

Analysis of the Gibbs sampler for hierarchical inverse problems

Sergios Agapiou

Mathematics Institute, University of Warwick
Coventry CV4 7AL, United Kingdom
S.Agapiou@warwick.ac.uk

Johnathan M. Bardsley

Department of Mathematics Sciences, University of Montana
Missoula, MT, 59812-0864 USA
BardsleyJ@mso.umt.edu

Omiros Papaspiliopoulos

Department of Economics, Universitat Pompeu Fabra
Ramon Trias Fargas 25-27, 08005 Barcelona, Spain
Omiros.Papaspiliopoulos@upf.edu

Andrew M. Stuart

Mathematics Institute, University of Warwick
Coventry CV4 7AL, United Kingdom
A.M.Stuart@warwick.ac.uk

Abstract Many inverse problems arising in applications come from continuum models where the unknown parameter is a field. In practice the unknown field is discretized resulting in a problem in \mathbb{R}^N , with an understanding that refining the discretization, that is increasing N , will often be desirable. In the context of Bayesian inversion this situation suggests the importance of two issues: (i) defining hyper-parameters in such a way that they are interpretable in the continuum limit $N \rightarrow \infty$ and so that their values may be compared between different discretization levels; (ii) understanding the efficiency of algorithms for probing the posterior distribution, as a function of large N . Here we address these two issues in the context of linear inverse problems subject to additive Gaussian noise within a hierarchical modelling framework based on a Gaussian prior for the unknown field and an inverse-gamma prior for a hyper-parameter, namely the amplitude of the prior variance. The structure of the model is such that the Gibbs sampler can be easily implemented for probing the posterior distribution. Subscribing to the dogma that one should think infinite-dimensionally before implementing in finite dimensions, we present function space intuition and provide rigorous theory showing that as N increases, the component of the Gibbs sampler for sampling the amplitude of the prior variance becomes increasingly slower. We discuss a reparametrization of the prior variance that is robust with respect to the increase in dimension; we give numerical experiments which exhibit that our reparametrization prevents the slowing down. Our intuition on the behaviour of the prior hyper-parameter, with and without reparametrization, is sufficiently general to include a broad class of nonlinear inverse problems as well as other families of hyper-priors.

Key words. Gaussian process priors, Markov chain Monte Carlo, inverse covariance operators, hierarchical models, diffusion limit.

2010 Mathematics Subject Classification. 62G20, 62C10, 62D05, 45Q05.

1 Introduction

We consider the possibly nonlinear inverse problem of recovering an unknown parameter $\mathbf{u} \in \mathcal{X}$ from a noisy indirect observation $\mathbf{y} \in \mathcal{Y}$. We work in a framework where \mathcal{X} is an infinite-dimensional separable Hilbert space with inner product $\langle \cdot, \cdot \rangle$ and norm $\|\cdot\|$, and \mathcal{Y} is also a separable Hilbert space. We will be especially interested in the case $\mathcal{Y} = \mathcal{X}$ or $\mathcal{Y} = \mathbb{R}^M$. The unknown parameter and the observation are related through an additive noise model

$$\mathbf{y} = \mathcal{G}(\mathbf{u}) + \boldsymbol{\eta}, \quad (1)$$

where $\mathcal{G} : \mathcal{X} \rightarrow \mathcal{Y}$ is the forward map which is assumed to be continuous, and $\boldsymbol{\eta}$ is Gaussian noise

$$\boldsymbol{\eta} \sim \mathcal{N}(0, \lambda^{-1} \mathcal{C}_1). \quad (2)$$

The linear operator $\mathcal{C}_1 : \mathcal{Y} \rightarrow \mathcal{Y}$ is bounded and positive definite and $\lambda > 0$ models the noise level; we do not enforce that \mathcal{C}_1 is trace-class, thereby allowing the case of Gaussian white noise where it is the identity.

We adopt a Bayesian approach with a Gaussian prior on the unknown parameter \mathbf{u}

$$\mathbf{u}|\delta \sim \mathcal{N}(0, \delta^{-1} \mathcal{C}_0), \quad (3)$$

where $\mathcal{C}_0 : \mathcal{X} \rightarrow \mathcal{X}$ is a positive definite and trace-class operator and $\delta > 0$ models the amplitude of the prior variance; the unknown \mathbf{u} is assumed to be independent of the noise $\boldsymbol{\eta}$. The trace-class assumption on \mathcal{C}_0 ensures that draws from the prior on $\mathbf{u}|\delta$ are in \mathcal{X} . For a fixed \mathbf{u} the likelihood is Gaussian, $\mathbf{y}|\mathbf{u}, \delta \sim \mathcal{N}(\mathcal{G}(\mathbf{u}), \lambda^{-1} \mathcal{C}_1)$. We work under certain regularity conditions on the forward map \mathcal{G} , which imply that the inverse problem is sufficiently ill-posed; in particular, for the noise model at hand, these conditions imply that the unknown \mathbf{u} is not perfectly identifiable from a single realization of the data. Under the additional assumption that the prior on $\mathbf{u}|\delta$ is such that the regularity conditions on \mathcal{G} are satisfied in its support, it can be shown that almost surely with respect to the data the posterior on $\mathbf{u}|\mathbf{y}, \delta$ is well defined, non-degenerate and absolutely continuous with respect to the prior on $\mathbf{u}|\delta$, [39].

In what follows, we consider the hyper-parameter δ as a part of the inference problem, that is, we endow it with a prior $\mathbb{P}(\delta)$; this leads to a hierarchical Bayesian model. The potential for the use of hierarchical priors in inverse problems has been highlighted in [23], where the authors express the conviction that *if a parameter is not known, it is a part of the inference problem*; see also [11, 10] where conditionally Gaussian hierarchical models have been considered in finite dimensional contexts. Returning to our setting, we note that of course in practice other aspects of the model, such as parameters that control the regularity of the draws from the prior, will also be part of the inference problem. Section 6 discusses how the results of this paper can be extended to such situations, but the focus here is the joint hierarchical inference on \mathbf{u} and δ . Statistical inference is achieved by Markov chain Monte Carlo sampling from the resulting full posterior on $\mathbf{u}, \delta|\mathbf{y}$, where by Bayes' rule

$$\mathbb{P}(\mathbf{u}, \delta|\mathbf{y}) \propto \mathbb{P}(\mathbf{y}|\mathbf{u}, \delta) \mathbb{P}(\mathbf{u}|\delta) \mathbb{P}(\delta) \propto \mathbb{P}(\mathbf{u}|\mathbf{y}, \delta) \mathbb{P}(\delta|\mathbf{y}).$$

A sufficient condition for this posterior to be well defined is that the prior $\mathbb{P}(\delta)$ is proper.

Due to the nature of the pair $(\mathbf{u}, \delta) \in \mathcal{X} \times [0, \infty)$, sampling can be achieved by a two-component Metropolis-within-Gibbs (MwG) algorithm. There is a range of possible parametrizations for this

MwG algorithm, perhaps the most natural of which is the so-called *centered algorithm* (CA), [34]. This scheme alternates between simulating from $\mathbf{u}|\mathbf{y}, \delta$ and $\delta|\mathbf{y}, \mathbf{u}$ using Metropolis-Hastings steps. Each pair of such simulations is one algorithmic iteration of a prescribed number k_{max} . For specific models the simulation from the two conditionals can be done directly, without Metropolis-Hastings, in which case the resultant algorithm is the Gibbs sampler. Note that the model structure implies that δ and \mathbf{y} are conditionally independent given \mathbf{u} , that is $\delta|\mathbf{y}, \mathbf{u} \equiv \delta|\mathbf{u}$. This is the defining property of the so-called *centered parameterisation* of a hierarchical model, [34].

In practice the inverse problem and the algorithm are discretized and Bayesian inference is implemented in finite dimensions. We then have two sources of error in the estimated posterior distribution: a) the approximation error due to the discretization of the unknown and the forward problem, that is the discretization bias, discussed in a general Bayesian (non-hierarchical) inverse problem setting in [13]; b) the Monte Carlo error due to the use of a Markov chain Monte Carlo method to sample the discretized posterior distribution. Assuming that the discretization level of the unknown is N , we have that the total error is of the order

$$\frac{1}{N^s} + \frac{C(N)}{\sqrt{k_{max}}}, \quad (4)$$

for some $s > 0$ which relates to the quality of approximation of the unknown and forward problem, and $C(N)$ which depends on the mixing properties of the particular algorithm used to probe the posterior. This picture allows the practitioner to get a rough idea how to distribute the computational budget by balancing investments in higher discretization levels with investments in longer chains in order to achieve the desired error level in the estimated posterior distribution. In reality, of course, the constants that multiply these rates will be relevant and hard to determine.

There are four principal motivations for formulating the inverse problem and the simulation algorithms in infinite dimensions, while using consistent discretizations (in the sense of numerical analysis, see subsection 1.2) for the numerical implementation. First, such formulation is often more faithful to the mathematical model that we wish to learn from the data. Second, it makes the inference comparable across different levels of discretization, so that the estimation of the model with increasing values of N corresponds to a reduction in the discretization bias at the cost of additional computation. Third, the prior distribution on hyperparameters, such as δ , represents the same prior beliefs across different levels of discretization. On the contrary, when the finite-dimensional model is not a consistent discretization of an infinite-dimensional one, the prior on the hyperparameters might contain an amount of information that depends on the level of discretization chosen; see for example the last paragraph in subsection 1.2.2 below. Finally, practically useful algorithms can be designed for moderate or even small values of N by studying their behaviour at the asymptotic limit $N \rightarrow \infty$. In fact, it is usually unrealistic to try to obtain practically useful theoretical results on the convergence of Markov chain Monte Carlo for sampling non-trivial targets, unless such asymptotic regimes are constructed and invoked. This is precisely the case with the Gibbs sampler and related MwG algorithms, which are particularly hard to analyse (see for example [32]). Similarly, conceiving of Metropolis-Hastings methods in the infinite-dimensional limit leads to algorithms with provably dimension-independent convergence properties, whilst standard methods have convergence properties which degenerate with increased refinement of the discretization; see [14] and discussion therein.

In this paper we investigate theoretically and numerically the performance of MwG algorithms in the asymptotic regime of large N . In order to have a mathematically tractable analysis, we focus on linear inverse problems, see subsection 1.1. For these models, and under a commonly adopted prior on δ , the MwG becomes a Gibbs sampler. We establish a result on the mean drift and diffusion of the δ -chain in CA, which has the informal interpretation that $C(N)$ is of the order $N^{1/2}$. An immediate consequence of this result, is that in order to minimize the total error in (4), k_{max} should be scaled like N^{1+2s} , whilst for algorithms for which $C(N)$ is uniformly bounded with respect to N , the same error level can be achieved by scaling k_{max} like N^{2s} ; we expect this to be the case for

the non-centered algorithm proposed later in this section. We emphasize that although we prove this result for the linear model and for a specific prior on δ , a detailed understanding of the ideas underlying our proofs, indicates that most of the details of the model, including linearity, and the prior used on δ , do not really affect the validity of our main finding, that is, that CA deteriorates with N . The fundamental reason why this algorithm becomes unusable for large N is an absolute continuity property, a high-level description of which we now provide. Note, however, that proving the result in such generality is definitely beyond the scope of this paper.

In the infinite-dimensional limit, δ is an almost sure property of $\mathbf{u}|\delta \sim \mathcal{N}(0, \delta^{-1}\mathbf{C}_0)$. This means that a single draw of \mathbf{u} contains infinite information about the value of δ that generated it. In measure-theoretic terms, it means that the prior measures $\mathbb{P}(\mathbf{u}|\delta)$ and $\mathbb{P}(\mathbf{u}|\delta')$ for $\delta \neq \delta'$ are mutually singular, [16, Remark 2.10]. Recalling that we work under assumptions which imply that $\mathbf{u}|\mathbf{y}, \delta$ is absolutely continuous with respect to $\mathbf{u}|\delta$, we deduce that δ is also an almost sure property of $\mathbf{u}|\mathbf{y}, \delta$. As a result, iterative simulation from the distributions, $\mathbf{u}|\mathbf{y}, \delta$ and $\delta|\mathbf{y}, \mathbf{u}$, will fail in ever changing the initial value of δ . On the other hand, recall that we also work under assumptions that imply that \mathbf{u} , hence δ , are not perfectly identifiable from the data. Therefore, $\delta|\mathbf{y}$ is non-degenerate (provided the prior is non-degenerate) and hence any single value of δ has zero probability under the data. Concatenating, we have that when iteratively simulating from $\mathbf{u}|\mathbf{y}, \delta$ and $\delta|\mathbf{y}, \mathbf{u}$, the values of \mathbf{u} will be changing along the iterations, but will be in fact sampled from a subspace which has probability zero under $\mathbb{P}(\mathbf{u}|\mathbf{y})$. In other words CA is *reducible* in infinite dimensions and will fail to sample from $\mathbf{u}, \delta|\mathbf{y}$. Iterative conditional sampling of the finite-dimensional approximation of $\mathbf{u}, \delta|\mathbf{y}$, will be able to obtain samples from the (approximated) posterior distribution of δ , but will suffer from increasingly slow *mixing* as the discretization level N increases. In fact, the dependence between the discretized unknown parameter u and δ increases with N , and becomes infinitely strong in the limit $N \rightarrow \infty$; it is this dependence that slows down the MwG.

In order to alleviate the undesirable effects of the strong dependence between the prior on \mathbf{u} and δ , using intuition from [34, 37], we reparametrize the prior by writing $\mathbf{u} = \delta^{-\frac{1}{2}}\mathbf{v}$ where $\mathbf{v} \sim \mathcal{N}(0, \mathbf{C}_0)$ and $\delta \sim \mathbb{P}(\delta)$. This results in a MwG algorithm which alternates between a step of updating $\mathbf{v}|\mathbf{y}, \delta$ and a step of updating $\delta|\mathbf{y}, \mathbf{v}$; this is an example of a *non-centered algorithm* (NCA), [34]. Since \mathbf{v} and δ are now a priori independent, and recalling that \mathbf{u} is not perfectly identified by the data, the dependence of these two parameters is not perfect conditionally on the data. Thus, the NCA is *irreducible* in infinite dimensions and is thus robust with respect to the discretization level N . Hence, for NCA we expect that $C(N)$ is uniformly bounded with respect to N ; we show numerical evidence in support of this statement in section 5.

1.1 The linear case - modelling and notation

We will concentrate on the linear inverse problem case with gamma priors on δ which has the convenient property of conditional conjugacy. Specifically, we restrict our attention to the case $\mathcal{G} = \mathbf{K}$ where $\mathbf{K} : \mathcal{X} \rightarrow \mathcal{Y}$ is a bounded linear operator. Then, the posterior distribution $\mathbf{u}|\mathbf{y}, \delta$ is also Gaussian

$$\mathbf{u}|\mathbf{y}, \delta \sim \mathcal{N}(\mathbf{m}_{\lambda, \delta}(\mathbf{y}), \mathbf{C}_{\lambda, \delta});$$

see [29, 27] where formulae for the posterior mean and covariance operator are provided. When the prior distribution and the noise are specified in terms of precision operators (that is, inverse covariance operators) the following expressions for the posterior mean and precision are known to hold in a range of situations in [3, 4]:

$$\mathbf{C}_{\lambda, \delta}^{-1} = \lambda \mathbf{K}^* \mathbf{C}_1^{-1} \mathbf{K} + \delta \mathbf{C}_0^{-1}, \quad (5)$$

$$\mathbf{C}_{\lambda, \delta}^{-1} \mathbf{m}_{\lambda, \delta}(\mathbf{y}) = \lambda \mathbf{K}^* \mathbf{C}_1^{-1} \mathbf{y}. \quad (6)$$

In order to introduce discretizations and their connection to the continuum limit we need some additional notation; subsection 1.2 gives specific examples of continuum models and their discretizations, where the notation introduced below is put into practice. In order to avoid a notational overload, in the development of the theory we assume that $\mathcal{X} = \mathcal{Y}$ and that the discretization levels of the unknown and the data are the same. This assumption is not crucial to our results and we refer to the PhD thesis [2, section 4.5] for the more general statements. Furthermore, in section 5, we present numerical examples corresponding to both $\mathcal{Y} = \mathcal{X}$ with an increasing discretization level which is the same for both the unknown and the data, and $\mathcal{Y} = \mathbb{R}^M$ for some fixed M , whilst the dimension of the discretization of the unknown is increased. The case $\mathcal{Y} = \mathcal{X}$ arises for example when we observe the whole unknown function subject to blurring and noise, while the case $\mathcal{Y} = \mathbb{R}^M$ can arise when we have available blurred and noisy observations of the unknown at only M spatial locations (see subsection 1.2.2). The two cases can also arise if we work in the spectral domain, depending on the availability of observations of a full or only a partial spectral expansion of a blurred noisy version of the unknown.

