

Probabilistic Iterative Methods for Linear Systems

Jon Cockayne* Ilse C.F. Ipsen† Chris. J. Oates‡
 Tim W. Reid§

December 24, 2020

This paper presents a probabilistic perspective on iterative methods for approximating the solution $\mathbf{x}_* \in \mathbb{R}^d$ of a nonsingular linear system $\mathbf{A}\mathbf{x}_* = \mathbf{b}$. In the approach a standard iterative method on \mathbb{R}^d is lifted to act on the space of probability distributions $\mathcal{P}(\mathbb{R}^d)$. Classically, an iterative method produces a sequence \mathbf{x}_m of approximations that converge to \mathbf{x}_* . The output of the iterative methods proposed in this paper is, instead, a sequence of probability distributions $\mu_m \in \mathcal{P}(\mathbb{R}^d)$. The distributional output both provides a “best guess” for \mathbf{x}_* , for example as the mean of μ_m , and also probabilistic uncertainty quantification for the value of \mathbf{x}_* when it has not been exactly determined. Theoretical analysis is provided in the prototypical case of a stationary linear iterative method. In this setting we characterise both the rate of contraction of μ_m to an atomic measure on \mathbf{x}_* and the nature of the uncertainty quantification being provided. We conclude with an empirical illustration that highlights the insight into solution uncertainty that can be provided by probabilistic iterative methods.

1. Introduction

The focus of this paper is on the numerical solution of a linear systems of equations

$$\mathbf{A}\mathbf{x}_* = \mathbf{b} \tag{1}$$

where $\mathbf{A} \in \mathbb{R}^{d \times d}$ is a given non-singular matrix, $\mathbf{b} \in \mathbb{R}^d$ is a non-zero vector and $\mathbf{x}_* \in \mathbb{R}^d$ is an unknown vector to be computed. The problem of solving linear systems is central

*The Alan Turing Institute (jcockayne@turing.ac.uk)

†North Carolina State University (ipsen@ncsu.edu)

‡Newcastle University and The Alan Turing Institute (chris.oates@newcastle.ac.uk)

§North Carolina State University (twreid@ncsu.edu)

to scientific computation Golub and Van Loan [2013, p103]. Solvers can be broadly categorized as either *direct*, meaning they compute \mathbf{x}_* by factorizing the matrix \mathbf{A} , or as *iterative* meaning they output a sequence of approximate solutions to \mathbf{x}_* . The focus of the present paper is on a probabilistic version of iterative methods.

There exist a wide variety of iterative methods, with the two main classes being the stationary iterative methods [Young, 1971], such as Richardson’s method and Jacobi’s method, and Krylov subspace methods [Liesen and Strakos, 2012] such as the conjugate gradient method (CG; Hestenes and Stiefel [1952]). In each case, the output of an iterative method is a sequence \mathbf{x}_m of approximations that converge, one hopes, to \mathbf{x}_* as m is increased. The methods we present in this paper lift iterative methods into probability space, replacing iterates $\mathbf{x}_m \in \mathbb{R}^d$ with iterates in $\mu_m \in \mathcal{P}(\mathbb{R}^d)$, the set of probability measures on \mathbb{R}^d .

An object of critical importance in iterative methods is the error $\mathbf{e}_m = \mathbf{x}_* - \mathbf{x}_m$, and since \mathbf{x}_* is inaccessible the error or its norm must be estimated. Error estimation for linear systems has a long history. For CG applied to a symmetric positive definite matrix \mathbf{A} , one typically bounds the \mathbf{A} -norm of the error $\|\mathbf{e}_m\|_{\mathbf{A}} = \sqrt{\mathbf{e}_m^\top \mathbf{A} \mathbf{e}_m}$. Estimates such as this may be of limited utility for two reasons: firstly, they are often conservative and may be complicated to compute, and secondly, a scalar-valued error indicator does not capture any structure that may be present in the error \mathbf{e}_m .

