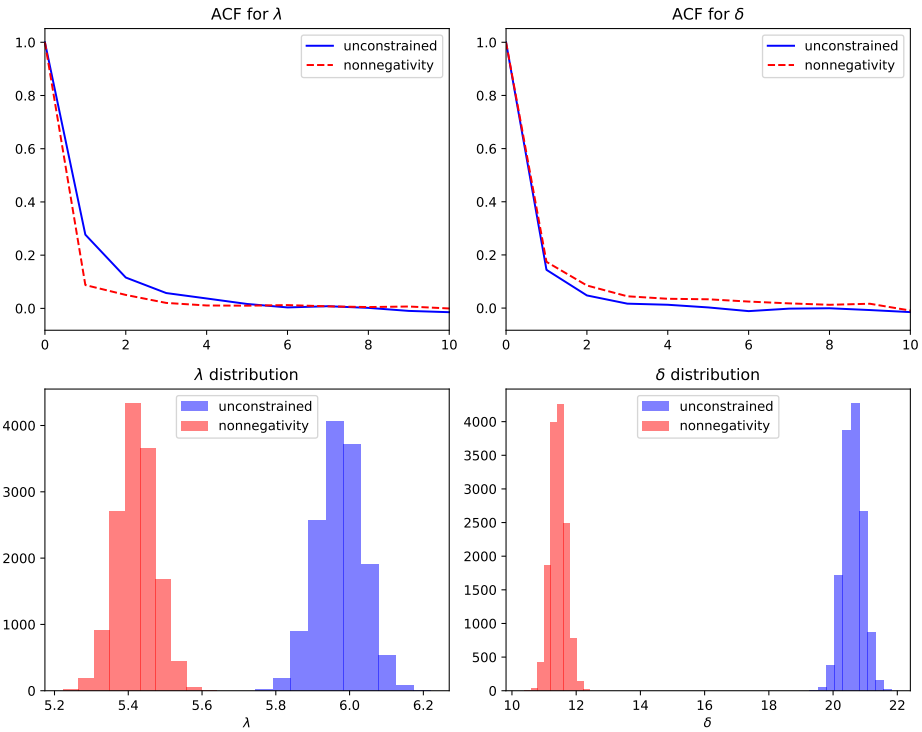

Figure 7: Autocorrelation functions (ACF) for the hyperparameters and hyperparameter distributions for unconstrained and nonnegative setting.

both settings, the noise parameter λ and the prior hyperparameter δ are noticeably smaller.

Figure 8 shows the component-wise median for both the unconstrained and nonnegative setting. The main difference between the two medians lies in the background of the phantom. The unconstrained median has a non-uniform background containing a lot of small artefacts, while the nonnegative median is, besides a few pixels, uniformly zero in the background.

Another difference between the unconstrained and nonnegative setting is the component-wise variation. Figure 9 shows the width of the component-wise 95% credible intervals together with the difference between the two settings. Note the different range of the unconstrained and nonnegative settings. Just like in the deblurring experiment, the width is greatly reduced in the components that are close to the extreme values. Therefore, the uncertainty can be used to easily distinguish between the actual object and the background. The uncertainty inside the object is slightly larger with nonnegativity constraints, which is due to the smaller λ and δ hyperparameters as observed in Figure 7.

5. Conclusion. We have presented a method for handling constraints in Bayesian inference that puts positive probability on the boundary of the constraint set. In general, the method works by projecting posterior samples from outside the constraint set onto the constraint set. Therefore, the method can be interpreted as post-processing the posterior by projecting the

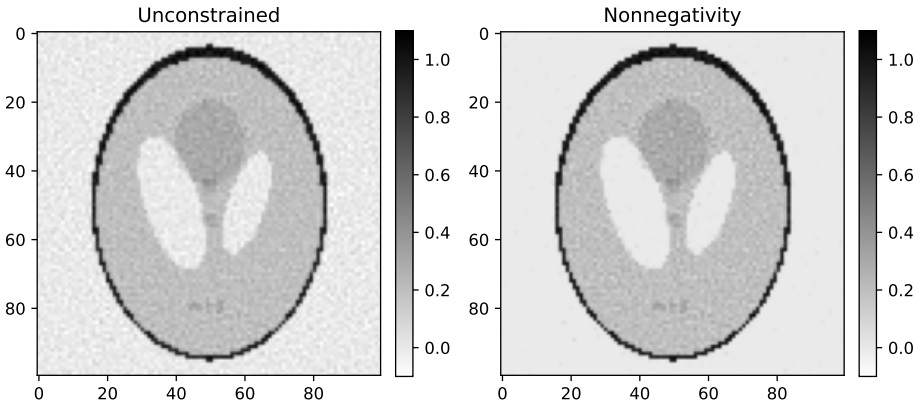

Figure 8: Component-wise median for the unconstrained (left) and nonnegative (right) settings. Note that the range is the same.