We denote by $\langle \cdot, \cdot \rangle_{\mathbb{R}^N}$ and $\|\cdot\|_{\mathbb{R}^N}$ the (possibly scaled) Euclidean inner product and norm in \mathbb{R}^N and by $\|\cdot\|_{2,N}$ the induced operator norm for $N \times N$ matrices. Throughout the paper we assume that this norm and inner product on \mathbb{R}^N are scaled so that, formally, the large N limit recovers the norm and inner product on the Hilbert space when, for example, spectral or finite difference approximations are made. Henceforward, we use boldface and regular typeface letters to distinguish between infinite and finite-dimensional objects respectively. We assume that we have a way of computing discretizations $y \in \mathbb{R}^N$ of the observation \mathbf{y} and replace the operators \mathbf{K}, \mathbf{C}_0 and \mathbf{C}_1 by the $N \times N$ matrices K, C_0 and C_1 respectively, which arise from a consistent, in the sense of numerical analysis, family of approximations of the corresponding operators. In this finite-dimensional setting, the unknown is $u \in \mathbb{R}^N$ and it is assigned a finite-dimensional Gaussian prior, $u|\delta \sim \mathcal{N}(0, \delta^{-1}C_0)$. The noise distribution has Lebesgue density and the corresponding log-likelihood is quadratic in u . Thus, standard Bayesian linear theory (see e.g. [28]) implies that the posterior is also Gaussian, $u|y, \delta \sim \mathcal{N}(m_{\lambda, \delta}(y), C_{\lambda, \delta})$, where $m_{\lambda, \delta}(y)$ and $C_{\lambda, \delta}^{-1}$ solve equations (5) and (6) where the boldface infinite-dimensional quantities are replaced by the corresponding finite-dimensional regular typeface quantities.

Bayesian modelling for finite-dimensional approximations of linear inverse problems using Gaussian priors and noise models was recently carried out in [6]. The approach consisted in simultaneous inference for the unknown u and the hyper-parameters λ and δ . We will concentrate on simultaneous inference on u and δ only, since λ can be efficiently estimated from a single high dimensional realization of the data, for example using quadratic variation. We again refer the interested reader to the PhD thesis [2, Chapter 4] for theoretical and numerical results on the large N behaviour of λ when considered as part of the inference problem; we stress here that for low-dimensional data, the inference on λ is non-trivial. In [6], a standard conditionally conjugate prior was used for the hyper-parameter, $\delta \sim \text{Gamma}(\alpha_0, \beta_0)$, which in this type of finite-dimensional Gaussian models is known to lead to a gamma conditional posterior distribution, [7, Chapter 5.2]

$$\delta|y, u \sim \text{Gamma}(\alpha_0 + \frac{N}{2}, \beta_0 + \frac{1}{2}\|C_0^{-\frac{1}{2}}u\|_{\mathbb{R}^N}^2). \quad (7)$$

The inference for this model was carried out using CA which in this case is a Gibbs sampler (see Algorithm 1 in section 2 below), since both conditional distributions $u|y, \delta$ and $\delta|y, u$ belong to known parametric families and can be sampled directly. One of the main aims of this paper is to analyze the convergence of this algorithm in the large N limit. We also aim to exhibit via numerical simulations, the deterioration of the performance of CA in the large N limit, as well as the benefits of reparametrizing the prior and using the corresponding NCA (see Algorithm 2 in section 2 below).

1.2 Examples of consistent discretizations

In order to aid the understanding of the paper and in anticipation of the subsequent developments, we briefly describe two methods for passing from the continuum infinite-dimensional model in \mathcal{X} to a discrete model in \mathbb{R}^N . Here and elsewhere in the paper, we define a Gaussian white noise in \mathbb{R}^N to be a random variable ζ given as $\zeta = \sum_{j=1}^N \zeta_j e_j$, where $\{e_j\}_{j=1}^N$ is a basis in \mathbb{R}^N which is orthonormal in the possibly scaled Euclidean inner product $\langle \cdot, \cdot \rangle_{\mathbb{R}^N}$, and $\{\zeta_j\}_{j \in \mathbb{N}}$ is a sequence of independent standard Gaussian random variables in \mathbb{R} .

1.2.1 Spectral truncation

Let $\{e_j\}_{j \in \mathbb{N}}$ be a complete orthonormal basis in \mathcal{X} . An element $\mathbf{w} \in \mathcal{X}$ can be identified with the sequence $\{w_j\}_{j \in \mathbb{N}}$ of coefficients $w_j := \langle \mathbf{w}, e_j \rangle$ and by Parseval's identity the Hilbert space norm of \mathbf{w} can be replaced by the ℓ_2 -norm of the sequence of coefficients (similarly for the inner product). One can then discretize \mathbf{w} by replacing it with $w \in \text{span}\{e_1, \dots, e_N\}$ which is identified with the truncated sequence of coefficients $\{w_1, \dots, w_N\} \in \mathbb{R}^N$. The ℓ_2 -norm and inner product are then replaced by the Euclidean norm and inner product. Let $\Sigma : \mathcal{X} \rightarrow \mathcal{X}$ be a bounded operator which is diagonalizable in $\{e_j\}_{j \in \mathbb{N}}$ with eigenvalues $\{\mu_j^\Sigma\}_{j \in \mathbb{N}}$. The operator Σ can be identified with the sequence $\{\mu_j^\Sigma\}_{j \in \mathbb{N}}$ and we can discretize Σ at level N by replacing it with the finite rank operator which is identified with the $N \times N$ diagonal matrix $\Sigma = \text{diag}(\mu_1^\Sigma, \dots, \mu_N^\Sigma)$. If $\mathbf{x} \sim \mathcal{N}(0, \Sigma)$ is a Gaussian random variable in \mathcal{X} , we can discretize by replacing \mathbf{x} with $x \in \text{span}\{e_1, \dots, e_N\}$ which is identified with a random variable with distribution $\mathcal{N}(0, \Sigma)$ in \mathbb{R}^N . Equivalently, x is identified with $\Sigma^{\frac{1}{2}} x_0$ where x_0 is a Gaussian white noise in \mathbb{R}^N with respect to the standard orthonormal basis of Euclidean space. For more details see subsection 4.1.

1.2.2 Finite differences approximation

Let $\mathcal{X} = L^2(\text{I})$, $\text{I} = (0, 1)$, and denote by \mathcal{A}_0 the negative Laplacian densely defined on \mathcal{X} with domain $H^2(\text{I}) \cap H_0^1(\text{I})$, that is with Dirichlet boundary conditions. We discretize the domain I using a grid of N equally spaced points $\{\frac{1}{N+1}, \dots, \frac{N}{N+1}\}$; we can restrict our attention to the interior points due to the Dirichlet boundary conditions. We define the inner product and norm in \mathbb{R}^N

$$\langle u, v \rangle_{\mathbb{R}^N} = \frac{1}{N+1} \sum_{j=1}^N u_j v_j \quad \text{and} \quad \|u\|_{\mathbb{R}^N} = \left(\frac{1}{N+1} \sum_{j=1}^N u_j^2 \right)^{\frac{1}{2}}.$$

Note that the natural orthonormal basis on the N -dimensional space of grid points with respect to the above norm and inner product is $\{e_j\}_{j=1}^N$, with $e_j = \{\sqrt{N+1} \delta_{ij}\}_{i=1}^N$, where δ_{ij} is Kronecker's delta. For a function \mathbf{u} in I which vanishes on the boundary, we consider its discretization on the grid, hence $u_j = \mathbf{u}(\frac{j}{N+1})$. We thus have a discrete approximation of \mathcal{X} with norm and inner product which are the discrete analogues of the L^2 -norm and inner product. We use finite differences to discretize \mathcal{A}_0 . In particular, we replace \mathcal{A}_0 by the $N \times N$ matrix

$$\mathcal{A}_0 = (N+1)^2 \begin{bmatrix} 2 & -1 & 0 & \dots & 0 \\ -1 & 2 & -1 & \ddots & \vdots \\ 0 & \ddots & \ddots & \ddots & 0 \\ \vdots & \ddots & -1 & 2 & -1 \\ 0 & \dots & 0 & -1 & 2 \end{bmatrix}.$$

If $\mathbf{z} \sim \mathcal{N}(0, \Sigma)$ is a Gaussian random variable in \mathcal{X} where Σ is a function of \mathcal{A}_0 (for example a power), we discretize \mathbf{z} by considering the N -dimensional random variable $z = \Sigma^{\frac{1}{2}} z_0$ defined on the

grid, where Σ is the corresponding function of the matrix \mathcal{A}_0 and z_0 is a Gaussian white noise with respect to $\{e_j\}_{j=1}^N$.

In subsection 5.2 we consider subsampling at a set of M equally spaced points amongst the N grid points, where $\frac{N+1}{M+1}$ is a nonnegative power of 2. To this end, we define the matrix $P \in \mathbb{R}^{M \times N}$ by

$$P_{i,j} = \begin{cases} 1, & \text{if } j = i \frac{N+1}{M+1} \\ 0, & \text{otherwise.} \end{cases}$$

The matrix P maps the vector of values on the fine grid $\{\mathbf{u}(\frac{j}{N+1})\}_{j=1}^N$ to the subsampled vector of the values on the coarse grid $\{\mathbf{u}(\frac{i}{M+1})\}_{i=1}^M$. If we fix M and let N increase, then P corresponds to a discretization of the operator $\mathbf{P} : C(\mathbf{I}) \rightarrow \mathbb{R}^M$ defined as M pointwise evaluations at the points $x_i = \frac{i}{M+1}$, $i = 1, \dots, M$, $(\mathbf{P}\mathbf{u})_i = \mathbf{u}(\frac{i}{M+1})$, for any continuous function \mathbf{u} . A formal calculation suggests that the adjoint of the pointwise evaluation operator at $x \in \mathbf{I}$, is an operator mapping $r \in \mathbb{R}$ to $r\delta_x$, where δ_x is the Dirac distribution at x . This suggests that $\mathbf{P}^* : \mathbb{R}^M \rightarrow C(\mathbf{I})$, maps $r \in \mathbb{R}^M$ to the linear combination of Dirac distributions $r_1\delta_{x_1} + \dots + r_M\delta_{x_M}$. At the same time the matrix $P^T \in \mathbb{R}^{N \times M}$ maps the vector of values on the coarse grid $\{\mathbf{u}(\frac{i}{M+1})\}_{i=1}^M$ to a vector in \mathbb{R}^N which is zero everywhere except from the $i \frac{N+1}{M+1}$ -th components where it is equal to $\mathbf{u}(\frac{i}{M+1})$, $i = 1, \dots, M$. Combining, and in order to capture the effect of the Dirac distribution at the locations $\frac{i}{M+1}$, we have that \mathbf{P}^* should be discretized using the matrix $(N+1)P^T$.

Note that if $\mathcal{N}(0, \delta^{-1}\mathcal{T}^{-1})$ is used as a prior on $u|\delta$ at level N , where \mathcal{T} is the $N \times N$ tridiagonal matrix in the definition of \mathcal{A}_0 , then this corresponds to having a prior with covariance matrix $(N+1)^2\delta^{-1}\mathcal{A}_0$. In particular, if $\delta \sim \text{Gamma}(\alpha_0, \beta_0)$, then we have that $\frac{1}{(N+1)^2}\delta \sim \text{Gamma}(\alpha_0, (N+1)^2\beta_0)$ where in the large N limit the last gamma distribution converges to a point mass at zero, while \mathcal{A}_0 approximates \mathcal{A}_0 . This means that as $N \rightarrow \infty$ the correlation structure of the prior is described by the limiting \mathcal{A}_0 but with an amplitude which becomes larger and larger with ever increasing confidence; in other words as N grows the prior on $u|\delta$ looks increasingly flatter.

1.3 Notation

We use subscripts to make explicit the dependence of the δ -chain on the discretization level N and superscripts to denote the iteration number in the Gibbs sampler. For a random variable x which depends on the mutually independent random variables z_1 and z_2 , we use $\mathbb{E}^{z_1}[x]$ to denote the expectation of x with respect to z_1 for fixed z_2 . We use $x_1 \stackrel{\mathcal{L}}{=} x_2$ to denote that the random variables x_1 and x_2 have the same law. Finally, for two sequences of positive numbers $\{s_j\}$ and $\{t_j\}$, we use the notation $s_j \asymp t_j$ to mean that s_j/t_j is bounded away from zero and infinity uniformly in j .

1.4 Paper structure

In the next section we present the centered Gibbs and non-centered MwG algorithms in our assumed linear conjugate setting; we also discuss the option of integrating u out of the data likelihood and the resulting marginal algorithm. In section 3 we present our main result on the deterioration of the centered Gibbs sampler which holds under certain assumptions made at the discrete level and which are stated explicitly in the same section. Our discrete level assumptions are typically inherited from Assumptions 3.1 on the underlying infinite-dimensional model also stated in section 3, when consistent numerical discretizations are used. In section 4 we exhibit three classes of linear inverse problems satisfying our assumptions on the underlying infinite-dimensional model. For the first two of these classes, that is a class of mildly ill-posed and a class of severely ill-posed linear inverse problems both in a simultaneously diagonalizable setting, we also explicitly prove that our discrete level assumptions are inherited from the infinite-dimensional assumptions when discretizing via spectral truncation (see subsections 4.1 and 4.2). In section 5 we present numerical evidence

supporting our theory and intuition on the deterioration of the centered algorithm and the merits of using the non-centered algorithm, using both spectral truncation (subsection 5.1) and discretization via finite differences and subsampling (subsection 5.2). The main body of the paper ends with concluding remarks in section 6, while the Appendix in section 7 contains the proof of our main result as well as several technical lemmas.

2 Sampling algorithms

We now present in more detail the different algorithms for sampling $u, \delta|y$ in linear hierarchical inverse problems, and provide a high-level comparison of their relative merits in the asymptotic regime of large N .

2.1 Centered Algorithm (CA)

We first provide pseudo-code for the most natural algorithm for sampling $u, \delta|y$ in this linear conjugate setting, that is the centered Gibbs sampler used in [6] and discussed in section 1.

Algorithm 1.

0. Initialize $\delta^{(0)}$ and set $k = 0$;
1. $u^{(k)} \sim \mathcal{N}(m_{\lambda, \delta^{(k)}}(y), \mathcal{C}_{\lambda, \delta^{(k)}})$;
2. $\delta^{(k+1)} \sim \text{Gamma}(\alpha_0 + \frac{N}{2}, \beta_0 + \frac{1}{2} \|C_0^{-\frac{1}{2}} u^{(k)}\|_{\mathbb{R}^N}^2)$;
3. Set $k = k + 1$. If $k < k_{max}$ return to step 1, otherwise stop.

2.2 Non-centered Algorithm (NCA)

We now formulate in more detail the non-centered algorithm introduced in section 1. We define the algorithm in the infinite-dimensional setting, and then discretize it. We reparametrize the prior by writing $\mathbf{u} = \delta^{-\frac{1}{2}} \mathbf{v}$, where now $\mathbf{v} \sim \mathcal{N}(0, \mathbf{C}_0)$, and the observation model becomes

$$\mathbf{y} = \delta^{-\frac{1}{2}} \mathbf{K} \mathbf{v} + \boldsymbol{\eta}. \quad (8)$$

The MwG sampler is used to sample $\mathbf{v}, \delta|\mathbf{y}$ by iteratively sampling from the two conditionals. Recall from the discussion on CA in section 1, that $\delta|\mathbf{y}, \mathbf{u} \equiv \delta|\mathbf{u}$ and note that $\delta|\mathbf{y}, \mathbf{v}$, no longer simplifies to $\delta|\mathbf{v}$, since even conditionally on \mathbf{v} , δ and \mathbf{y} are dependent; this is the non-centered property in the hierarchical model, [34]. Additionally, note that a practically useful way to sample from $\mathbf{v}|\mathbf{y}, \delta$, which recycles available code for CA, is to first sample $\mathbf{u}|\mathbf{y}, \delta$, as in CA, and then transform \mathbf{u} to \mathbf{v} via $\mathbf{v} = \delta^{\frac{1}{2}} \mathbf{u}$. Finally, for reasons of efficiency described below, we prefer to sample $\tau = \delta^{-\frac{1}{2}}$ instead of δ directly. In order to obtain the same Bayesian model as the one before the transformation, the prior distribution for τ should be the one obtained from the prior on δ after the $1/\sqrt{\delta}$ transformation, that is a square root of an inverse-gamma distribution. Of course, we can deterministically calculate $\delta = 1/\tau^2$ after each such update, to get δ -samples and proceed to the next conditional simulation in the algorithm.