For a concrete motivating example, suppose that the value of \mathbf{x}_* is the input to some further computation, denoted abstractly as $F(\mathbf{x}_*)$ for $F : \mathbb{R}^d \rightarrow \mathbb{R}$, and suppose that one wishes to characterise the error $|F(\mathbf{x}_*) - F(\mathbf{x}_m)|$. It is not trivial to transfer a bound on $\|\mathbf{x}_* - \mathbf{x}_m\|$ into a practically useful estimate of this error, particularly when F is not analytically tractable. For example, if F depends only on a subset of the entries of \mathbf{x}_* for which the iterative method converges rapidly, while the other entries converge slowly, a bound on $|F(\mathbf{x}_*) - F(\mathbf{x}_m)|$ that is a function only of $\|\mathbf{x}_* - \mathbf{x}_m\|$ is likely to be extremely conservative.

The methods described herein fall in the class of Probabilistic numerical methods (PNM; [Larkin, 1972, Diaconis, 1988, Hennig et al., 2015, Cockayne et al., 2019b]). PNM are numerical methods that accept some initial distribution μ_0 as input, along with the usual quantities that specify the numerical problem (in our case, \mathbf{A} and \mathbf{b}), and produce a probability distribution μ_m as their output. The distribution μ_0 provides a means by which *a priori* information can be incorporated into a numerical method. In the case of Eq. (1), this distribution is intended to assign probability mass to subsets of \mathbb{R}^d in which \mathbf{x}_* is believed to be located, prior to any computations being performed. This information may be elicited from a domain expert or obtained in an objective manner, for instance by performing additional computations pertaining to the numerical task.

The output distribution μ_m provides a structured description of the error \mathbf{e}_m . Returning to our motivating example, this description may trivially be propagated through F without incurring a significant loss of information, for example by sampling from μ_m and passing the samples through F . The resulting probability distribution provides a kind of *uncertainty quantification* (UQ) for the unknown quantity of interest $F(\mathbf{x}_*)$, and may not suffer the same degree of conservatism of the norm-based estimators that we briefly described. While the UQ we produce has a different flavour to that traditionally

explored in the field (e.g. in Smith [2014]) it is philosophically similar; see Hennig et al. [2015] and Cockayne et al. [2019b] for a discussion of the contrast.

1.1. Related Work

There has been recent interest in the construction of PNM for the solution of Eq. (1) with contributions in Hennig et al. [2015], Bartels and Hennig [2016], Bartels et al. [2019], Cockayne et al. [2019a], Reid et al. [2020], Wegner and Hennig [2020]. With the exception of Bartels et al. [2019] these works have predominantly focused on replicating CG, and so a positive-definite \mathbf{A} is assumed. Each constructed a PNM in the Bayesian statistical framework, where the distribution μ_0 has the interpretation of a *prior* posited over some quantity related to Eq. (1) at the outset, and this distribution is updated based on the limited computations that are performed. The updating is achieved using Bayes’ theorem and the result is a *posterior* or *conditional* distribution μ_m that forms the output of the method; it is a distribution over the unknown \mathbf{x}_* that quantifies uncertainty given the limited computation performed. In Hennig et al. [2015], Bartels and Hennig [2016], Wegner and Hennig [2020] the prior was placed on the entries of \mathbf{A}^{-1} (or jointly on \mathbf{A} and \mathbf{A}^{-1}), while in Bartels et al. [2019], Cockayne et al. [2019a] the prior was placed directly on the unknown solution of Eq. (1). In each case computation consisted of projecting Eq. (1) against a set of search directions \mathbf{s}_i , $i = 1, \dots, m$, i.e. by computing $\mathbf{s}_i^\top \mathbf{A} \mathbf{x}_* = \mathbf{s}_i^\top \mathbf{b}$ and the output of the PNM was a distribution that contracts to a point mass at \mathbf{x}_* in an appropriate computational limit.

Each of these methods exploited conjugacy of Gaussian distributions under linear transformations to condition on the linear information provided by the pairs $(\mathbf{s}_i, \mathbf{s}_i^\top \mathbf{b})$, $i = 1, \dots, m$. This use of Bayes’ theorem is justified when the search directions \mathbf{s}_i are not themselves dependent on \mathbf{x}_* , the solution of Eq. (1). However, in practice these authors advocated the use of search directions generated using a Lanczos-style recursion [Liesen and Strakos, 2012, Section 2.4], meaning that the \mathbf{s}_i depend on \mathbf{x}_* via \mathbf{b} and the required assumption is violated. As remarked in Bartels et al. [2019], Cockayne et al. [2019a], this violation leads to PNM that are not strictly Bayesian and are unacceptably conservative, meaning that the “width” of the probability distribution μ_m produced by the PNM can be a gross over-estimate of the actual error, as quantified by the difference between the mean of μ_m and \mathbf{x}_* . Reid et al. [2020] addressed this deficiency by constructing a prior which corrects for the over-confidence in an empirical Bayesian fashion, though with such a prescribed prior it is difficult for other problem-specific information to be incorporated. It therefore remains an open problem to develop a PNM for the solution of Eq. (1) that allows a generic initial distribution $\mu_0 \in \mathcal{P}(\mathbb{R}^d)$ to be used and ensures the distributional output $\mu_m \in \mathcal{P}(\mathbb{R}^d)$ of the PNM is an accurate reflection of the error \mathbf{e}_m .