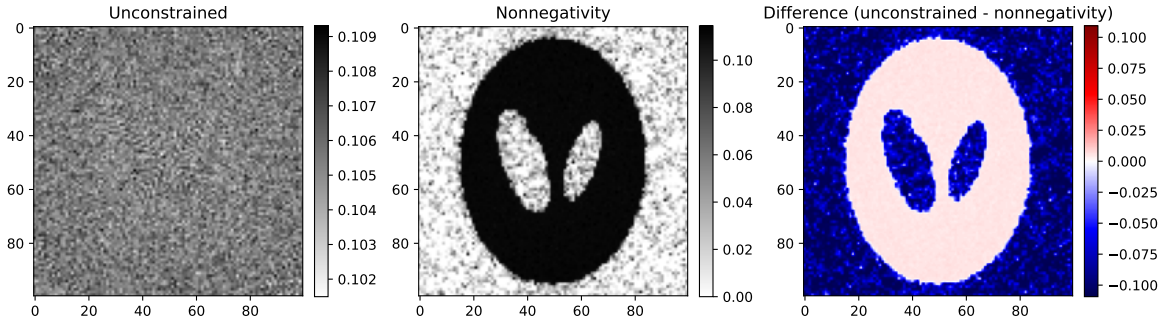

Figure 9: Component-wise width of 95% credible intervals for the unconstrained (left) and nonnegative (middle) settings, together with their difference (right).

density onto the constraint set.

The posterior sampling and projection steps can be combined if the posterior distribution is Gaussian and the projection is with respect to the posterior precision matrix. Samples from such a post-processed posterior can be obtained by solving perturbed constrained quadratic optimization problems. Although this distribution is difficult to describe in closed form for general constraint sets, we were able to characterize the distribution when the constraint set is a polyhedral set. We have proven that the projected posterior on a polyhedral set consists of densities on each of the faces and these densities are proportional to the unprojected posterior density.

To apply the theory, we considered Bayesian linear inverse problems. For such problems,

sampling from the projected posterior can be achieved by solving perturbed constrained linear least squares problems. Furthermore, we considered a Bayesian hierarchical model and derived a Gibbs sampler for the constrained problem when the constraint set is a polyhedral cone.

We tested the projection method on deblurring and CT test cases, for which component-wise bounds are natural constraints. These numerical experiments have shown that the projected posterior can greatly reduce the uncertainty of vector components that are close to their bounds. Furthermore, the experiments suggest that using the oblique projection instead of the Euclidean projection gives better results.

One major issue with sampling by means of solving perturbed constrained least squares problems is the computational cost. The computational cost of solving these optimization problems accurately can be very high, yet it is still unknown to what extent accurate solutions are necessary. Therefore, further research should focus on determining what the effect that inaccurately solving the optimization problems has on the distribution of the samples.

Another topic of further study is identifying more techniques for efficient post-processing of posteriors beyond constraints. For example, adding penalization functions to the randomized least squares problems results in new modified posteriors. The computational cost of sampling from this new posterior is similar to the projected posterior, but it can introduce different regularization-like effects.

REFERENCES

[1] K. O. Bangsgaard and M. S. Andersen, *A statistical reconstruction model for absorption CT with source uncertainty*, Inverse Problems, 37 (2021), p. 085009.

[2] J. M. Bardsley, *Computational uncertainty quantification for inverse problems*, vol. 19, SIAM, 2018.

[3] J. M. Bardsley and C. Fox, *An MCMC method for uncertainty quantification in nonnegativity constrained inverse problems*, Inverse Problems in Science and Engineering, 20 (2012), pp. 477–498.