The finite-dimensional discretization of the algorithm is obtained in the same way as CA. We notice that the log-likelihood is quadratic in τ , for given \mathbf{v} . We can exploit this property to sample τ efficiently. The conditional posterior $\tau|\mathbf{y}, \mathbf{v}$ is not Gaussian, because the prior on τ is not Gaussian,

hence for our numerical results we replace direct simulation from the conditional with a Metropolis-Hastings step that targets the conditional. Given that the conditional posterior is the product of the prior and the conditional likelihood, and we expect the likelihood to be the dominant term of the two, we use the likelihood, seen as a function of τ , as a proposal density in the Metropolis-Hastings step. The likelihood as a function of τ is Gaussian $\mathcal{N}(r_{\lambda,v}, q_{\lambda,v}^2)$, where

$$\frac{1}{q_{\lambda,v}^2} = \lambda \|\mathcal{C}_1^{-\frac{1}{2}} K v\|_{\mathbb{R}^N}^2, \quad \frac{r_{\lambda,v}}{q_{\lambda,v}^2} = \lambda \langle K^* \mathcal{C}_1^{-1} y, v \rangle_{\mathbb{R}^N}, \quad (9)$$

hence easy to simulate from. Proposals generated in this way are immediately rejected if negative, and if not they are accepted according to the Metropolis-Hastings ratio that by construction only involves the prior density. Note that the same complication would arise had we chosen to work with δ instead of τ , since $\delta|y, v$ is also not a known distribution. The difference in that case is that there is no apparent good proposal density for the Metropolis-Hastings step, since the likelihood is not a known distribution as a function of δ .

We use the following Gibbs sampler, where $p(\cdot)$ denotes the density of the square root of the inverse-gamma distribution with parameters α_0, β_0 :

Algorithm 2.

- 0) Initialize $\tau^{(0)}$, calculate $\delta^{(0)} = 1/(\tau^{(0)})^2$ and set $k = 0$;
- 1) $u^{(k)} \sim \mathcal{N}(m_{\lambda, \delta^{(k)}}(y), \mathcal{C}_{\lambda, \delta^{(k)}})$;
 $v^{(k)} = (\delta^{(k)})^{\frac{1}{2}} u^{(k)}$;
- 2) propose $\tau \sim \mathcal{N}(r_{\lambda, v^{(k)}}, q_{\lambda, v^{(k)}}^2)$;
 if $\tau \leq 0$ reject; if $\tau > 0$ accept with probability $\frac{p(\tau)}{p(\tau^{(k)})} \wedge 1$ otherwise reject;
 if τ accepted set $\tau^{(k+1)} = \tau$, otherwise set $\tau^{(k+1)} = \tau^{(k)}$;
 $\delta^{(k+1)} = 1/(\tau^{(k+1)})^2$;
- 3) Set $k = k + 1$. If $k < k_{max}$ return to step 1, otherwise stop.

2.3 Marginal Algorithm (MA)

Given that \mathbf{u} (hence $\mathbf{K}\mathbf{u}$) and $\boldsymbol{\eta}$ are independent Gaussian random variables, the marginal distribution of the data \mathbf{y} given δ is also Gaussian,

$$\mathbf{y}|\delta \sim \mathcal{N}(0, \delta^{-1} \mathbf{K} \mathcal{C}_0 \mathbf{K} + \lambda^{-1} \mathcal{C}_1).$$

One can then use Bayes' theorem to get that

$$\mathbb{P}(\delta|\mathbf{y}) \propto \mathbb{P}(\mathbf{y}|\delta) \mathbb{P}(\delta).$$

This distribution can be sampled using the random walk Metropolis (RWM) algorithm. In order to get samples from $\mathbf{u}, \delta|\mathbf{y}$, we alternate between drawing $\delta|\mathbf{y}$ and updating $\mathbf{u}|\mathbf{y}, \delta$. Furthermore, it is beneficial to the performance of the RWM, to sample $\log(\delta)|\mathbf{y}$ instead of $\delta|\mathbf{y}$; of course, samples from $\log(\delta)|\mathbf{y}$ can be deterministically transformed to samples from $\delta|\mathbf{y}$. The resultant algorithm is what we call the *marginal algorithm* (MA). MA in the discrete level is as follows, where $p(\cdot)$ now denotes the density of the logarithm of a gamma distribution with parameters α_0, β_0 and $\rho = \log(\delta)$:

Algorithm 3.

- 0) Initialize $\rho^{(0)}$ and set $k = 0$;
- 1) $u^{(k)} \sim \mathcal{N}(m_{\lambda, \delta^{(k)}}(y), \mathcal{C}_{\lambda, \delta^{(k)}})$;
- 2) propose $\rho \sim \mathcal{N}(\rho^{(k)}, s^2)$;
 accept with probability $\frac{\mathbb{P}(y|\exp(\rho))p_0(\rho)}{\mathbb{P}(y|\exp(\rho^{(k)}))p_0(\rho^{(k)})} \wedge 1$ otherwise reject;
 if ρ accepted set $\rho^{(k+1)} = \rho$, otherwise set $\rho^{(k+1)} = \rho^{(k)}$;
 set $\delta^{(k+1)} = \exp(\rho^{(k+1)})$;
- 3) Set $k = k + 1$. If $k < k_{max}$ return to step 1, otherwise stop.

We follow the rule-of-thumb proposed in [19] and choose the RWM proposal variance s^2 to achieve an acceptance probability around 44%.

2.4 Contrasting the methods

As discussed in section 1, and is formally shown in section 3, CA will deteriorate as the discretization level of the unknown, N , becomes larger. To get a first understanding of this phenomenon in the linear-conjugate setting, note that the $\text{Gamma}(\alpha_0, \beta_0)$ distribution has mean and variance $\alpha_0\beta_0^{-1}$ and $\alpha_0\beta_0^{-2}$ respectively. Hence, for any $\mu > 0$, as N grows, the random variable $\text{Gamma}(\alpha_0 + \frac{N}{2}, \beta_0 + \mu\frac{N}{2})$ behaves like a Dirac distribution centred on μ^{-1} . Furthermore, we will show that, because of the consistency of the approximation of the operators defining the Bayesian inverse problem, together with scaling of the norms on \mathbb{R}^N to reproduce the Hilbert space norm limit, it is natural to assume that

$$\|\mathcal{C}_0^{-\frac{1}{2}} u^{(k)}\|_{\mathbb{R}^N}^2 \simeq (\delta^{(k)})^{-1} N.$$

Using the limiting behaviour of the gamma distribution described above, this means that as the dimension N increases, we have $\delta^{(k+1)} \simeq \delta^{(k)}$ hence the δ -chain makes very small moves and slows down.

In contrast, both conditionals $\mathbf{u}|\mathbf{y}, \delta$ and $\delta|\mathbf{y}, \mathbf{v}$ sampled in NCA are non-degenerate even in the infinite-dimensional limit. Our numerical results show that this reparametrization is indeed robust with respect to the increase in dimension (see section 5), although establishing formally that a spectral gap exists for NCA in this limit is beyond the scope of this paper.

Similarly, both distributions $\mathbf{u}|\mathbf{y}, \delta$ and $\delta|\mathbf{y}$ sampled in MA are non-degenerate in the continuum limit, hence MA is robust with respect to N . Moreover, MA is optimal with respect to the dependence between the two components of the algorithm, since the δ -chain is independent of the u -draws; there is a loss of efficiency due to the use of RWM to sample $\delta|\mathbf{y}$, but provided the proposal variance is optimally tuned, this will only have a minor effect in the performance of MA. For these reasons, in section 5 we use the optimally tuned MA as the gold standard with which we compare the performance of CA and NCA. Nevertheless, we stress here that:

- i) MA requires at each iteration the potentially computationally expensive calculation of the square root and the determinant of the precision matrix of $y|\delta$. This makes the implementation of MA in large scale linear inverse problems less straightforward compared to CA and NCA.
- ii) even though we view MA as a gold, albeit potentially expensive, standard in our linear setting, for nonlinear problems MA is not available. On the contrary, CA and NCA are straightforward to extend to the nonlinear case (see Section 6); this is one of the principal motivations for studying the optimal parametrization of Gibbs sampling in this context.

3 Theory

In this section we present our theory concerning the behaviour of CA as the discretization level increases, in the linear inverse problem setting introduced in subsection 1.1. We first formulate our assumptions on the underlying infinite-dimensional model as well as a corresponding set of discrete-level assumptions, before presenting our main result on the large N behaviour of Algorithm 1.

3.1 Assumptions

We work under the following assumptions on the underlying infinite-dimensional linear inverse problem:

Assumptions 3.1.

- i) For any $\lambda, \delta > 0$, we have $\mathbf{m}_{\lambda, \delta}(\mathbf{y}) \in \mathcal{D}(\mathcal{C}_0^{-\frac{1}{2}})$ \mathbf{y} -almost surely; that is, the posterior mean belongs to the Cameron-Martin space of the prior on $\mathbf{u}|\delta$;
- ii) $\mathcal{C}_1^{-\frac{1}{2}} \mathbf{K} \mathcal{C}_0 \mathbf{K}^* \mathcal{C}_1^{-\frac{1}{2}}$ is trace-class; that is, the prior is sufficiently regularizing.

Assumption 3.1(ii) implies the second and third conditions of the Feldman-Hajek theorem [17, Theorem 2.23]. Together with Assumption 3.1(i), they thus imply that \mathbf{y} -almost surely $\mathbf{u}|\mathbf{y}, \delta$ is absolutely continuous with respect to $\mathbf{u}|\delta$ and hence the infinite-dimensional intuition on the behaviour of CA described in section 1 applies.

In the following, we assume that \mathcal{C}_0 and \mathcal{C}_1 are positive definite $N \times N$ matrices which are the discretizations of the positive definite operators \mathbf{C}_0 and \mathbf{C}_1 respectively, and the $N \times N$ matrix \mathbf{K} is the discretization of the bounded operator \mathbf{K} . Our analysis of the δ -chain is valid under the following assumptions at the discrete level:

Assumptions 3.2.

- i) For almost all data \mathbf{y} , for any $\lambda, \delta > 0$, there exists a constant $c_1 = c_1(\mathbf{y}; \lambda, \delta) \geq 0$, independent of N , such that

$$\|\mathcal{C}_0^{-\frac{1}{2}} \mathbf{m}_{\lambda, \delta}(\mathbf{y})\|_{\mathbb{R}^N} \leq c_1;$$

- ii) there exists a constant $c_2 \geq 0$, independent of N and \mathbf{y} , such that

$$\text{Tr}(\mathcal{C}_1^{-\frac{1}{2}} \mathbf{K} \mathcal{C}_0 \mathbf{K}^* \mathcal{C}_1^{-\frac{1}{2}}) \leq c_2.$$

These assumptions are typically inherited from Assumptions 3.1 when consistent discretizations are used; see subsection 1.2 and section 4 for more details and examples.

3.2 Main Result

We now present our main result on the behaviour of Algorithm 1 in the asymptotic regime of large N . We start by noting that the two steps of updating $u|\mathbf{y}, \delta$ and $\delta|\mathbf{y}, u$ in Algorithm 1, can be compressed to give one step of updating δ and involving the noise in the u update. Indeed, we denote by $\delta_N^{(k+1)}$ the δ -draw in the $k+1$ iteration of the Gibbs sampler where the problem is discretized in \mathbb{R}^N . This draw is made using the previous draw of $u|\mathbf{y}, \delta$, which assuming that $\delta_N^{(k)} = \delta$, is denoted by $u_\delta^{(k)}$ and can be written as

$$u_\delta^{(k)} = m_{\lambda, \delta}(\mathbf{y}) + \mathcal{C}_{\lambda, \delta}^{\frac{1}{2}} \zeta, \quad (10)$$

where ζ is an N -dimensional Gaussian white noise representing the fluctuation in step 1, and $\mathcal{C}_{\lambda,\delta}, m_{\lambda,\delta}$ are given by the formulae (5), (6) respectively. Hence we have

$$\delta_N^{(k+1)} \sim \text{Gamma}(\alpha_0 + \frac{N}{2}, \beta_0 + \frac{1}{2} \|\mathcal{C}_0^{-\frac{1}{2}} u_\delta^{(k)}\|_{\mathbb{R}^N}^2). \quad (11)$$

Assumptions 3.2 ensure that the squared norm appearing in (11) behaves like $\delta^{-1}N$, as assumed in the discrete level intuition discussed in subsection 2.4. This is made precise in the following lemma which forms the backbone of our analysis and is proved in subsection 7.2.

Lemma 3.3. *Under Assumptions 3.2, for any $\lambda, \delta > 0$ we have,*

$$\beta_0 + \frac{1}{2} \|\mathcal{C}_0^{-\frac{1}{2}} u_\delta^{(k)}\|_{\mathbb{R}^N}^2 = \delta^{-1} \frac{N}{2} + \delta^{-1} \sqrt{\frac{N}{2}} W_{1,N} + F_N(\delta), \quad (12)$$

where *i) $W_{1,N}$ only depends on the white noise ζ in (10), has mean zero and variance one, higher order moments which are bounded uniformly in N , and converges weakly to a standard normal random variable as $N \rightarrow \infty$; ii) $F_N(\delta)$ depends on the data y and y -almost surely has finite moments of all positive orders uniformly in N (where the expectation is taken with respect to ζ).*

Combining with the scaling property of the gamma distribution as in the intuition described in subsection 2.4, we show that as the dimension increases the δ -chain makes increasingly smaller steps, and quantify the scaling of this slowing down. Indeed, we prove that for large N the δ -chain makes moves which on average are of order N^{-1} with fluctuations of order $N^{-\frac{1}{2}}$. As a result, it takes $\mathcal{O}(N)$ steps for the δ -chain to move by $\mathcal{O}(1)$.

Theorem 3.4. *Let $\lambda > 0$ and consider Algorithm 1 under Assumptions 3.2. In the limit $N \rightarrow \infty$, we have almost surely with respect to y and where all the expectations are taken with respect to the randomness in the algorithm:*

i) the expected step in the δ -chain scales like $\frac{2}{N}$, that is, for any $\delta > 0$,

$$\frac{N}{2} \mathbb{E} \left[\delta_N^{(k+1)} - \delta_N^{(k)} | \delta_N^{(k)} = \delta \right] = (\alpha_0 + 1)\delta - f_N(\delta; y)\delta^2 + \mathcal{O}(N^{-\frac{1}{2}}),$$

where $f_N(\delta; y)$ is bounded uniformly in N . In particular, if there exists $f(\delta; y) \in \mathbb{R}$ such that $f_N(\delta; y) \rightarrow f(\delta; y)$, then

$$\frac{N}{2} \mathbb{E} \left[\delta_N^{(k+1)} - \delta_N^{(k)} | \delta_N^{(k)} = \delta \right] = (\alpha_0 + 1)\delta - f(\delta; y)\delta^2 + \mathcal{O}(1);$$

ii) the variance of the step also scales like $\frac{2}{N}$ and in particular for any $\delta > 0$,

$$\frac{N}{2} \text{Var} \left[\delta_N^{(k+1)} - \delta_N^{(k)} | \delta_N^{(k)} = \delta \right] = 2\delta^2 + \mathcal{O}(N^{-\frac{1}{2}}).$$

Remark 3.5.

i) The proof of Theorem 3.4 can be found in subsection 7.1 in the Appendix. Equation (17) is a key identity, as it very clearly separates the three sources of fluctuation in the draw $\delta_N^{(k+1)}$, that is, the fluctuation in the Gaussian-draw $u|y, \delta$, the fluctuation in the gamma-draw $\delta|y, u$ and the fluctuation in the data.

ii) $f_N(\delta; y) := \mathbb{E}^\zeta[F_N(\delta; y)]$, where F_N is defined in the proof of Lemma 3.3. The assumption on the convergence of $f_N(\delta; y)$ is trivially satisfied under Assumptions 3.2, if the discretization scheme used is such that if the vector $x \in \mathbb{R}^N$ and the $N \times N$ matrix T are the discretizations at level N of $x \in \mathcal{X}$ and the linear operator \mathbf{T} respectively, then $\|Tx\|_{\mathbb{R}^N}$ is a non-decreasing sequence. This is the case for example in spectral truncation methods, when \mathbf{T} is diagonalizable in the orthonormal basis used (see subsection 1.2.1).