1.2. Contributions

This paper adopts a different strategy to the aforementioned work. Instead of applying Bayes’ theorem, we first posit an initial distribution μ_0 and iteratively update this

distribution using a transformation derived from an iterative method for the solution of Eq. (1). The initial distribution μ_0 is loosely analogous to the prior in a Bayesian approach, but since no analogue of the Bayesian update occurs these methods are not Bayesian in the standard sense. We thus refer to μ_m as a *belief distribution* rather than a posterior, following the contemporary literature on generalised Bayesian inference [Bisiri et al., 2016]. In departing from an established statistical paradigm one is required to justify, mathematically, the sense in which the uncertainty quantification provided by μ_m is meaningful. For this we purpose we leverage the recent work of Graham et al. [2020], who argued that non-Bayesian procedures can be justified if they are *calibrated*, meaning that \boldsymbol{x}_* is indistinguishable in a certain, precise sense, from any other sample drawn independently from μ_m . The contributions of this paper are therefore as follows:

- We introduce *probabilistic iterative methods*, a class of PNM derived from iterative methods for solving linear systems such as Eq. (1). These methods can be interpreted as a lifting of any standard iterative method into probability space, and are equivalent to randomising the initial iterate in a standard iterative method.
- A detailed theoretical analysis of the convergence properties of these new PNM is conducted for the class of linear stationary iterative methods, in which the next iterate is obtained by an affine transformation of the previous iterate. We prove that in this case the iterates produced are *strongly calibrated* in the sense of Graham et al. [2020] and hence can be thought of as providing reliable uncertainty quantification despite not existing in the Bayesian paradigm.
- We describe statistical tests for *weak calibration*, again following Graham et al. [2020], which can be used to test calibration of probabilistic iterative methods based on more complex iterative methods, such as Krylov methods.
- A simulation study is conducted to analyse the performance of probabilistic iterative methods in a regression problem. We examine the convergence and calibration of both linear and nonlinear probabilistic iterative methods, and highlight how the output of the method may be used to gain insight into potential numerical issues with the regression problem, such as the areas of the space in which higher error is incurred as a result of slow convergence of the iterative method.

1.3. Structure of the Paper

In Section 2 we introduce iterative methods for linear systems and describe how these may be lifted into algorithms that operate on probability space. Theoretical results concerning the convergence and uncertainty calibration of a class of analytically tractable probabilistic iterative methods are presented in Section 3, and in Section 4 we consider the general case, presenting a statistical test that can be used to assess whether the UQ provided by a probabilistic iterative method is meaningful in the sense of Graham et al. [2020]. In Section 5 we apply probabilistic iterative methods to solve a linear system arising in a regression problem. Lastly, in Section 6 we discuss the results presented and the outlook for this new class of methods.

1.4. Notation

Here the notation for the paper is established. We will work in the measurable space $(\mathbb{R}^d, \mathcal{B}(\mathbb{R}^d))$ where $\mathcal{B}(\mathbb{R}^d)$ is the standard Borel sigma-algebra for \mathbb{R}^d . Let $\mathcal{P}(\mathbb{R}^d)$ denote the set of all probability measures on $(\mathbb{R}^d, \mathcal{B}(\mathbb{R}^d))$. Bold lower-case italic letters (i.e. \mathbf{x}) will be used to denote vectors in \mathbb{R}^d and bold capital italic letters to denote matrices in $\mathbb{R}^{d \times d}$ (i.e. \mathbf{M}). Non-bold capital italic letters will denote random variables on \mathbb{R}^d (i.e. X) and lower-case Greek letters will be used to denote elements of $\mathcal{P}(\mathbb{R}^d)$.