[4] J. M. Bardsley and P. C. Hansen, *MCMC algorithms for computational UQ of nonnegativity constrained linear inverse problems*, SIAM Journal on Scientific Computing, 42 (2020), pp. A1269–A1288.

[5] J. M. Bardsley, J. K. Merikoski, and R. Vio, *The stabilizing properties of nonnegativity constraints in least-squares image reconstruction*, International Journal of Pure and Applied Mathematics, 43 (2008), p. 95.

[6] A. Beck and M. Teboulle, *A fast iterative shrinkage-thresholding algorithm for linear inverse problems*, SIAM journal on imaging sciences, 2 (2009), pp. 183–202.

[7] A. Gelman, J. B. Carlin, H. S. Stern, and D. B. Rubin, *Bayesian data analysis*, Chapman and Hall/CRC, 1995.

[8] J. F. Geweke, *Bayesian inference for linear models subject to linear inequality constraints*, in Modelling and Prediction Honoring Seymour Geisser, Springer, 1996, pp. 248–263.

[9] P. C. Hansen, J. Jørgensen, and W. R. Lionheart, *Computed Tomography: Algorithms, Insight, and Just Enough Theory*, SIAM, 2021.

[10] P. C. Hansen and J. S. Jørgensen, *AIR Tools II: algebraic iterative reconstruction methods, improved implementation*, Numerical Algorithms, 79 (2018), pp. 107–137.

[11] J.-B. Hiriart-Urruty and C. Lemaréchal, *Fundamentals of convex analysis*, Springer Science & Business Media, 2004.

[12] A. Jeffrey and H. H. Dai, *Handbook of mathematical formulas and integrals*, Elsevier, 2008.

[13] J. Kaipio and E. Somersalo, *Statistical and computational inverse problems*, vol. 160, Springer Science & Business Media, 2006.

[14] D. Sen, S. Patra, and D. Dunson, *Constrained inference through posterior projections*, arXiv preprint arXiv:1812.05741, (2018).

Appendix A. Miscellaneous proofs and computations.

A.1. Linear algebra lemma.

Lemma A.1. *If $\boldsymbol{u}_1, \dots, \boldsymbol{u}_k, \boldsymbol{u}_{k+1}, \dots, \boldsymbol{u}_n \in \mathbb{R}^n$ is an orthonormal basis and Σ is a positive definite matrix, then $\boldsymbol{u}_1, \dots, \boldsymbol{u}_k, \Sigma\boldsymbol{u}_{k+1}, \dots, \Sigma\boldsymbol{u}_n \in \mathbb{R}^n$ is a basis for $\mathbb{R}^n$.*

Proof. Let $U = [U_1, U_2]$ where U_1 and U_2 are matrices with the first k and last $n-k$ vectors u_i as columns respectively. We will proof the claim by showing that $\hat{U} = [U_1, \Sigma U_2]$ has full rank. Note that

$$U^T\hat{U} = \begin{bmatrix} I & U_1^T\Sigma U_2 \\ 0 & U_2^T\Sigma U_2 \end{bmatrix},$$

has full rank, because $U_2^T\Sigma U_2$ is positive definite. Therefore, $\hat{U} = U(U^T\hat{U})$ also has full rank.

A.2. Analytic computations.

A.2.1. Halfspace. Define the halfspace $\mathbf{C} = \{\boldsymbol{x} \in \mathbb{R}^m \,|\, \boldsymbol{a}^T\boldsymbol{x} \leqslant b\}$, where $\boldsymbol{a} \in \mathbb{R}^n$ is a nonzero normal vector and $b \in \mathbb{R}$. Denote by $F \in \mathbb{R}^{n\times(n-1)}$ a matrix whose columns form an orthonormal basis for $\text{Null}(\boldsymbol{a}^T)$, then the boundary of the halfspace defined can be parameterized by $\boldsymbol{x}_0 + F\boldsymbol{u}$ for $\boldsymbol{u} \in \mathbb{R}^{n-1}$. Let $\boldsymbol{x}^\star \sim \mathcal{N}(\boldsymbol{\mu}, \Sigma)$ and $\mathbf{E} \subseteq \text{bd}(\mathbf{C})$ be measurable,