Theorem 3.4 suggests that

$$\delta_N^{(k+1)} - \delta_N^{(k)} \approx \frac{2}{N} \left((\alpha_0 + 1) \delta_N^{(k)} - f_N(\delta_N^{(k)}; y) (\delta_N^{(k)})^2 \right) + \frac{2\delta_N^{(k)}}{\sqrt{N}} \Xi, \quad (13)$$

where Ξ is a real random variable with mean zero and variance one. In the case where f_N has a limit, the last expression looks like the Euler-Maruyama discretization of the stochastic differential equation

$$d\delta = (\alpha_0 + 1 - f(\delta; y)\delta)dt + \sqrt{2}\delta dW, \quad (14)$$

where $W = W(t)$ is a standard Brownian motion, with time step $\frac{2}{N}$. This is another manifestation of the fact that it takes $\mathcal{O}(N)$ steps for the δ -chain to make a move of $\mathcal{O}(1)$ size.

Note that (13) implies that the *expected square jumping distance* of the Markov chain for δ generated by CA is $\mathcal{O}(1/N)$. Recall (see for example [38] for a recent account) that this distance is defined as $\mathbb{E}[(\delta_N^{(k+1)} - \delta_N^{(k)})^2]$, where $\delta_N^{(k)}$ is drawn from the stationary distribution. Hence, it is the expected squared step of the chain in stationarity. It is easy to check that it equals $2\text{Var}(\delta_N^{(k)})(1 - \text{Corr}(\delta_N^{(k)}, \delta_N^{(k+1)}))$, where again all quantities are computed in stationarity. Although the expected square jumping distance is a sensible and practically useful measure of efficiency of a Markov chain, there is no explicit result that links it to the variance of Monte Carlo averages formed by using the output of the chain. This variance will not only depend on autocorrelation at other lags, but also on the function being averaged. Still, it gives a rough idea: if the autocorrelation function associated to the identity function is geometrically decaying, with lag-1 autocorrelation ρ_N , then the variance of the sample average of k_{max} , $\delta_N^{(k)}$ values in stationarity will be $\text{Var}(\delta_N^{(k)})(1 + \rho_N)/((1 - \rho_N)k_{max})$. The point here is that ρ_N behaves like $1 - c/N$, for some c , but $\text{Var}(\delta_N^{(k)})$ is $\mathcal{O}(1)$. Hence, the Monte Carlo error associated with k_{max} draws in stationarity is $\mathcal{O}(\sqrt{N/k_{max}})$.

4 Examples satisfying our assumptions

We now present three families of linear inverse problems satisfying Assumptions 3.1 on the underlying continuum model: a family of mildly ill-posed inverse problems, where the operators defining the problem are simultaneously diagonalizable, [25]; a family of severely ill-posed inverse problems again in a diagonal setting, [26, 4]; and a family of mildly ill-posed inverse problems in a nondiagonal setting, [3]. We expect that Assumptions 3.2, will be satisfied by consistent discretizations of these models. Indeed, we show that our discrete level assumptions are satisfied if we discretize the two diagonal examples using spectral truncation (see subsection 1.2.1). Furthermore, in section 5 we provide numerical evidence that our ideas also apply in nondiagonal settings and when using other discretization schemes, in particular discretization via finite difference approximations (see subsection 1.2.2). We do not prove that discretization via finite differences satisfies our discrete level assumptions as it is beyond the scope of this paper; we expect however this to be the case.

4.1 Linear mildly ill-posed simultaneously diagonalizable inverse problem

We consider the linear inverse problem setting of subsection 1.1, where \mathbf{K}, \mathbf{C}_0 and \mathbf{C}_1 commute with each other and $\mathbf{K}^* \mathbf{K}, \mathbf{C}_0$ and \mathbf{C}_1 are simultaneously diagonalizable with common complete orthonormal eigenbasis $\{\mathbf{e}_j\}_{j \in \mathbb{N}}$. Note that we do not assume that \mathbf{K} and \mathbf{C}_1 are compact, but we do assume that $\mathbf{K}^* \mathbf{K}$ and \mathbf{C}_1 are both diagonalizable in $\{\mathbf{e}_j\}_{j \in \mathbb{N}}$; in particular, we allow for \mathbf{K} and \mathbf{C}_1 to be the identity. For any $\mathbf{w} \in \mathcal{X}$, let $w_j := \langle \mathbf{w}, \mathbf{e}_j \rangle$. Let Σ be a positive definite and trace class operator in \mathcal{X} which is diagonalizable in the orthonormal basis $\{\mathbf{e}_j\}_{j \in \mathbb{N}}$, with eigenvalues $\{\mu_j^\Sigma\}_{j \in \mathbb{N}}$.

Then for any $\boldsymbol{\rho} \in \mathcal{X}$, we can write a draw $\mathbf{x} \sim \mathcal{N}(\boldsymbol{\rho}, \boldsymbol{\Sigma})$ as

$$\mathbf{x} = \boldsymbol{\rho} + \sum_{j=1}^{\infty} \sqrt{\mu_j^{\boldsymbol{\Sigma}}} \gamma_j \mathbf{e}_j,$$

where γ_j are independent standard normal random variables in \mathbb{R} ; this is the Karhunen-Loeve expansion [1, Chapter III.3]. In fact, the Karhunen-Loeve expansion makes sense even if $\mu_j^{\boldsymbol{\Sigma}}$ are not summable, that is if $\boldsymbol{\Sigma}$ is not trace class in \mathcal{X} ; the expansion then defines a Gaussian measure in a bigger space than \mathcal{X} in which $\boldsymbol{\Sigma}$ is trace class. This expansion suggests that since we are in a simultaneously diagonalizable setting we can use the Parseval identity and work entirely in the frequency domain as in subsection 1.2.1. Indeed, we identify an element $\mathbf{w} \in \mathcal{X}$ with the sequence of coefficients $\{w_j\}_{j \in \mathbb{N}}$, and the norm and inner product in \mathcal{X} with the ℓ^2 -norm and inner product. Furthermore, we identify the operators $\mathbf{C}_0, \mathbf{C}_1$ and \mathbf{K} with the sequences of their eigenvalues $\{\mu_j^{\mathbf{C}_0}\}_{j \in \mathbb{N}}, \{\mu_j^{\mathbf{C}_1}\}_{j \in \mathbb{N}}$ and $\{\mu_j^{\mathbf{K}}\}_{j \in \mathbb{N}}$ respectively. Algebraic operations on the operators $\mathbf{C}_0, \mathbf{C}_1, \mathbf{K}$ are defined through the corresponding operations on the respective sequences.

We make the following assumptions on the spectral decay of \mathbf{K}, \mathbf{C}_0 and \mathbf{C}_1 :

Assumptions 4.1. *The eigenvalues of $\mathbf{K}^* \mathbf{K}, \mathbf{C}_0$ and \mathbf{C}_1 , denoted by $(\mu_j^{\mathbf{K}})^2, \mu_j^{\mathbf{C}_0}, \mu_j^{\mathbf{C}_1}$, respectively, satisfy¹*

- $(\mu_j^{\mathbf{K}})^2 \asymp j^{-4\ell}$, $\ell \geq 0$;
- $\mu_j^{\mathbf{C}_0} \asymp j^{-2\alpha}$, $\alpha > \frac{1}{2}$;
- $\mu_j^{\mathbf{C}_1} \asymp j^{-2\beta}$, $\beta \geq 0$.

Let ν be the joint distribution of \mathbf{y} and \mathbf{u} , where $\mathbf{u}|\delta \sim \mathcal{N}(0, \delta^{-1} \mathbf{C}_0)$ and $\mathbf{y}|\mathbf{u}, \delta \sim \mathcal{N}(\mathbf{K}\mathbf{u}, \lambda^{-1} \mathbf{C}_1)$. Then in this diagonal case, it is straightforward to show in the infinite-dimensional setting that the conditional posterior $\mathbf{u}|\mathbf{y}, \delta$ is ν -almost surely Gaussian, $\mathcal{N}(\mathbf{m}_{\lambda, \delta}(\mathbf{y}), \mathbf{C}_{\lambda, \delta})$, where $\mathbf{C}_{\lambda, \delta}$ and $\mathbf{m}_{\lambda, \delta}(\mathbf{y})$ satisfy (5) and (6) respectively. We make the following additional assumption:

Assumption 4.2. *The parameters α, β, ℓ in Assumptions 4.1 satisfy $2\alpha + 4\ell - 2\beta > 1$.*

We show that under Assumptions 4.1 and 4.2, Assumptions 3.1 on the underlying infinite-dimensional model are satisfied ν -almost surely. Without loss of generality assume $\delta = \lambda = 1$. For Assumption 3.1(i), we have using the Karhunen-Loeve expansion and Assumption 4.1,

$$\mathbb{E}^{\nu} \|\mathbf{C}_0^{-\frac{1}{2}} \mathbf{m}(\mathbf{y})\|^2 \leq c \mathbb{E}^{\nu} \sum_{j=1}^{\infty} \frac{j^{2\alpha - 4\ell + 4\beta}}{(j^{-4\ell + 2\beta} + j^{2\alpha})^2} (j^{-2\ell - \alpha} \zeta_j + j^{-\beta} \xi_j)^2,$$

where $\{\zeta_j\}_{j \in \mathbb{N}}, \{\xi_j\}_{j \in \mathbb{N}}$ are two independent sequences of independent standard normal random variables. The assumption $2\alpha + 4\ell - 2\beta > 1$ secures that the right hand side is finite, hence $\mathbf{m}(\mathbf{y}) \in \mathcal{D}(\mathbf{C}_0^{-\frac{1}{2}})$ ν -almost surely. For Assumption 3.1(ii), the operator $\mathbf{C}_1^{-\frac{1}{2}} \mathbf{K} \mathbf{C}_0 \mathbf{K}^* \mathbf{C}_1^{-\frac{1}{2}}$ has eigenvalues that decay like $j^{-2\alpha - 4\ell + 2\beta}$ and hence are summable by Assumption 4.2.

We define the Sobolev-like spaces $\mathcal{H}^t, t \in \mathbb{R}$: for $t \geq 0$, define

$$\mathcal{H}^t := \{\mathbf{u} \in \mathcal{X} : \|\mathbf{u}\|_{\mathcal{H}^t} := \sum_{j=1}^{\infty} j^{2t} \langle u_j, \mathbf{e}_j \rangle^2 < \infty\},$$

and for $t < 0$, $\mathcal{H}^t := (\mathcal{H}^{-t})^*$. We assume to have data of the following form:

¹ α, β not to be confused with α, β used respectively as shape and rate parameters of the gamma distribution.

Assumption 4.3. $\mathbf{y} = \mathbf{K}\mathbf{u}^\dagger + \lambda^{-\frac{1}{2}}\mathcal{C}_1^{\frac{1}{2}}\boldsymbol{\xi}$, where $\mathbf{u}^\dagger \in \mathcal{H}^{\beta-2\ell}$ is the underlying true solution and $\boldsymbol{\xi}$ is a Gaussian white noise, $\boldsymbol{\xi} \sim \mathcal{N}(0, I)$.

Note that under Assumptions 4.1, 4.2 and 4.3, it is straightforward to check that Assumption 3.1(i) is also satisfied $\boldsymbol{\xi}$ -almost surely. Indeed, using the Karhunen-Loeve expansion we have

$$\mathbb{E}\|\mathcal{C}_0^{-\frac{1}{2}}\mathbf{m}(\mathbf{y})\|^2 \leq c\mathbb{E}\sum_{j=1}^{\infty} \frac{j^{2\alpha-4\ell+4b}}{(j^{-4\ell+2\beta} + j^{2\alpha})^2} (j^{-2\ell}(u_j^\dagger)^2 + \lambda^{-\frac{1}{2}}j^{-\beta}\xi_j)^2,$$

where $\{\xi_j\}_{j \in \mathbb{N}}$ is a sequence of independent standard normal random variables. The assumption $2\alpha + 4\ell - 2\beta > 1$ together with $\mathbf{u}^\dagger \in \mathcal{H}^{\beta-2\ell}$ secure that the right hand side is finite. Assumption 3.1(ii) is independent of \mathbf{y} , hence also holds by our previous considerations.

A natural way to discretize random draws in this setup is by truncating the Karhunen-Loeve expansion which is equivalent to the spectral truncation in subsection 1.2.1. We assume to have discrete data of the form

$$y = Ku^\dagger + \lambda^{-\frac{1}{2}}\mathcal{C}_1^{\frac{1}{2}}\xi,$$

where $K, \mathcal{C}_1, u^\dagger$ and ξ are discretized as in subsection 1.2.1. The prior is also discretized using spectral truncation, $u \sim \mathcal{N}(0, \mathcal{C}_0)$. We show that Assumptions 3.2 are satisfied under Assumptions 4.1 and 4.2, for this data and discretization scheme.

By Assumption 4.1, there exists a constant $c \geq 0$ independent of N , such that

$$\mathbb{E}\|\mathcal{C}_0^{-\frac{1}{2}}m(y)\|_{\mathbb{R}^N}^2 \leq c\mathbb{E}\sum_{j=1}^N \frac{j^{2\alpha-4\ell+4b}}{(j^{-4\ell+2\beta} + j^{2\alpha})^2} (j^{-2\ell}u_j^\dagger + j^{-\beta}\xi_j)^2,$$

where the right hand side is bounded uniformly in N , since we are summing nonnegative numbers and we have seen that under Assumptions 4.2 and 4.3 the corresponding infinite series is summable. Furthermore, again by Assumption 4.1, there exists another constant $c \geq 0$ independent of N , such that

$$\text{Tr}(\mathcal{C}_1^{-\frac{1}{2}}K\mathcal{C}_0K^*\mathcal{C}_1^{-\frac{1}{2}}) \leq c\sum_{j=1}^N j^{-2\alpha-4\ell+2\beta},$$

where the right hand side is bounded uniformly in N , since we have seen that under Assumption 4.2 the corresponding infinite series is summable.

4.2 Linear severely ill-posed simultaneously diagonalizable inverse problem

We consider the setting of [26, 4], that is, a similar situation with the previous example, where instead of having $(\mu_j^K)^2 \asymp j^{-4\ell}$ we now have $(\mu_j^K)^2 \asymp e^{-2sj^b}$, for $b, s > 0$. The proof of the validity of Assumptions 3.1 ν -almost surely is identical to the proof in the previous example, where we now have the added advantage of the exponential decay of the eigenvalues of $\mathbf{K}^*\mathbf{K}$. We can also prove that for data of the form $\mathbf{y} = \mathbf{K}\mathbf{u}^\dagger + \lambda^{-\frac{1}{2}}\mathcal{C}_1^{\frac{1}{2}}\boldsymbol{\xi}$, where now it suffices to have $\mathbf{u}^\dagger \in \mathcal{X}$, Assumption 3.1 is satisfied $\boldsymbol{\xi}$ -almost surely. Finally, in a similar way to the previous example, Assumptions 3.2 are valid if we discretize this setup by spectral truncation (subsection 1.2.1).

4.3 Nondiagonal linear inverse problem

We consider the setting of [3], that is the linear inverse problem setting of subsection 1.1, where $\mathbf{K}^*\mathbf{K}, \mathcal{C}_0$ and \mathcal{C}_1 are not necessarily simultaneously diagonalizable but they are related to each other

via a range of norm equivalence assumptions expressing that $\mathbf{K} \simeq \mathbf{C}_0^\ell$ and $\mathbf{C}_1 \simeq \mathbf{C}_0^\beta$ for some $\ell, \beta \geq 0$ (see [3, Assumption 3.1]). Here \simeq is used loosely to indicate two operators which induce equivalent norms. As before let ν be the joint distribution of \mathbf{y} and \mathbf{u} , where $\mathbf{u}|\delta \sim \mathcal{N}(0, \delta^{-1}\mathbf{C}_0)$ and $\mathbf{y}|\mathbf{u}, \delta \sim \mathcal{N}(\mathbf{K}\mathbf{u}, \lambda^{-1}\mathbf{C}_1)$. Then as in the simultaneously diagonalizable case examined above, we have that the conditional posterior $\mathbf{u}|\mathbf{y}, \delta$ is ν -almost surely $\mathcal{N}(\mathbf{m}_{\lambda,\delta}(\mathbf{y}), \mathbf{C}_{\lambda,\delta})$, where $\mathbf{C}_{\lambda,\delta}$ and $\mathbf{m}_{\lambda,\delta}(\mathbf{y})$ satisfy (5) and (6) respectively (see [3, Theorem 2.1]). It is implicit in [3, Theorem 2.1] that $\mathbf{m}_{\lambda,\delta}(\mathbf{y}) \in \mathcal{D}(\mathbf{C}_0^{-\frac{1}{2}})$ ν -almost surely, hence Assumption 3.1(i) holds ν -almost surely. Assumption 3.1(ii) also holds ν -almost surely since if $\{\phi_j\}_{j \in \mathbb{N}}$ is a complete orthonormal system of eigenfunctions of \mathbf{C}_0 and $\{\mu_j^{C_0}\}_{j \in \mathbb{N}}$ the corresponding eigenvalues, by [3, Assumption 3.1(3)] we have $\|\mathbf{C}_1^{-\frac{1}{2}} \mathbf{K} \mathbf{C}_0^{\frac{1}{2}} \phi_j\|^2 \leq c \|\mathbf{C}_0^{-\frac{\beta}{2} + \ell + \frac{1}{2}} \phi_j\|^2 = c(\mu_j^{C_0})^{-\beta + 2\ell + 1}$ which is summable by [3, Assumption 3.1(1) and 3.1(2)]. Hence, $\mathbf{C}_1^{-\frac{1}{2}} \mathbf{K} \mathbf{C}_0^{\frac{1}{2}}$ is Hilbert-Schmidt thus $\mathbf{C}_1^{-\frac{1}{2}} \mathbf{K} \mathbf{C}_0 \mathbf{K}^* \mathbf{C}_1^{-\frac{1}{2}}$ is trace-class. We believe that Assumptions 3.2 on the discrete level are also satisfied in this example if consistent discretization methods are used, however proving this is beyond the scope of the present paper.