Throughout it will be assumed that $\|\cdot\|$ is a fixed but arbitrary norm on \mathbb{R}^d . One important example is the vector \mathcal{L} -norm, given by

$$\|\mathbf{x}\|_p = \left(\sum_{i=1}^d |x_i|^p \right)^{\frac{1}{p}},$$

though we note that many of the results presented herein do not assume any particular norm, and where a specific norm *is* required this will be emphasised. This notation will also be used for the induced norm on $\mathbb{R}^{d \times d}$, given by

$$\|\mathbf{M}\| = \sup_{\|\mathbf{x}\|=1} \|\mathbf{M}\mathbf{x}\|.$$

Recall that all induced norms are sub-multiplicative, meaning that $\|\mathbf{M}\mathbf{x}\| \leq \|\mathbf{M}\| \|\mathbf{x}\|$. Let $\rho(\mathbf{M})$ denote the spectral radius of \mathbf{M} , let \mathbf{M}^\dagger denote the Moore-Penrose pseudo-inverse of \mathbf{M} , let $\text{range}(\mathbf{M})$ denote its range and $\ker(\mathbf{M})$ its kernel (or null space). For a symmetric matrix \mathbf{M} let $\lambda_{\min}(\mathbf{M})$ and $\lambda_{\max}(\mathbf{M})$ denote the smallest and largest eigenvalue of \mathbf{M} . For a positive-definite matrix \mathbf{M} , we define the weighted norm $\|\mathbf{x}\|_{\mathbf{M}} = (\mathbf{x}^\top \mathbf{M} \mathbf{x})^{\frac{1}{2}}$. Let $\mathbf{M}^{\frac{1}{2}}$ denote a matrix for which $\mathbf{M} = (\mathbf{M}^{\frac{1}{2}})^\top \mathbf{M}^{\frac{1}{2}}$. Note that this is not the typical notion of a square root, in that it will *not* be required that $(\mathbf{M}^{\frac{1}{2}})^\top = \mathbf{M}^{\frac{1}{2}}$.

For a measurable map $S : \mathbb{R}^d \rightarrow \mathbb{R}^d$ and a set $B \subset \mathbb{R}^d$, $S^{-1}[B]$ will be used to denote the preimage of B under S , i.e.

$$S^{-1}[B] = \{\mathbf{x} \in \mathbb{R}^d \text{ s.t. } S(\mathbf{x}) \in B\}.$$

For a distribution $\mu \in \mathcal{P}(\mathbb{R}^d)$, recall that the pushforward distribution $S_{\#}\mu$ is defined as $(S_{\#}\mu)(B) = \mu(S^{-1}[B])$ for each $B \in \mathcal{B}(\mathbb{R}^d)$, and thus $S_{\#}\mu \in \mathcal{P}(\mathbb{R}^d)$. The notation $\mathcal{N}(\mathbf{x}, \mathbf{\Sigma})$ will be used to denote the multivariate Gaussian distribution with mean \mathbf{x} and positive semi-definite covariance $\mathbf{\Sigma}$. The notation χ_d^2 will denote the chi-squared distribution with $d \in \mathbb{N}$ degrees of freedom. Recall that if $X \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d)$ then $\|X\|_2^2 \sim \chi_d^2$.

2. Probabilistic Iterative Methods

In this section we will first introduce iterative methods by way of the taxonomy introduced in Young [1971], before proceeding to the general definition of probabilistic iterative methods.

2.1. Iterative Methods

A general iterative method $\mathcal{I}(\mathbf{A}, \mathbf{b})$ is defined [Young, 1971, Section 3.1] by a sequence of maps $\mathcal{I} = (P_m)_{m \geq 1}$, for which $\mathbf{x}_m = P_m(\mathbf{x}_0, \dots, \mathbf{x}_{m-1}; \mathbf{A}, \mathbf{b})$. The iterative method \mathcal{I} is said to be *linear* if each P_m is linear in $\mathbf{x}_0, \dots, \mathbf{x}_{m-1}$. It is said to be *of degree s* if for all $m \geq s$ we have that P_m depends only on the s previous iterates, i.e. $P_m(\mathbf{x}_0, \dots, \mathbf{x}_{m-1}; \mathbf{A}, \mathbf{b}) = P_m(\mathbf{x}_{m-s}, \dots, \mathbf{x}_{m-1}; \mathbf{A}, \mathbf{b})$. Lastly, the method is said to be *stationary* if the maps P_m are independent of m . We will generally suppress dependence of \mathcal{I} and the P_m on \mathbf{A} and \mathbf{b} for notational convenience.