then

$$\mathbb{P}\left(\Pi_{\mathbf{C}}^{\Sigma^{-1}}(\boldsymbol{x}^\star) \in \mathbf{E}\right) = |\det\left(\begin{bmatrix}\Sigma\boldsymbol{a} & F\end{bmatrix}\right)| \int_{\mathbb{R}^{n-1}} \int_0^\infty \pi_{\boldsymbol{x}^\star}(\boldsymbol{x}_0 + F\boldsymbol{u} + t\Sigma\boldsymbol{a})\mathrm{d}t\,\mathrm{d}\boldsymbol{u}$$
$$= \int_{\mathbb{R}^{n-1}} \frac{\boldsymbol{a}^T\Sigma\boldsymbol{a}}{\|\boldsymbol{a}\|_2} \int_0^\infty \pi_{\boldsymbol{x}^\star}(\boldsymbol{x}_0 + F\boldsymbol{u} + t\Sigma\boldsymbol{a})\mathrm{d}t\,\mathrm{d}\boldsymbol{u},$$

hence the density on the boundary can be written as

$$\pi_{\mathrm{bd}}(\boldsymbol{u}) = \frac{\boldsymbol{a}^T\Sigma\boldsymbol{a}}{\|\boldsymbol{a}\|_2} \int_0^\infty \pi_{\boldsymbol{x}^\star}(\boldsymbol{x}_0 + F\boldsymbol{u} + t\Sigma\boldsymbol{a})\mathrm{d}t.$$

The integral can be computed using the following integral identity [12]: for constants $a < 0$ and $b, c \in \mathbb{R}$,

$$\int_0^\infty \exp(at^2 + bt + c)\mathrm{d}t = \frac{\sqrt{\pi}}{2\sqrt{-a}} \exp\left(-\frac{b^2}{4a} + c\right) \operatorname{erfc}\left(-\frac{b}{2\sqrt{-a}}\right).$$

A.2.2. Disc. Suppose that $\mathbf{C}$ is a unit disc defined by $\mathbf{C} = \{\boldsymbol{x} \in \mathbb{R}^2 \,|\, \|\boldsymbol{x}\|_2 \leqslant 1\}$ with boundary parameterization $\boldsymbol{n}(u) := (\cos(u), \sin(u))^T$ for $u \in [0, 2\pi)$. For the two-dimensional ball, we can derive an exact distribution on its boundary in a similar way to the boundary of a halfspace. Let $\mathbf{E} \subseteq \mathrm{bd}(\mathbf{C})$ be measurable, then

$$\mathbb{P}\left(\Pi_{\mathbf{C}}^{\Sigma^{-1}}(\boldsymbol{x}^\star) \in \mathbf{E}\right) = \int_{\mathbf{E}} \int_0^\infty \pi_{\boldsymbol{x}^\star}(\boldsymbol{n}(u) + t\Sigma\boldsymbol{n}(u))|J(t,u)|\mathrm{d}t\,\mathrm{d}u,$$

where

$$|J(t,\theta)| = \det\begin{bmatrix}\Sigma\boldsymbol{n}(u) & (I + t\Sigma)R\boldsymbol{n}(u)\end{bmatrix}, \quad \text{with} \quad R = \begin{bmatrix}0 & -1\\ 1 & 0\end{bmatrix},$$

or equivalently

$$|J(t,\theta)| = \det\begin{bmatrix}\Sigma\boldsymbol{n}(u) & R\boldsymbol{n}(u)\end{bmatrix} + t\det(\Sigma) =: K(u) + t\det(\Sigma).$$

Therefore, the resulting boundary distribution is given by

$$\pi_{\mathrm{bd}}(u) = \int_0^\infty (K(u) + t\det(\Sigma))\pi_{\boldsymbol{x}^\star}(\boldsymbol{n}(u) + t\Sigma\boldsymbol{n}(u))\mathrm{d}t.$$

The integral can be computed using the following integral identity [12]: for constants $a > 0$ and $b, c, d, f \in \mathbb{R}$,

$$\int_0^\infty (d+ft)\exp(-\frac{1}{2}(at^2+bt+c))\mathrm{d}t = \frac{e^{-c/2}}{4a^{3/2}}\left(4f\sqrt{a} + \sqrt{2\pi}(2ad - bf)\exp\left(\frac{b^2}{8a}\right)\operatorname{erfc}\left(\frac{b}{\sqrt{8a}}\right)\right).$$