5 Numerical Results

We now present numerical simulations supporting our result in section 3 on the large N behaviour of CA described in subsection 2.1 and our intuition contained in subsection 2.4 on the benefits of the reparametrization described in subsection 2.2. We consider an instance and a modification of the mildly ill-posed diagonal setting presented in subsection 4.1. In subsection 5.1 we use spectral truncation (see subsections 1.2.1, 4.1) and in subsection 5.2 we use finite differences approximation (see subsection 1.2.2).

5.1 Signal in white noise model using truncated Karhunen-Loeve expansion

We consider the simultaneously diagonalizable setup described in subsection 4.1, where $\mathcal{X} = L^2(\mathbf{I})$, $\mathbf{I} = (0, 1)$. We consider the orthonormal basis $e_j(x) = \sqrt{2} \sin(j\pi x)$, $x \in \mathbf{I}$, and define the operators \mathbf{K} , \mathbf{C}_0 and \mathbf{C}_1 directly through their eigenvalues $\mu_j^K = 1$, $\mu_j^{C_0} = j^{-3}$ and $\mu_j^{C_1} = 1$, for all $j \in \mathbb{N}$, respectively. In particular, this is the *normal mean model*, in which one assumes observations of the form

$$y_i = u_i + \eta_j, \quad j \in \mathbb{N},$$

where $\eta_j \sim \mathcal{N}(0, \lambda^{-1})$ and the unknown is $\{u_j\}_{j \in \mathbb{N}} \in \ell^2$. This model is clearly equivalent to the *white noise model*,

$$\mathbf{y} = \mathbf{u} + \boldsymbol{\eta}, \tag{15}$$

where $\boldsymbol{\eta} = \lambda^{-\frac{1}{2}} \boldsymbol{\xi}$ and $\boldsymbol{\xi}$ is an unobserved Gaussian white noise, see subsection 1.2.1. Note that $\boldsymbol{\xi}$ whose covariance function is a Dirac delta function, is not realizable in the basic Hilbert space \mathcal{X} (instead \mathcal{X} is the corresponding Cameron-Martin space), but can be interpreted in process form as for example in [8, 12] in the context of inverse problems. Although it can be argued that white noise data models are unrealistic at the very smallest scales, they are a useful idealization of noise which is best thought of as a continuous process with very short correlation lengthscales; in particular if the correlation lengthscale is much smaller than the grid scale used, then it is reasonable to use a white noise model. The white noise model (15) is an important statistical model which is known to be asymptotically equivalent to several standard statistical models, for example nonparametric regression, [9, 42]. It is also practically relevant, since it is a nontrivial special case of the deconvolution inverse problem, [22, 36]. Finally, it gives rise to Gaussian posterior distributions which are well studied in the sense of posterior consistency, see [25, 3, 36].

Defining \mathcal{A}_0 to be the negative Laplace operator in I with Dirichlet boundary conditions, we recognize that we use a Gaussian prior with covariance operator \mathcal{C}_0 proportional to $\mathcal{A}_0^{-\frac{3}{2}}$. Assumptions 4.1 are satisfied with $\alpha = 1.5$ and $\beta = \ell = 0$; since $2\alpha + 4\ell - 2\beta = 3 > 1$, Assumption 4.2 is also satisfied. We assume that we have data produced from the underlying true signal $\mathbf{u}^\dagger(x) = \sum_{j=1}^{\infty} u_j^\dagger \sqrt{2} \sin(j\pi x)$, for $x \in I$, where $u_j^\dagger = j^{-2.25} \sin(10j)$ and $\lambda = 200$, and in particular we have that the coefficients of \mathbf{y} , are given as

$$y_j = u_j^\dagger + \lambda^{-\frac{1}{2}} \xi_j,$$

where ξ_j are standard normal random variables. It is straightforward to check that $\mathbf{u}^\dagger \in \mathcal{H}^t$ for any $t < 1.75$, hence Assumption 4.3 is also satisfied. According to the considerations in subsection 4.1, we thus have that Assumptions 3.2 hold when using the spectral truncation discretization method. This example is studied in [40] where the interest is in studying the asymptotic performance of the posterior in the small noise limit (see section 6).

We use the hierarchical setup presented in subsection 1.1 and implement Algorithms 1 (CA), 2 (NCA) and 3 (MA) contained in section 2 at discretization levels $N = 32, 512, 8192$, with hyperparameters $\alpha_0 = 1, \beta_0 = 10^{-4}$, chosen to give uninformative hyper-priors, that is, hyper-priors whose variance is much larger than their mean. Following the discussion in subsection 2.4, we view MA as the gold standard and benchmark CA and NCA against it. We use 10^4 iterations and choose $\delta^{(0)} = 1$ in all cases. In order to have fair comparisons, we use a fixed burn-in time of 10^3 iterations. We take the viewpoint that we have a fixed computational budget, hence we choose not to increase the burn-in time as N increases as one can do if infinite resources are available.

In Figure 1 we plot the true solution, the noisy data and the sample means and credibility bounds using CA and NCA for $N = 8192$. The sample means and credibility bounds at other discretization levels of the unknown are similar and are therefore omitted.




Figure 1: Left: true solution (dashed black) and noisy data (blue continuous). Middle and right: true solution (dashed black), sample mean (red continuous) and 87.5% credibility bounds (shaded area) for CA (middle) and NCA (right). Dimension is $N = 8192$.

In Figure 2 we see that for CA, in small dimensions the δ -chain has a healthy mixing, however as predicted by Theorem 3.4, as N increases it becomes increasingly slower and exhibits diffusive behaviour. This is also reflected in the density plots where we observe that as N increases, the kernel density estimates computed using CA look less and less like the density estimates computed using MA which we consider to be optimal in this setting. In Figure 3 we see that for NCA as expected, the δ -chain appears to be robust with respect to the increase in dimension; this is also reflected in the density estimates using NCA which now look very close to the ones obtained using MA for all discretization levels.

Our observations in Figures 2 and 3 are supported by the autocorrelation plots presented in Figure 4. The rate of decay of correlations in the δ -chain in CA appears to decrease as the dimension increases, and in particular for $N = 8192$ the correlations seem not to decay at all. On the contrary, the rate of decay of correlation in the δ -chain in NCA appears not to be affected by the increase in dimension and is very similar to the one in MA.




Figure 2: CA: δ -chains (top) and kernel density estimates of the posterior on δ (bottom) for dimensions $N = 32, 512$ and 8192 left to right. In dashed red in the density plots is the density estimate using MA, considered as the gold standard.

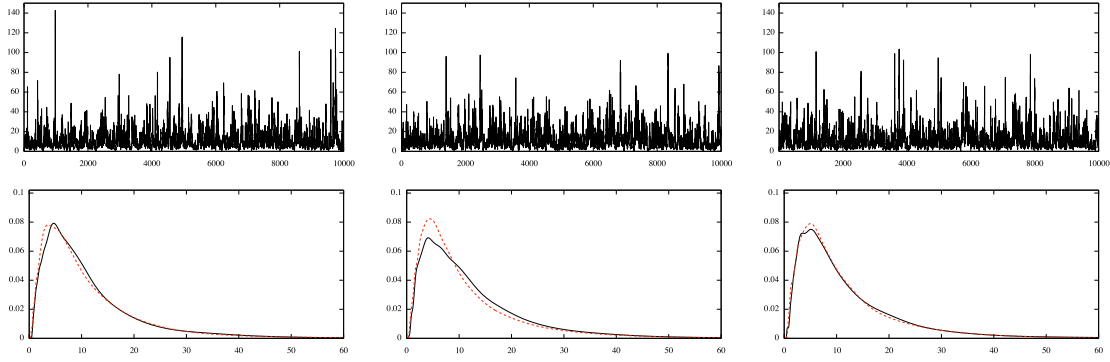


Figure 3: NCA: δ -chains (top) and kernel density estimates of the posterior on δ (bottom) for dimensions $N = 32, 512$ and 8192 left to right. In dashed red in the density plots is the density estimate using MA, considered as a gold standard.




Figure 4: Autocorrelation functions of δ -chain, dimensions 32 (black), 512 (red) and 8192 (blue); left for MA, middle for CA, right for NCA.

5.2 Linear Bayesian inverse problem with coarse data using finite difference discretization

We consider a modification of the simultaneously diagonalizable setup described in subsection 4.1, where $\mathcal{X} = L^2(I)$, $I = (0, 1)$ and we allow \mathbf{K} to map \mathcal{X} into \mathbb{R}^M and hence have data $y \in \mathbb{R}^M$. This setting is not directly covered by the theoretical analysis presented in section 3, however our theory readily generalizes to cover this setting; we refer the interested reader to the PhD thesis [2, section 4.5] for more details. The generalized analysis holds again under Assumptions 3.2 on the discrete-level based on intuition which holds for problems satisfying Assumptions 3.1 on the underlying continuum model for the unknown \mathbf{u} .

In particular, we consider the problem of recovering a true signal \mathbf{u}^\dagger , by observing a blurred version of it at M equally spaced points $\{\frac{1}{M+1}, \dots, \frac{M}{M+1}\}$, polluted by additive independent Gaussian noise of constant variance λ^{-1} . We define \mathcal{A}_0 to be the negative Laplacian with Dirichlet boundary conditions in I . We let \mathbf{P} be defined as in subsection 1.2.2 and define $\tilde{\mathbf{K}} = (\mathbf{I} + \frac{1}{100\pi^2}\mathcal{A}_0)^{-1}$, and consider the case $\mathbf{K} = \mathbf{P}\tilde{\mathbf{K}}$, $\mathcal{C}_0 = \mathcal{A}_0^{-1}$ and $\mathcal{C}_1 = I_M$ in the setting of subsection 1.1 and where I_M is the $M \times M$ identity matrix. Notice that due to the smoothing effect of $\tilde{\mathbf{K}}$, the operator \mathbf{K} is bounded in \mathcal{X} . However, due to the presence of \mathbf{P} , \mathbf{K} is not simultaneously diagonalizable with \mathcal{C}_0 .

We now check that this problem satisfies Assumptions 3.1. Indeed, assuming without loss of generality that $\lambda = \delta = 1$, by [39, Example 6.23] we have that the posterior covariance and mean satisfy (5) and (6), hence $\mathcal{C}_0^{-\frac{1}{2}}\mathbf{m}(y) = \mathcal{C}_0^{-\frac{1}{2}}(\mathcal{C}_0^{-1} + \mathbf{K}^*\mathbf{K})^{-1}\mathbf{K}^*y = (\mathbf{I} + \mathcal{C}_0^{\frac{1}{2}}\mathbf{K}^*\mathbf{K}\mathcal{C}_0^{\frac{1}{2}})^{-1}\mathcal{C}_0^{\frac{1}{2}}\mathbf{K}^*y$, where $\mathcal{C}_0^{\frac{1}{2}}\mathbf{K}^*y \in \mathcal{X}$, and $(\mathbf{I} + \mathcal{C}_0^{\frac{1}{2}}\mathbf{K}^*\mathbf{K}\mathcal{C}_0^{\frac{1}{2}})^{-1}$ is bounded in \mathcal{X} by the nonnegativity of $\mathcal{C}_0^{\frac{1}{2}}\mathbf{K}^*\mathbf{K}\mathcal{C}_0^{\frac{1}{2}}$. Furthermore, we have that $\text{Tr}(\mathcal{C}_1^{-\frac{1}{2}}\mathbf{K}\mathcal{C}_0\mathbf{K}^*\mathcal{C}_1^{-\frac{1}{2}}) = \text{Tr}(\mathbf{K}\mathcal{C}_0\mathbf{K}^*)$, which is finite since $\mathbf{K}\mathcal{C}_0\mathbf{K}^*$ is an $M \times M$ matrix.

We discretize this setup at level N , using the finite differences approximation as explained in subsection 1.2.2. In particular, we discretize \mathcal{A}_0 , \mathbf{P} and \mathbf{P}^* by replacing them with the matrices \mathcal{A}_0 , P and $(N+1)P^T$ respectively as in subsection 1.2.2; this induces a discretization of the operators \mathbf{K} and \mathcal{C}_0 by replacing them with the corresponding matrices K and \mathcal{C}_0 calculated through the appropriate functions of \mathcal{A}_0 and P . In defining K , we also replace the identity operator by the $N \times N$ identity matrix. We do not prove that this discretization scheme satisfies Assumptions 3.2, however we expect this to be the case.

We assume that we have data produced from the underlying true signal $\mathbf{u}^\dagger(x) = 0.75 \cdot \mathbf{1}_{[0.1, 0.25]}(x) + 0.25 \cdot \mathbf{1}_{[0.35, 0.38]} + \sin^4(2\pi x) \cdot \mathbf{1}_{[0.5, 1]}(x)$, $x \in I$. In particular, we construct data of the form

$$y = \mathbf{K}\mathbf{u}^\dagger + \lambda^{-\frac{1}{2}}\mathcal{C}_1^{\frac{1}{2}}\xi,$$

where $\lambda = 100$ and using a discretization level $N_c = 8192$ for the unknown; we treat this discretization level as the continuum limit.

We implement Algorithms 1 (CA), 2 (NCA) and 3 (MA) for constant number of observation points $M = 15$, and for discretization levels of the unknown $N = 15, 127, 1023$, with hyperparameters $\alpha_0 = 1, \beta_0 = 10^{-4}$, chosen to give uninformative hyper-priors, that is, hyper-priors whose variance is much larger than their mean. Following the discussion in subsection 2.4, we view MA as the gold standard and benchmark CA and NCA against it. We use 10^4 iterations and choose $\delta^{(0)} = 1$ in all cases. We again use a constant burn-in time of 10^3 iterations.

In Figure 5 we plot the true solution, the noisy data and the sample means and credibility bounds using CA and NCA for $N = 1023$. The sample means and credibility bounds at other discretization levels of the unknown are similar and are therefore omitted.

In Figure 6 we see that for CA, in small dimensions the δ -chain has a healthy mixing, however as predicted by our theory, as N increases it becomes increasingly slower and exhibits diffusive behaviour. This is also reflected in the density plots where we observe that as N increases, the kernel density estimates computed using CA look less and less like the density estimates computed




Figure 5: Left: true solution (dashed black) and discrete blurred noisy data (blue asterisks). Middle and right: true solution (dashed black), sample mean (red continuous) and 87.5% credibility bounds (shaded area) for CA (middle) and NCA (right). Dimensions of true solution and observed data are $N = 1023$ and $M = 15$ respectively.

using MA which we consider to be optimal in this setting. In Figure 7 we see that for NCA the δ -chain appears to be robust with respect to the increase in dimension; this is also reflected in the density estimates using NCA which now look very close to the ones obtained using MA for all discretization levels.




Figure 6: CA: δ -chains (top) and kernel density estimates of the posterior on δ (bottom) for dimensions $N = 15, 127$ and 1023 left to right. In dashed red in the density plots is the density estimate using MA, considered as a gold standard.

Our observations in Figures 6 and 7 are supported by the autocorrelation plots presented in Figure 8. The rate of decay of correlations in the δ -chain in CA appears to decrease as the dimension increases, and in particular for large N the correlations seem to decay very slowly. On the contrary, the rate of decay of correlations in the δ -chain in NCA appears not to be affected by the increase in dimension and is relatively close to the one in MA.