Many of the most widely used iterative methods can be expressed as methods of degree $s = 1$. For simplicity we will present the majority of the material in this paper in these terms, though the core ideas readily generalise to higher degree methods as will be discussed in Section 4 and explored in Section 5. Any iterative method \mathcal{I} of degree $s = 1$ implies a map P^m that acts only on the first s iterates to produce iterate m , as follows:

$$P^m(\mathbf{x}_0) = (P_m \circ \dots \circ P_1)(\mathbf{x}_0).$$

For first degree methods each P_m is generally a *contraction map* with fixed point \mathbf{x}_* , i.e. $P_m(\mathbf{x}_*) = \mathbf{x}_*$. Thus when the iterative method is stationary it amounts to applying a single fixed contraction map to an initial iterate until convergence.

We now present several classical examples of first-degree iterative methods; for each see Young [1971, Section 3.3]. These methods are seldom used as linear solvers in contemporary applications, but are still sometimes used in conjunction with other methods [Saad, 2003, p103].

Example 2.1 (Stationary Richardson method). This method adopts the following iteration

$$\mathbf{x}_m = \mathbf{x}_{m-1} + \omega(\mathbf{b} - \mathbf{A}\mathbf{x}_{m-1}), \quad m \geq 1$$

where $\omega > 0$ is a parameter of the method. The method is stationary and linear, with each map P_m of the form

$$P_m(\mathbf{x}) = P(\mathbf{x}) = \mathbf{G}\mathbf{x} + \mathbf{f} \tag{2}$$

where $\mathbf{G} = \mathbf{I} - \omega\mathbf{A}$ and $\mathbf{f} = \omega\mathbf{b}$.

Example 2.2 (Jacobi's method). In *Jacobi's method* it is assumed that the diagonal elements of \mathbf{A} are nonzero. The iteration takes the form

$$\mathbf{x}_m = \mathbf{D}^{-1}(\mathbf{b} - (\mathbf{A} - \mathbf{D})\mathbf{x}_{m-1}) + \mathbf{x}_{m-1}, \quad m \geq 1$$

where $\mathbf{D} = \text{diag}(\mathbf{A})$. The method is again stationary and linear. In the notation of Eq. (2). we have that $\mathbf{G} = \mathbf{I} - \omega\mathbf{D}^{-1}\mathbf{A}$ and $\mathbf{f} = \omega\mathbf{D}^{-1}\mathbf{b}$.

The next method, CG, sees significantly more use, particularly in the solution of large sparse linear systems. Whereas the above two methods are based on matrix splittings, in CG the solution \mathbf{x}_* is instead projected into a sequence of *Krylov subspaces* [Liesen and Strakos, 2012, Section 2.2] of increasing dimension. As a result it is not traditionally

viewed within the classification of Young [1971]¹. Nevertheless CG is currently seen as an iterative method and may be categorised within the taxonomy presented above, albeit rather degenerately since CG provably converges in a finite number $m' \leq d$ of iterations, and P_m is undefined for $m > m'$.

Example 2.3 (Conjugate Gradient Method). In CG the iteration is of the form

$$\begin{aligned}\mathbf{x}_m &= \mathbf{x}_{m-1} + \alpha_m \mathbf{s}_m \\ \alpha_m &= \frac{\mathbf{s}_m^\top \mathbf{r}_m}{\mathbf{s}_m^\top \mathbf{A} \mathbf{s}_m} \\ \mathbf{s}_{m+1} &= \mathbf{r}_m + \beta_m \mathbf{s}_m \\ \beta_m &= \frac{\mathbf{r}_m^\top \mathbf{r}_m}{\mathbf{r}_{m-1}^\top \mathbf{r}_{m-1}}\end{aligned}$$

where the initial direction \mathbf{s}_0 is taken to be \mathbf{r}_0 , and we recall that recall that $\mathbf{r}_m = \mathbf{b} - \mathbf{A}\mathbf{x}_m$. From Saad [2003, Algorithm 6.19] CG may be expressed as a three-term recurrence. Examining this we see that it is neither stationary nor linear, and is of second degree. Nevertheless in terms of its *implementation*, the algorithm requires only the storage of \mathbf{x}_m and \mathbf{r}_m to compute \mathbf{x}_{m+1} .