6 Conclusions

We considered a hierarchical Bayesian approach to the function-space general inverse problem (1), with Gaussian priors on the unknown function \mathbf{u} which depend on a variance-scaling parameter δ also endowed with a prior. We studied the finite-dimensional implementation of this setup and in particular, examined the mixing properties of MwG algorithms for sampling the posterior, as the discretization level N of the unknown increases. We provided measure-theoretic intuition suggesting that under natural assumptions on the underlying function space model, as N increases, CA, which is the most natural algorithm in this setting, deteriorates (see section 1). We then used this intuition to propose a reparametrization of the prior for which the resultant algorithm, NCA, is expected to




Figure 7: NCA: δ -chains (top) and kernel density estimates of the posterior on δ (bottom) for dimensions $N = 15, 127$ and 1023 left to right. In dashed red in the density plots is the density estimate using MA, considered as a gold standard.




Figure 8: Autocorrelation functions of δ -chain, dimensions 15 (black), 127 (red) and 1023 (blue); left for MA, middle for CA, right for NCA.

be robust with respect to N . In the linear-conjugate setting we formulated rigorous theory which quantifies the deterioration of CA in the asymptotic regime of large N (see section 3).

This theory holds under assumptions on the discrete level (Assumptions 3.2) which we expect to be inherited from our assumptions on the function-space model (Assumptions 3.1) when consistent discretizations are used. Indeed, we provided three families of linear inverse problems satisfying our assumptions on the underlying infinite-dimensional model (section 4), and for two of them, which are families of mildly and severely ill-posed problems in a simultaneously diagonal setting, we also showed that a spectral truncation method based on the common eigenbasis satisfies our discrete level assumptions (subsections 4.1 and 4.2). It would be interesting to show that discretization via finite differences of these examples also satisfies our discrete assumptions.

Our numerical results confirmed our theory on the deterioration of CA as well as our intuition about the robustness of NCA in the large N limit. However, for NCA the δ -chain slows down in the small noise limit. This is because even though v and δ are a priori independent, they both need to explain the data, and this creates an increasingly severer constraint as λ becomes large. Hence, δ and v concentrate near a lower dimensional manifold, where $\delta^{-\frac{1}{2}}Kv \approx y$, and the Gibbs sampler mixes poorly (see Figure 9 for a numerical illustration of this effect in the example of subsection 5.1). Although MA is robust in both the large N and the small noise limit, it can be prohibitively expensive for large scale inverse problems; new work is required to produce effective hierarchical algorithms in this small noise limit, when N is large. We have considered the interweaving method of [41], which combines in each iteration centered and non-centered draws of δ , and the partially non-centered parametrizations of [33], in which the prior is reparametrized as $u = \delta^{-\frac{1}{2}}v_t$ where $v_t \sim \mathcal{N}(0, \delta^{t-1}\mathcal{C}_0)$, for some $t \in [0, 1]$. Our numerical experimentation did not suggest significant benefits from their use, hence we do not report them here, but further investigation of these issues would be of interest.




Figure 9: Signal in white noise model - NCA for small noise, $\lambda = 200^2$, and dimension $N = 512$: δ -chain (left) and kernel density estimate of posterior on δ (right, black). In dashed red in right plot is the density estimate using MA, considered as a gold standard.

In addition to [6], a similar hierarchical setup has been considered in [40] in the *signal in white noise* model (see subsection 5.1). The authors of [40] study a different aspect of the problem, namely the asymptotic performance of the posterior distribution in the small noise limit. This is motivated by results on posterior consistency suggesting that the optimal rates of contraction are achieved by rescaling the prior depending on the size of the noise [25, 3, 26, 4]. They also study an empirical Bayes method for estimating the value of the prior scaling from the data and show that both methods achieve optimal posterior contraction rates over a range of regularity classes of the true solution. However, we have seen in this paper that the implementation of the hierarchical Bayesian method in the large dimensional limit is problematic. On the other hand, while the empirical Bayes method is appealing because of the lack of mixing issues, it involves solving an optimization problem which in more complicated models can be computationally demanding, and it does not provide uncertainty quantification of the prior scaling which may be desirable. Again we highlight the need for more research and new ideas in the small noise limit, when N is large.

An asymptotic regime which we have not investigated yet, is the case where we have a sequence of N -dimensional linear inverse problems, with the relevant matrices being consistent discretizations of

linear operators and where the size of the noise decreases as N grows larger, that is $\lambda = \lambda(N) \rightarrow \infty$ as $N \rightarrow \infty$. This is the limit of an infinite dimensional unknown which is also identifiable from the data. Since in this regime, as N grows larger the supports of both $\delta|y, u$ and $\delta|y$ shrink to zero, we expect that there will be an optimal relationship between λ and N , for which CA will not deteriorate for large N .

Our theory on the slowing down of the δ -chain can be extended to cover nonlinear Gaussian-conjugate Bayesian inverse problems and in particular the nonparametric drift estimation in SDE's setting considered in [31, 35, 30]; see [2, Chapter 4.5]. Again the main result holds under assumptions on the discrete level which we expect to be inherited by consistent discretizations from natural assumptions on the underlying infinite-dimensional model which express that the posterior is absolutely continuous with respect to the prior.

Furthermore, our infinite-dimensional intuition extends to hierarchical setups for inference on other hyper-parameters, for instance the prior regularity parameter α , where $\mathcal{C}_0 = \mathcal{A}_0^{-\alpha}$, as studied in [24]. In Figure 10 we plot autocorrelation functions for the centered MwG algorithm used in this setting and the corresponding version of the non-centered algorithm; as before we also implemented the corresponding marginal algorithm and use it as the gold standard. The underlying truth, the noise distribution and the discretization method are the same as in subsection 5.1 and we use an exponential hyper-prior on α . The idea is the same as the intuition presented in section 1, since in infinite dimensions two Gaussian measures $\mathcal{N}(0, \Sigma_1)$ and $\mathcal{N}(0, \Sigma_2)$, where Σ_1 and Σ_2 are simultaneously diagonalizable with eigenvalues $\{j^{-\alpha_1}\}_{j \in \mathbb{N}}$ and $\{j^{-\alpha_2}\}_{j \in \mathbb{N}}$ respectively, are mutually singular unless $\alpha_1 = \alpha_2$. Indeed, our numerical simulations confirm again the deterioration of the centered algorithm and the robustness of the non-centered algorithm, in the large N limit. More generally, as suggested in section 1, our intuition holds for inference on any prior on \mathbf{u} which depends on a hyper-parameter θ , when it holds that $\mathbf{u}|\mathbf{y}, \theta$ is absolutely continuous with respect to $\mathbf{u}|\theta$ almost surely with respect to the data, while $\mathbf{u}|\theta$ and $\mathbf{u}|\theta'$ are mutually singular when $\theta \neq \theta'$.




Figure 10: Autocorrelation functions of α -chain, dimensions 32 (black), 512 (red) and 8192 (blue); left for marginal, middle for centered, right for non-centered.

Returning to the general nonlinear setting discussed in section 1, we note that both Algorithms 1 and 2 are straightforward to generalize, however with a certain loss of efficiency compared to the linear-conjugate setting. The distribution of $\mathbf{u}|\mathbf{y}, \delta$ no longer belongs to a known parametric family of distributions, and thus has to be sampled using a Metropolis-Hastings (for example one based on Langevin diffusion) step. Moreover, for nonlinear inverse problems there is no longer an easy way of finding the marginal distribution $\mathbf{y}|\delta$, hence MA will not be an option. The so-called *pseudo-marginal algorithm* [5], might be an alternative for non-linear problems, and has been recently employed to perform Bayesian inference using Gaussian process priors in [18]. An interesting research direction is the comparison of the performance of the two MwG algorithms with the pseudo-marginal algorithm in both the large N and the small noise limits.

Finally, our research agenda includes the extension to the hierarchical setting of the present paper, of the analysis contained in [13] of the bias in the estimated posterior distribution due to the

discretization of the unknown and forward problem.

7 Appendix

In this section we present the proof of Theorem 3.4, as well as several technical results and lemmas. Subsection 7.1 contains the proof of Theorem 3.4, the backbone of which is Lemma 3.3 proved in subsection 7.2. In subsection 7.3 we state and prove a lemma on the negative moments of the rate parameter in the δ draw (11), which allows us to control the lower order terms arising in the proof of Theorem 3.4. Finally, in subsection 7.4, we prove several probability and linear algebra lemmas, which are useful in our analysis.

7.1 Proof of Theorem 3.4

We now prove Theorem 3.4 under Assumptions 3.2. Using the scaling property of the gamma distribution, $\text{Gamma}(\alpha, \beta) \stackrel{\mathcal{L}}{=} \beta^{-1} \text{Gamma}(\alpha, 1)$, and multiplying and dividing by $\frac{2}{N}\delta$, we can write the $\delta_N^{(k+1)}$ draw in (11) as

$$\delta_N^{(k+1)} \stackrel{\mathcal{L}}{=} \delta \frac{\Gamma_{0,N}}{\frac{2}{N}\delta(\beta_0 + \frac{1}{2}\|C_0^{-\frac{1}{2}}u_\delta^{(k)}\|_{\mathbb{R}^N}^2)} \quad (16)$$

where $\Gamma_{0,N} \sim \text{Gamma}(\alpha_0 + \frac{N}{2}, \frac{N}{2})$ is independent of y and $u_\delta^{(k)}$.

Defining $W_{2,N} := \frac{\Gamma_{0,N} - 1 - \frac{2\alpha_0}{N}}{\sqrt{\frac{2}{N} + \frac{4\alpha_0}{N^2}}}$, we have

$$\Gamma_{0,N} = 1 + \frac{2\alpha_0}{N} + \sqrt{\frac{2}{N} + \frac{4\alpha_0}{N^2}} W_{2,N},$$

where for every N , the random variable $W_{2,N}$ has mean zero and variance one, third and fourth moments bounded uniformly in N (see Lemma 7.5), and is independent of the data y and ζ , the Gaussian white noise expressing the fluctuation in $u_\delta^{(k)}$. Concatenating we get

$$\delta_N^{(k+1)} \stackrel{\mathcal{L}}{=} \delta \frac{1 + \frac{2\alpha_0}{N} + \sqrt{\frac{2}{N} + \frac{4\alpha_0}{N^2}} W_{2,N}}{1 + \sqrt{\frac{2}{N}} W_{1,N} + \frac{2}{N} F_N(\delta) \delta}, \quad (17)$$

and we are now ready to prove Theorem 3.4:

Proof. By the independence of $W_{2,N}$ and ζ and since $\mathbb{E}[W_{2,N}] = 0$, we have

$$\begin{aligned} \mathbb{E}[\delta_N^{(k+1)} - \delta_N^{(k)} | \delta_N^{(k)} = \delta] &= \delta \mathbb{E} \left[\frac{1 + \frac{2\alpha_0}{N} + \sqrt{\frac{2}{N} + \frac{4\alpha_0}{N^2}} W_{2,N}}{1 + \sqrt{\frac{2}{N}} W_{1,N} + \frac{2F_N\delta}{N}} - 1 \right] \\ &= \delta \mathbb{E}^\zeta \left[\frac{\frac{2\alpha_0}{N} - \sqrt{\frac{2}{N}} W_{1,N} - \frac{2F_N\delta}{N}}{1 + \sqrt{\frac{2}{N}} W_{1,N} + \frac{2F_N\delta}{N}} \right]. \end{aligned}$$

Using the identity $\frac{1}{1+x} = 1 - x + \frac{x^2}{1+x}$ we get

$$\begin{aligned} &\mathbb{E}[\delta_N^{(k+1)} - \delta_N^{(k)} | \delta_N^{(k)} = \delta] \\ &= \delta \mathbb{E}^\zeta \left[\left(\frac{2(\alpha_0 - F_N\delta)}{N} - \sqrt{\frac{2}{N}} W_{1,N} \right) \left(1 - \sqrt{\frac{2}{N}} W_{1,N} - \frac{2F_N\delta}{N} \right) \right] + \mathbb{E}^\zeta[e_{1,N}], \end{aligned}$$

where

$$e_{1,N} = \delta \frac{\left(\frac{2(\alpha_0 - F_N \delta)}{N} - \sqrt{\frac{2}{N}} W_{1,N} \right) \left(\frac{2W_{1,N}^2}{N} + \frac{4F_N^2 \delta^2}{N^2} + \frac{4\sqrt{2} F_N W_{1,N} \delta}{N^{\frac{3}{2}}} \right)}{1 + \sqrt{\frac{2}{N}} W_{1,N} + \frac{2F_N \delta}{N}}.$$

Using Hölder's inequality and the fact that F_N and $W_{1,N}$ have moments of all positive orders which are bounded uniformly in N , we get

$$\mathbb{E}[\delta_N^{(k+1)} - \delta_N^{(k)} | \delta_N^{(k)} = \delta] = \frac{2}{N} ((\alpha_0 + 1)\delta - \mathbb{E}^\zeta[F_N]\delta^2) + \mathcal{O}(N^{-\frac{3}{2}}) + \mathbb{E}^\zeta[e_{1,N}],$$

almost surely with respect to \mathbf{y} . For the residual $e_{1,N}$, by Cauchy-Schwarz inequality and (12), we have

$$\begin{aligned} \mathbb{E}^\zeta[e_{1,N}] &= \mathbb{E}^\zeta \left[\frac{\left(\frac{2(\alpha_0 - F_N \delta)}{N} - \sqrt{\frac{2}{N}} W_{1,N} \right) \left(W_{1,N}^2 + \frac{2}{N} F_N^2 \delta^2 + \frac{2\sqrt{2}}{N^{\frac{1}{2}}} F_N W_{1,N} \delta \right)}{\frac{N}{2\delta} (1 + \sqrt{\frac{2}{N}} W_{1,N} + \frac{2F_N \delta}{N})} \right] \\ &\leq \left(\mathbb{E} \left[\left(\frac{2(\alpha_0 - F_N \delta)}{N} - \sqrt{\frac{2}{N}} W_{1,N} \right)^2 \left(W_{1,N}^2 + \frac{2F_N^2 \delta^2}{N} + \frac{2\sqrt{2} F_N W_{1,N} \delta}{N^{\frac{1}{2}}} \right)^2 \right] \right)^{\frac{1}{2}} \\ &\quad \cdot \left(\mathbb{E} \left[\left(\beta_0 + \frac{1}{2} \|\mathcal{C}_0^{-\frac{1}{2}} u_\delta^{(k)}\|_{\mathbb{R}^N}^2 \right)^{-2} \right] \right)^{\frac{1}{2}}. \end{aligned}$$

The square root of the first expectation on the right hand side of the inequality is of order $N^{-\frac{1}{2}}$, while by Lemma 7.1 the square root of the second expectation is of order N^{-1} for almost all \mathbf{y} . Combining we get that $\mathbb{E}^\zeta[e_{1,N}] = \mathcal{O}(N^{-\frac{3}{2}})$, almost surely with respect to \mathbf{y} , hence

$$\mathbb{E}[\delta_N^{(k+1)} - \delta_N^{(k)} | \delta_N^{(k)} = \delta] = \frac{2}{N} ((1 + \alpha_0)\delta - \mathbb{E}^\zeta[F_N]\delta^2) + \mathcal{O}(N^{-\frac{3}{2}}),$$

\mathbf{y} -almost surely.