2.2. Lifting to Probability Space

We now introduce the central definition of this paper, that of a probabilistic iterative method. As noted above, the definition is presented in terms of a method of degree $s = 1$; extension to higher degree is considered in Section 3.3.

Definition 2.4. Let $\mathcal{I} = (P_m)_{m \geq 1}$ be an iterative method of first degree. Then the maps $P_m : \mathbb{R}^d \rightarrow \mathbb{R}^d$ can be lifted to maps $(P_m)_\# : \mathcal{P}(\mathbb{R}^d) \rightarrow \mathcal{P}(\mathbb{R}^d)$ operating on elements of $\mathcal{P}(\mathbb{R}^d)$. Then we say that $\mathcal{I}_\# = ((P_m)_\#)_{m \geq 1}$ is a *probabilistic iterative method*.

Thus probabilistic iterative methods are a class of PNMs that take as input an initial distribution $\mu_0 \in \mathcal{P}(\mathbb{R}^d)$ and return a sequence of iterates $\mu_m = (P_m)_\#(\mu_0)$. Again we note that $\mathcal{I}_\#$, and therefore μ_m , each formally depend on \mathbf{A} and \mathbf{b} , but this dependence will generally be suppressed.

The distribution μ_0 should be thought of as an initial belief about where the solution \mathbf{x}_* to the linear system might lie in \mathbb{R}^d . Thus μ_0 has a similar role to the prior distribution in the Bayesian setting. However, the iterates μ_m do not arise as a conditional distribution, and so the output from probabilistic iterative methods does not have a classical Bayesian interpretation. It is therefore crucial to ensure that the UQ provided by the method is meaningful, which will be discussed in detail in Section 3.2 (for stationary linear methods) and Section 4.2 (for general methods).

¹The discussion in Liesen and Strakos [2012, Section 2.5.7] highlights that, when the book was written, CG was still often considered a direct method owing to its convergence in $m' \leq d$ iterations; its attractive properties as an iterative method were not understood by the community until Reid [1971], who studied its use as an iterative method for large sparse linear systems. This likely explains why Young [1971] does not attempt to categorise it within his taxonomy.

Compared to the Bayesian approach to PNM [Cockayne et al., 2019b], probabilistic iterative methods are significantly easier to implement. For example, an algorithm for producing a sample from μ_m is to sample $\mathbf{x} \sim \mu_0$ and compute $P^m(\mathbf{x})$. Thus sampling from the output of a probabilistic iterative method inherits the computational efficiency and stability of the underlying classical iterative method, only multiplying the cost by the number of samples required. Conversely Bayesian PNM generally require new and often highly computationally expensive code to be developed, which can bring its own share of computational stability issues.

Unlike in a Bayesian PNM however, it is straightforward to construct a counterexample showing that the support of μ_m need not be a subset of the support of μ_0 . Thus, even if μ_0 encodes properties of the solution that are expected to hold with probability one, (for example, positivity of the elements), μ_m is not guaranteed to inherit those properties. This emphasises the need for careful analysis of the UQ of probabilistic iterative methods, which we provide in Section 3.2.

Our first theoretical result shows that if the iterates \mathbf{x}_m contract to the truth at a certain rate then the distributions μ_m contract at the same rate.

Proposition 2.5. *Let \mathcal{I} be an iterative method of first degree for solution of Eq. (1). Suppose that each P_m has error controlled by the bound*

$$\|\mathbf{x}_* - P_m(\mathbf{x})\| \leq \varphi(m)\|\mathbf{x}_* - \mathbf{x}_0\|, \quad m \geq 1 \quad (3)$$

where $\varphi : \mathbb{N} \rightarrow \mathbb{R}$ is some function independent of \mathbf{x}_0 , such that $\varphi(m) \rightarrow 0$ as $m \rightarrow \infty$. Then for any $k > 0$ and $\delta > 0$,

$$\mu_m(B_\delta^c(\mathbf{x}_*)) \leq \left(\frac{\varphi(m)}{\delta}\right)^k \int_{\mathbb{R}^d} \|\mathbf{x}_* - \mathbf{x}\|^k d\mu_0(\mathbf{x}),$$

where $B_\delta^c(\mathbf{x}_*)$ represents the complement of a $\|\cdot\|$ -ball of radius δ about \mathbf{x}_* , i.e.