For the variance of the step, we have

$$\text{Var} [\delta_N^{(k+1)} - \delta_N^{(k)} | \delta_N^{(k)} = \delta] = \mathbb{E} [(\delta_N^{(k+1)} - \delta_N^{(k)})^2 | \delta_N^{(k)} = \delta] - \mathbb{E} [\delta_N^{(k+1)} - \delta_N^{(k)} | \delta_N^{(k)} = \delta]^2,$$

where by the first part of the proof the second term is $\mathcal{O}(N^{-2})$. Thus, we need only consider the first term, which will be shown to be $\mathcal{O}(N^{-1})$. By equation (17) we have

$$\begin{aligned} \mathbb{E} [(\delta_N^{(k+1)} - \delta_N^{(k)})^2 | \delta_N^{(k)} = \delta] &= \delta^2 \mathbb{E} \left[\left(\frac{\frac{2\alpha_0}{N} + \sqrt{\frac{2}{N} + \frac{4\alpha_0}{N^2}} W_{2,N} - \sqrt{\frac{2}{N}} W_{1,N} - \frac{2F_N \delta}{N}}{1 + \sqrt{\frac{2}{N}} W_{1,N} + \frac{2F_N \delta}{N}} \right)^2 \right] \\ &= \delta^2 \mathbb{E} \left[\frac{\frac{2W_{2,N}^2}{N} + \frac{2W_{1,N}^2}{N} + \frac{V_N}{N^{\frac{3}{2}}}}{\left(1 + \sqrt{\frac{2}{N}} W_{1,N} + \frac{2F_N \delta}{N} \right)^2} \right], \end{aligned}$$

where the random variable V_N depends only on $W_{1,N}$ and F_N and has higher order moments which are bounded uniformly in N , \mathbf{y} -almost surely (the dependence on $W_{2,N}$ disappears by the independence of $W_{2,N}$ and ζ and the fact that both $W_{2,N}$ has mean zero and variance one). Using the

identity $\frac{1}{(1+x)^2} = 1 - 2x + \frac{3x^2+2x^3}{(1+x)^2}$, we get

$$\begin{aligned} & \mathbb{E} \left[(\delta_N^{(k+1)} - \delta_N^{(k)})^2 | \delta_N^{(k)} = \delta \right] \\ &= \delta^2 \mathbb{E} \left[\left(\frac{2W_{2,N}^2}{N} + \frac{2W_{1,N}^2}{N} + \frac{V_N}{N^{\frac{3}{2}}} \right) \left(1 - 2\sqrt{\frac{2}{N}}W_{1,N} - \frac{4}{N}F_N\delta \right) \right] + \mathbb{E}[e_{2,N}], \end{aligned}$$

where

$$\begin{aligned} e_{2,N} &= \delta^2 \left(\frac{2W_{2,N}^2}{N} + \frac{2W_{1,N}^2}{N} + \frac{V_N}{N^{\frac{3}{2}}} \right) \frac{3 \left(\sqrt{\frac{2}{N}}W_{1,N} + \frac{2F_N\delta}{N} \right)^2 + 2 \left(\sqrt{\frac{2}{N}}W_{1,N} + \frac{2F_N\delta}{N} \right)^3}{\left(1 + \sqrt{\frac{2}{N}}W_{1,N} + \frac{2F_N\delta}{N} \right)^2} \\ &:= \frac{E_N\delta^2}{\left(1 + \sqrt{\frac{2}{N}}W_{1,N} + \frac{2F_N\delta}{N} \right)^2}. \end{aligned}$$

Using the fact that \mathbf{y} -almost surely $W_{1,N}$, F_N and V_N have moments of all positive orders which are bounded uniformly in N , by Hölder inequality (we do not need to consider higher order moments for $W_{2,N}$ here, because it is independent with $W_{1,N}$ and F_N , hence bounding terms involving $W_{2,N}$ does not require the use of Hölder's inequality which needs higher moments), we get that

$$\mathbb{E}[(\delta_N^{(k+1)} - \delta_N^{(k)})^2 | \delta_N^{(k)} = \delta] = \frac{2\delta^2}{N} (\mathbb{E}[W_{2,N}^2] + \mathbb{E}[W_{1,N}^2]) + \mathcal{O}(N^{-\frac{3}{2}}) + \mathbb{E}[e_{2,N}],$$

\mathbf{y} -almost surely. For the residual $e_{2,N}$, as before using Cauchy-Schwarz inequality and (12),

$$\mathbb{E}[e_{2,N}] \leq \frac{N^2}{4} (\mathbb{E}[E_N^2])^{\frac{1}{2}} \left(\mathbb{E}[(\beta_0 + \frac{1}{2} \|C_0^{-\frac{1}{2}} u_\delta^{(k)}\|_{\mathbb{R}^N}^2)^{-4}] \right)^{\frac{1}{2}}.$$

Since by Lemma 7.5 the first four moments of $W_{2,N}$ are also bounded uniformly in N , the square root of the first expectation on the right hand side is of order N^{-2} , while by Lemma 7.1 the square root of the second expectation is of order N^{-2} , for almost all \mathbf{y} . Combining we get $\mathbb{E}^\zeta[e_{2,N}] = \mathcal{O}(N^{-2})$, almost surely with respect to \mathbf{y} , hence since $\mathbb{E}[W_{1,N}^2] = \mathbb{E}[W_{2,N}^2] = 1$,

$$\mathbb{E}[(\delta_N^{(k+1)} - \delta_N^{(k)})^2 | \delta_N^{(k)} = \delta] = \frac{4\delta^2}{N} + \mathcal{O}(N^{-\frac{3}{2}}),$$

\mathbf{y} -almost surely. Concatenating, we get the result. \square

7.2 Proof of Lemma 3.3

Proof. Let $\{e_j\}_{j=1}^N$ be any orthonormal basis of \mathbb{R}^N (with respect to the possibly scaled norm $\|\cdot\|_{\mathbb{R}^N}$) and for any $w \in \mathbb{R}^N$ write $w_j := \langle w, e_j \rangle_{\mathbb{R}^N}$. We then have that $\zeta = \sum_{j=1}^N \zeta_j e_j$ where $\{\zeta_j\}_{j=1}^N$ is a sequence of independent standard normal random variables. Using (10) we have,

$$\begin{aligned} \|C_0^{-\frac{1}{2}} u_\delta^{(k)}\|_{\mathbb{R}^N}^2 &= \|C_0^{-\frac{1}{2}} m_{\lambda,\delta}(y)\|_{\mathbb{R}^N}^2 + \|C_0^{-\frac{1}{2}} C_{\lambda,\delta}^{\frac{1}{2}} \zeta\|_{\mathbb{R}^N}^2 + 2 \langle C_0^{-\frac{1}{2}} m_{\lambda,\delta}(y), C_0^{-\frac{1}{2}} C_{\lambda,\delta}^{\frac{1}{2}} \zeta \rangle_{\mathbb{R}^N} \\ &:= A_N + B_N + C_N. \end{aligned}$$

Under Assumptions 3.2, we can analyze each term as follows:

- A) by Assumption 3.2(i), for almost all data \mathbf{y} , this term and all its positive integer powers are bounded uniformly in N .

B) the second term can be written as

$$\begin{aligned}\|\mathcal{C}_0^{-\frac{1}{2}}\mathcal{C}_{\lambda,\delta}^{\frac{1}{2}}\zeta\|_{\mathbb{R}^N}^2 &= \langle \mathcal{C}_0^{-\frac{1}{2}}\mathcal{C}_{\lambda,\delta}^{\frac{1}{2}}\zeta, \mathcal{C}_0^{-\frac{1}{2}}\mathcal{C}_{\lambda,\delta}^{\frac{1}{2}}\zeta \rangle_{\mathbb{R}^N} = \langle \mathcal{C}_{\lambda,\delta}^{\frac{1}{2}}\mathcal{C}_0^{-1}\mathcal{C}_{\lambda,\delta}^{\frac{1}{2}}\zeta, \zeta \rangle_{\mathbb{R}^N} \\ &= \delta^{-1} \langle \mathcal{C}_{\lambda,\delta}^{\frac{1}{2}}(\mathcal{C}_{\lambda,\delta}^{-1} - \lambda K^* \mathcal{C}_1^{-1} K) \mathcal{C}_{\lambda,\delta}^{\frac{1}{2}}\zeta, \zeta \rangle_{\mathbb{R}^N} = \delta^{-1} \|\zeta\|_{\mathbb{R}^N}^2 - \delta^{-1} \lambda \|\mathcal{C}_1^{-\frac{1}{2}} K \mathcal{C}_{\lambda,\delta}^{\frac{1}{2}}\zeta\|_{\mathbb{R}^N}^2 \\ &:= b_{1,N} - b_{2,N},\end{aligned}$$

where

- b1) $b_{1,N} = \delta^{-1} \|\zeta\|_{\mathbb{R}^N}^2 = \frac{N}{\delta} + \frac{1}{\delta} \sum_{j=1}^N (\zeta_j^2 - 1) := \frac{N}{\delta} + \frac{\sqrt{2N}}{\delta} W_{1,N}$, where as $N \rightarrow \infty$, $W_{1,N} = \frac{1}{\sqrt{2N}} \sum_{j=1}^N (\zeta_j^2 - 1)$ converges weakly to a standard normal random variable by the Central Limit Theorem and by Lemma 7.4 has all positive integer moments bounded uniformly in N ;
- b2) for $b_{2,N}$ we have by Lemma 7.6(ii) that $\mathbb{E}^\zeta[b_{2,N}]$ is uniformly bounded in N . In fact using Lemma 7.3 together with Lemma 7.6(ii), we get that for any $p \in \mathbb{N}$, $\mathbb{E}^\zeta[b_{2,N}^p]$ is bounded independently of N .

C) for the third term we have

$$\langle \mathcal{C}_0^{-\frac{1}{2}} m_{\lambda,\delta}(y), \mathcal{C}_0^{-\frac{1}{2}} \mathcal{C}_{\lambda,\delta}^{\frac{1}{2}} \zeta \rangle_{\mathbb{R}^N} = \langle (\mathcal{C}_0^{-\frac{1}{2}} \mathcal{C}_{\lambda,\delta}^{\frac{1}{2}})^* \mathcal{C}_0^{-\frac{1}{2}} m_{\lambda,\delta}(y), \zeta \rangle_{\mathbb{R}^N} = \sum_{j=1}^N ((\mathcal{C}_0^{-\frac{1}{2}} \mathcal{C}_{\lambda,\delta}^{\frac{1}{2}})^* \mathcal{C}_0^{-\frac{1}{2}} m_{\lambda,\delta}(y))_j \zeta_j.$$

It holds that

$$\sum_{j=1}^N ((\mathcal{C}_0^{-\frac{1}{2}} \mathcal{C}_{\lambda,\delta}^{\frac{1}{2}})^* \mathcal{C}_0^{-\frac{1}{2}} m_{\lambda,\delta}(y))_j^2 = \|(\mathcal{C}_0^{-\frac{1}{2}} \mathcal{C}_{\lambda,\delta}^{\frac{1}{2}})^* \mathcal{C}_0^{-\frac{1}{2}} m_{\lambda,\delta}(y)\|_{\mathbb{R}^N}^2,$$

and we claim that the norm on the right hand side is uniformly bounded in N \mathbf{y} -almost surely. Indeed, by (5), the Cauchy-Schwarz inequality and the non-negative definiteness of the matrix $\mathcal{C}_0^{\frac{1}{2}} K^* \mathcal{C}_1^{-1} K \mathcal{C}_0^{\frac{1}{2}}$, we have

$$\begin{aligned}\|(\mathcal{C}_0^{-\frac{1}{2}} \mathcal{C}_{\lambda,\delta}^{\frac{1}{2}})^* u\|_{\mathbb{R}^N}^2 &= \langle \mathcal{C}_0^{-\frac{1}{2}} \mathcal{C}_{\lambda,\delta}^{\frac{1}{2}} u, \mathcal{C}_0^{-\frac{1}{2}} u \rangle_{\mathbb{R}^N} = \langle \delta^{-1} (I + \frac{\lambda}{\delta} \mathcal{C}_0^{\frac{1}{2}} K^* \mathcal{C}_1^{-1} K \mathcal{C}_0^{\frac{1}{2}})^{-1} u, u \rangle_{\mathbb{R}^N} \\ &\leq \|\delta^{-1} (I + \frac{\lambda}{\delta} \mathcal{C}_0^{\frac{1}{2}} K^* \mathcal{C}_1^{-1} K \mathcal{C}_0^{\frac{1}{2}})^{-1} u\|_{\mathbb{R}^N} \|u\|_{\mathbb{R}^N} \leq \delta^{-1} \|u\|_{\mathbb{R}^N}^2.\end{aligned}$$

Combining with Assumption 3.2(i) we get the claim and therefore by Lemma 7.2 below we get that the third term has \mathbf{y} -almost surely all even moments uniformly bounded in N .

We define $F_N = \beta_0 + \frac{A_N - b_{2,N} + C_N}{2}$ and observe that since all terms have bounded moments of every order uniformly in N \mathbf{y} -almost surely, Hölder's inequality secures that F_N also has bounded moments of every order uniformly in N almost surely with respect to \mathbf{y} . \square

7.3 Negative moments of the rate parameter in the δ draw

Lemma 7.1. *Let $u_\delta^{(k)}$ be as in (10), for any $\delta, \lambda > 0$. Under Assumptions 3.2, we have*

$$\mathbb{E}^\zeta \left[\left(\beta_0 + \frac{1}{2} \|\mathcal{C}_0^{-\frac{1}{2}} u_\delta^{(k)}\|_{\mathbb{R}^N}^2 \right)^{-2i} \right] = \mathcal{O}(N^{-2i}),$$

as $N \rightarrow \infty$, almost surely with respect to \mathbf{y} , for $i = 1, 2$.

Proof. Without loss of generality we consider the case $\delta = \lambda = 1$ and drop the λ and δ dependence in u, m and \mathcal{C} . To de-clutter our notation we also drop the dependence of m on the data. Since $\beta_0 \geq 0$ it suffices to show it for $\beta_0 = 0$. Formally, the random variable $\|\mathcal{C}_0^{-\frac{1}{2}}u^{(k)}\|_{\mathbb{R}^N}^2$ behaves like a chi-squared random variable with N degrees of freedom. We estimate the squared norm by a random variable Y_N of known moment generating function $M_{Y_N}(t)$, and use the following formula from [15] for the calculation of negative moments of nonnegative random variables

$$\mathbb{E}[Y_N^{-l}] = \Gamma(l)^{-1} \int_0^\infty t^{l-1} M_{Y_N}(-t) dt, \quad l \in \mathbb{N}. \quad (18)$$

We begin by showing that there exists a constant $c > 0$ independent of N such that $\|\mathcal{C}^{-\frac{1}{2}}\mathcal{C}_0^{\frac{1}{2}}v\|_{\mathbb{R}^N} \leq c\|v\|_{\mathbb{R}^N}$ for any $v \in \mathbb{R}^N$. We have,

$$\begin{aligned} \|\mathcal{C}^{-\frac{1}{2}}\mathcal{C}_0^{\frac{1}{2}}v\|_{\mathbb{R}^N}^2 &= \langle \mathcal{C}_0^{\frac{1}{2}}\mathcal{C}^{-1}\mathcal{C}_0^{\frac{1}{2}}v, v \rangle_{\mathbb{R}^N} = \langle (I + \mathcal{C}_0^{\frac{1}{2}}K^*\mathcal{C}_1^{-1}K\mathcal{C}_0^{\frac{1}{2}})v, v \rangle_{\mathbb{R}^N} \\ &= \|v\|_{\mathbb{R}^N}^2 + \|\mathcal{C}_1^{-\frac{1}{2}}K\mathcal{C}_0^{\frac{1}{2}}v\|_{\mathbb{R}^N}^2 \leq (1 + c_2)\|v\|_{\mathbb{R}^N}^2, \end{aligned}$$

by Lemma 7.6(iii). The proved claim gives the estimate

$$\|\mathcal{C}_0^{-\frac{1}{2}}u^{(k)}\|_{\mathbb{R}^N}^2 = \|\mathcal{C}_0^{-\frac{1}{2}}(m + \mathcal{C}^{\frac{1}{2}}\zeta)\|_{\mathbb{R}^N}^2 = \|\mathcal{C}_0^{-\frac{1}{2}}\mathcal{C}^{\frac{1}{2}}(\mathcal{C}^{-\frac{1}{2}}m + \zeta)\|_{\mathbb{R}^N}^2 \geq c^{-1}\|\mathcal{C}^{-\frac{1}{2}}m + \zeta\|_{\mathbb{R}^N}^2,$$

hence it suffices to show that almost surely with respect to \mathbf{y} we have $\mathbb{E}^\zeta[Y_N^{-2i}] = \mathcal{O}(N^{-2i})$, for $Y_N := \|\mathcal{C}^{-\frac{1}{2}}m + \zeta\|_{\mathbb{R}^N}^2$. Indeed, let $\{e_j\}_{j=1}^N$ be any orthonormal basis of \mathbb{R}^N (with respect to the possibly scaled norm $\|\cdot\|_{\mathbb{R}^N}$), and define $w_j := \langle w, e_j \rangle$ for any $w \in \mathbb{R}^N$. Then we have

$$Y_N = \sum_{j=1}^N ((\mathcal{C}^{-\frac{1}{2}}m)_j + \zeta_j)^2,$$

where $\zeta_j \sim \mathcal{N}(0, 1)$ are the mutually independent components of the white noise ζ and $(\mathcal{C}^{-\frac{1}{2}}m)_j$ are independent of ζ , therefore Y_N is a non-central chi-squared random variable with N degrees of freedom and non-centrality parameter $p_N := \sum_{j=1}^N (\mathcal{C}^{-\frac{1}{2}}m)_j^2 \geq 0$. The definition and properties of the non-central chi-squared distribution can be found in [21], where in particular, we find the moment generating function of Y_N

$$M_{Y_N}(t) = (1 - 2t)^{-\frac{N}{2}} \exp\left(\frac{p_N t}{1 - 2t}\right),$$

hence using (18) we have for $i = 1, 2$,

$$\begin{aligned} \mathbb{E}^\zeta[Y_N^{-2i}] &= \Gamma(2i)^{-1} \int_0^\infty t^{2i-1} (1 + 2t)^{-\frac{N}{2}} \exp\left(\frac{-p_N t}{1 + 2t}\right) dt \\ &\leq c \int_0^\infty t^{2i-1} (1 + 2t)^{-\frac{N}{2}} dt = \mathcal{O}(N^{-2i}), \end{aligned}$$

provided $N > 4i$, where the last integral can be calculated by integration by parts. \square

7.4 Technical lemmas

Lemma 7.2. *Let $\{X_j\}$ be a sequence of random variables, such that $X_j = c_j Y_j$, where the Y_j , $j \in \mathbb{N}$ are independent and identically distributed random variables with finite even moments up to order*

$2r \in \mathbb{N}$ and zero odd moments, and the c_j , $j \in \mathbb{N}$ are deterministic real numbers. Then for any $N \in \mathbb{N}$,

$$\mathbb{E}[(\sum_{j=1}^N X_j)^{2r}] \leq \kappa (\sum_{j=1}^N c_j^2)^r,$$

where $\kappa = \mathbb{E}[Y_1^{2r}] > 0$ is independent of N .

Proof. Denote by m_n the $2n$ -th moment of Y_1 , $m_n = \mathbb{E}[Y_1^{2n}]$. Observe that since by Hölder's inequality for $0 < s \leq t$, $\mathbb{E}[|Y_1|^s]^{\frac{1}{s}} \leq \mathbb{E}[|Y_1|^t]^{\frac{1}{t}}$, we have that for $n_1, \dots, n_q > 0$ such that $n_1 + \dots + n_q = r$

$$m_{n_1} \dots m_{n_q} \leq \mathbb{E}[Y_1^{2r}]^{\frac{n_1 + \dots + n_q}{r}} = \mathbb{E}[Y_1^{2r}].$$

Combining with the fact that the random variables Y_j are independent with zero odd moments,

$$\begin{aligned} \mathbb{E}[(\sum_{j=1}^N X_j)^{2r}] &= \sum_{j=1}^N c_j^{2r} m_r + \sum_{j_1 \neq j_2}^N c_{j_1}^{2(r-1)} m_{r-1} c_{j_2}^2 m_1 + \sum_{j_1 \neq j_2}^N c_{j_1}^{2(r-2)} m_{r-2} c_{j_2}^4 m_2 \\ &\quad + \dots + \sum_{j_1 \neq j_2 \neq \dots \neq j_r}^N c_{j_1}^2 c_{j_2}^2 \dots c_{j_r}^2 m_1^r \leq m_r (\sum_{j=1}^N c_j^2)^r. \end{aligned}$$

□

Lemma 7.3. For any $p \in \mathbb{N}$, there exists a constant $c = c(p) \geq 0$, independent of N such that for any centered Gaussian random variable x_N in \mathbb{R}^N , it holds

$$\mathbb{E}[\|x_N\|_{\mathbb{R}^N}^{2p}] \leq c(p) (\mathbb{E}[\|x_N\|_{\mathbb{R}^N}^2])^p.$$

Proof. Direct consequence of [17, Corollary 2.17].

□

Lemma 7.4. Let $(\gamma_j)_{j \in \mathbb{N}}$ be a sequence of independent standard normal random variables and define $G_N := \frac{1}{\sqrt{2N}} \sum_{j=1}^N (\gamma_j^2 - 1)$. Then all the positive integer moments of G_N are bounded uniformly in N .

Proof. For $k \in \mathbb{N}$, we have $\mathbb{E}[G_N^k] = \frac{1}{(2N)^{\frac{k}{2}}} \sum_{j_1, \dots, j_k}^N \mathbb{E}[(\gamma_{j_1}^2 - 1) \dots (\gamma_{j_k}^2 - 1)]$. Since $\gamma_j^2 - 1$ are independent and identically distributed with finite moments of every order, the sum on the right hand side has a dependence on N determined by the total number of non zero terms in the summation. By independence and the fact that $\mathbb{E}[\gamma_j^2 - 1] = 0$, all the terms in the sum which contain a term with an index j_i which occurs only once in the product are equal to zero. We thus have that if k is even the sum on the right hand side is of order $N^{\frac{k}{2}}$, while if k is odd it is of order $N^{\frac{k-1}{2}}$. In both cases the k -th moment of G_N is bounded uniformly in N .

□

Lemma 7.5. Let $\Gamma_N \sim \text{Gamma}(\alpha + \frac{N}{2}, \frac{N}{2})$, for $\alpha > 0$, and define

$$\Theta_N := \frac{\Gamma_N - 1 - \frac{2\alpha}{N}}{\sqrt{\frac{2}{N} + \frac{4\alpha}{N^2}}}.$$

Then the first four moments of Θ_N are bounded uniformly in N .

Proof. The random variable $\text{Gamma}(a, 1)$ has mean and variance a and third and fourth central moments $2a$ and $3a^2 + 6a$ respectively, [20]. Hence by the scaling property of the gamma distribution, $\Gamma_N \stackrel{\mathcal{L}}{=} \frac{2}{N} \text{Gamma}(\alpha + \frac{N}{2}, 1)$ has mean $1 + \frac{2\alpha}{N}$, variance $\frac{2}{N} + \frac{4\alpha}{N^2}$, and third and fourth central moments which are both of order N^{-2} . It is thus straightforward to see that Θ_N has mean zero, variance equal to one, and since the denominator in Θ_N is of order $N^{-\frac{1}{2}}$ it has third and fourth moments which are $\mathcal{O}(N^{-\frac{1}{2}})$ and $\mathcal{O}(1)$ respectively.

□

Lemma 7.6. *Under Assumptions 3.2, we have that for any $\lambda, \delta > 0$,*

- i) $\text{Tr}(\mathcal{C}_1^{-\frac{1}{2}} K \mathcal{C}_{\lambda, \delta} K^* \mathcal{C}_1^{-\frac{1}{2}}) \leq c_2 \delta^{-1}$;
- ii) $\mathbb{E}^\theta \|\mathcal{C}_1^{-\frac{1}{2}} K \mathcal{C}_{\lambda, \delta}^{\frac{1}{2}} \theta\|_{\mathbb{R}^N}^2 \leq c_2 \delta^{-1}$, where θ is a Gaussian white noise in \mathbb{R}^N ;
- iii) $\|\mathcal{C}_1^{-\frac{1}{2}} K \mathcal{C}_0^{\frac{1}{2}}\|_{2, N} \leq \sqrt{c_2}$;

where c_2 is defined in Assumption 3.2(ii).

Proof.

- i) By (5), we have

$$\mathcal{C}_1^{-\frac{1}{2}} K \mathcal{C}_{\lambda, \delta} K^* \mathcal{C}_1^{-\frac{1}{2}} = \delta^{-1} \mathcal{C}_1^{-\frac{1}{2}} K \mathcal{C}_0^{\frac{1}{2}} (I + \frac{\lambda}{\delta} \mathcal{C}_0^{\frac{1}{2}} K^* \mathcal{C}_1^{-1} K \mathcal{C}_0^{\frac{1}{2}})^{-1} \mathcal{C}_0^{\frac{1}{2}} K^* \mathcal{C}_1^{-\frac{1}{2}},$$

hence the fact that for any matrix $A \in \mathbb{R}^{N \times N}$ it holds $\text{Tr}(A(I + cA^*A)A^*) \leq \text{Tr}(AA^*)$ for any $c > 0$, together with Assumption 3.2(ii) give the claim.

- ii) It is well known that for $x \sim \mathcal{N}(0, \Sigma)$, $\mathbb{E}\|x\|_{\mathbb{R}^N}^2 = \text{Tr}(\Sigma)$. Since for $\theta \sim \mathcal{N}(0, I)$ we have $\mathcal{C}_1^{-\frac{1}{2}} K \mathcal{C}_{\lambda, \delta}^{\frac{1}{2}} \theta \sim \mathcal{N}(0, \mathcal{C}_1^{-\frac{1}{2}} K \mathcal{C}_{\lambda, \delta} K^* \mathcal{C}_1^{-\frac{1}{2}})$, the claim follows from part (i).
- iii) It is well known that for any matrix $A \in \mathbb{R}^{N \times N}$, the Euclidean norm satisfies $\|A\|_{2, N} = \|A^*\|_{2, N} = \sqrt{\rho(A^*A)} \leq \sqrt{\text{Tr}(A^*A)}$ where $\rho(B)$ is the spectral radius of the matrix B . Hence we have $\|\mathcal{C}_1^{-\frac{1}{2}} K \mathcal{C}_0^{\frac{1}{2}}\|_{2, N} \leq \sqrt{\text{Tr}(\mathcal{C}_1^{-\frac{1}{2}} K \mathcal{C}_0 K^* \mathcal{C}_1^{-\frac{1}{2}})} \leq \sqrt{c_2}$, by Assumption 3.2(ii).

□

References

- [1] R. J. Adler. *An introduction to continuity, extrema, and related topics for general Gaussian processes*. Institute of Mathematical Statistics Lecture Notes—Monograph Series, 12. Institute of Mathematical Statistics, Hayward, CA, 1990.
- [2] S. Agapiou. Aspects of bayesian inverse problems. *Ph.D. Thesis, University of Warwick*, 2013.
- [3] S. Agapiou, S. Larsson, and A. M. Stuart. Posterior contraction rates for the Bayesian approach to linear ill-posed inverse problems. *Stochastic Processes and their Applications*, 2013.
- [4] S. Agapiou, A. M. Stuart, and Y. X. Zhang. Bayesian posterior contraction rates for linear severely ill-posed inverse problems. *arxiv:1210.1563*, 2012.
- [5] C. Andrieu and G. O. Roberts. The pseudo-marginal approach for efficient Monte Carlo computations. *The Annals of Statistics*, pages 697–725, 2009.
- [6] J. M. Bardsley. MCMC-based image reconstruction with uncertainty quantification. *SIAM J. Sci. Comput.*, 34(3):A1316–A1332, 2012.
- [7] J. M. Bernardo and A. F. M. Smith. *Bayesian theory*, volume 405. Wiley, 2009.
- [8] N. Bissantz, T. Hohage, A. Munk, and F. Ruymgaart. Convergence rates of general regularization methods for statistical inverse problems and applications. *SIAM Journal on Numerical Analysis*, 45(6):2610–2636, 2007.

- [9] L. D. Brown and M. G. Low. Asymptotic equivalence of nonparametric regression and white noise. *The Annals of Statistics*, 24(6):2384–2398, 1996.
- [10] D. Calvetti, H. Hakula, S. Pursiainen, and E. Somersalo. Conditionally Gaussian hypermodels for cerebral source localization. *SIAM Journal on Imaging Sciences*, 2(3):879–909, 2009.
- [11] D. Calvetti and E. Somersalo. Hypermodels in the Bayesian imaging framework. *Inverse Problems*, 24(3):034013, 2008.
- [12] L. Cavalier. Nonparametric statistical inverse problems. *Inverse Problems*, 24(3):034004, 2008.
- [13] S. L. Cotter, M. Dashti, and A. M. Stuart. Approximation of Bayesian inverse problems for pdes. *SIAM Journal on Numerical Analysis*, 48(1):322–345, 2010.
- [14] S. L. Cotter, G. O. Roberts, A. M. Stuart, and D. White. MCMC methods for functions: modifying old algorithms to make them faster. *Statistical Science*, 28(3):424–446, 2013.
- [15] N. Cressie, A. S. Davis, J. L. Folks, and G. E. Policello II. The moment-generating function and negative integer moments. *The American Statistician*, 35(3):148–150, 1981.
- [16] G. Da Prato. *An introduction to infinite-dimensional analysis*. Springer, 2005.
- [17] G. Da Prato and J. Zabczyk. Stochastic equations in infinite dimensions, volume 44 of encyclopedia of mathematics and its applications, 1992.
- [18] M. Filippone and M. Girolami. Pseudo-Marginal Bayesian inference for Gaussian processes. 10 2013.
- [19] A. Gelman, G. O. Roberts, and W. R. Gilks. Efficient Metropolis jumping rules. *Bayesian statistics*, 5:599–608, 1996.
- [20] N. L. Johnson, S. Kotz, and N. Balakrishnan. *Continuous univariate distributions. Vol. 1*. Wiley Series in Probability and Mathematical Statistics: Applied Probability and Statistics. John Wiley & Sons Inc., New York, second edition, 1994. A Wiley-Interscience Publication.
- [21] N. L. Johnson, S. Kotz, and N. Balakrishnan. *Continuous univariate distributions. Vol. 2*. Wiley Series in Probability and Mathematical Statistics: Applied Probability and Statistics. John Wiley & Sons Inc., New York, second edition, 1995. A Wiley-Interscience Publication.
- [22] I. M. Johnstone, G. Kerkycharian, D. Picard, and M. Raimondo. Wavelet deconvolution in a periodic setting. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)*, 66(3):547–573, 2004.
- [23] J. P. Kaipio and E. Somersalo. *Statistical and computational inverse problems*, volume 160. Springer, 2005.
- [24] B. T. Knapik, B. T. Szabó, A. W. van der Vaart, and J. H. van Zanten. Bayes procedures for adaptive inference in nonparametric inverse problems. *arXiv:1209.3628*, 2012.
- [25] B. T. Knapik, A. W. van Der Vaart, and J. H. van Zanten. Bayesian inverse problems with Gaussian priors. *The Annals of Statistics*, 39(5):2626–2657, 2011.
- [26] B. T. Knapik, A. W. van der Vaart, and J. H. van Zanten. Bayesian recovery of the initial condition for the heat equation. *Communications in Statistics-Theory and Methods*, 42(7):1294–1313, 2013.
- [27] M. S. Lehtinen, L. Paivarinta, and E. Somersalo. Linear inverse problems for generalised random variables. *Inverse Problems*, 5(4):599, 1989.

- [28] D. V. Lindley and A. F. M. Smith. Bayes estimates for the linear model. *Journal of the Royal Statistical Society. Series B (Methodological)*, pages 1–41, 1972.
- [29] A. Mandelbaum. Linear estimators and measurable linear transformations on a Hilbert space. *Zeitschrift für Wahrscheinlichkeitstheorie und Verwandte Gebiete*, 65(3):385–397, 1984.
- [30] F. van der Meulen, M. Schauer, and J. H. van Zanten. Reversible jump MCMC for nonparametric drift estimation for diffusion processes. *Computational Statistics & Data Analysis*, 2013.
- [31] O. Papaspiliopoulos, Y. Pokern, G. O. Roberts, and A. M. Stuart. Nonparametric estimation of diffusions: a differential equations approach. *Biometrika*, 99(3):511–531, 2012.
- [32] O. Papaspiliopoulos and G. O. Roberts. Stability of the Gibbs sampler for Bayesian hierarchical models. *Ann. Statist.*, 36(1):95–117, 2008.
- [33] O. Papaspiliopoulos, G. O. Roberts, and M. Sköld. Non-centered parameterisations for hierarchical models and data augmentation. In *Bayesian Statistics 7: Proceedings of the Seventh Valencia International Meeting*. Oxford University Press, USA, 2003.
- [34] O. Papaspiliopoulos, G. O. Roberts, and M. Sköld. A general framework for the parametrization of hierarchical models. *Statistical Science*, pages 59–73, 2007.
- [35] Y. Pokern, A. M. Stuart, and J. H. van Zanten. Posterior consistency via precision operators for Bayesian nonparametric drift estimation in SDEs. *Stochastic Process. Appl.*, 123(2):603–628, 2013.
- [36] K. Ray. Bayesian inverse problems with non-conjugate priors. *Electronic Journal of Statistics*, 7:2516–2549, 2013.
- [37] G. O. Roberts and O. Stramer. On inference for partially observed nonlinear diffusion models using the Metropolis–Hastings algorithm. *Biometrika*, 88(3):603–621, 2001.
- [38] C. Sherlock and G. O. Roberts. Optimal scaling of the random walk Metropolis on elliptically symmetric unimodal targets. *Bernoulli*, 15(3):774–798, 2009.
- [39] A. M. Stuart. Inverse problems: a Bayesian perspective. *Acta Numer.*, 19:451–559, 2010.
- [40] B. T. Szabó, A. W. van der Vaart, and J. H. van Zanten. Empirical Bayes scaling of Gaussian priors in the white noise model. *Electronic Journal of Statistics*, 7:991–1018, 2013.
- [41] Y. Yu and X. L. Meng. To center or not to center: that is not the question—an ancillarity–sufficiency interweaving strategy (ASIS) for boosting MCMC efficiency. *Journal of Computational and Graphical Statistics*, 20(3):531–570, 2011.
- [42] L. H. Zhao. Bayesian aspects of some nonparametric problems. *Annals of statistics*, pages 532–552, 2000.