$$B_\delta^c(\mathbf{x}_*) = \{\mathbf{x} \in \mathbb{R}^d : \|\mathbf{x}_* - \mathbf{x}\| \geq \delta\}.$$

Remark 2.6. Since \mathbf{x}_* is generally a fixed point of P_m the function $\varphi(m)$ in Proposition 2.5 is similar to a Lipschitz constant, though it need not be independent of \mathbf{x}_* for the result to hold.

Proof. For any $k > 0$,

$$\begin{aligned} \int_{\mathbb{R}^d} \|\mathbf{x}_* - \mathbf{x}\|^k d\mu_m(\mathbf{x}) &= \int_{\mathbb{R}^d} \|\mathbf{x}_* - P^m(\mathbf{x})\|^k d\mu_0(d\mathbf{x}) \\ &\leq \varphi(m)^k \int_{\mathbb{R}^d} \|\mathbf{x}_* - \mathbf{x}\|^k d\mu_0(d\mathbf{x}) \end{aligned}$$

where the second line follows from applying the assumed bound for the error and extracting terms independent of \mathbf{x} from the integral. Now, recall from Chebyshev's inequality

[Kallenberg, 2002, Lemma 3.1] we have that for a measure μ on \mathbb{R}^d , a μ -measurable function $f : \mathbb{R}^d \rightarrow [0, \infty)$ and scalars $\delta \in [0, \infty)$, $k \in (0, \infty)$ it holds that

$$\mu(\{\mathbf{x} \in \mathbb{R}^d : f(\mathbf{x}) \geq \delta\}) = \mu(\{\mathbf{x} \in \mathbb{R}^d : f(\mathbf{x})^k \geq \delta^k\}) \leq \frac{1}{\delta^k} \int_{\mathbb{R}^d} f(\mathbf{x})^k d\mu(\mathbf{x}).$$

Applying this in the present setting with $f(\mathbf{x}) = \|\mathbf{x}_* - \mathbf{x}\|$ we therefore have

$$\mu_m(B_\delta^c(\mathbf{x}_*)) \leq \left(\frac{\varphi(m)}{\delta}\right)^k \int_{\mathbb{R}^d} \|\mathbf{x}_* - \mathbf{x}\|^k d\mu_0(\mathbf{x})$$

as required. \square

Thus the probability mass assigned by μ_m to the region outside of a ball $B_\delta(\mathbf{x}_*)$ centred on the true solution \mathbf{x}_* vanishes as $m \rightarrow \infty$. Moreover, and again asymptotically as $m \rightarrow \infty$, the probability mass outside $B_\delta(\mathbf{x}_*)$ vanishes more rapidly when high-order moments of μ_0 exist (i.e. large k). However, Proposition 2.5 does not imply that the UQ provided by μ_m is meaningful, or even that \mathbf{x}_* is in the support of μ_m . For the UQ to be meaningful further assumptions are required on \mathcal{I} , such as those made in Section 3.2.

3. Linear Probabilistic Iterative Methods

In this section we restrict attention to *linear, stationary* iterative methods of first degree, as a richer set of theoretical results can readily be developed for this restricted set of methods. In Section 3.1 we review some classical results and describe how the probabilistic iterates μ_m can be exactly computed when μ_0 is Gaussian. In Section 3.2 we prove that these methods are strongly calibrated in the sense of Graham et al. [2020], and in Section 3.3 we discuss relaxing the stationarity and first degree assumptions.

3.1. Linear and Stationary Probabilistic Iterative Methods

In linear stationary iterative methods of first degree $P_m(\mathbf{x}_0, \dots, \mathbf{x}_{m-1}) = P(\mathbf{x}_{m-1})$, as described in Example 2.1, where

$$P(\mathbf{x}) = \mathbf{G}\mathbf{x} + \mathbf{f} \tag{4}$$

for some $\mathbf{G} \neq \mathbf{0} \in \mathbb{R}^{d \times d}$ and $\mathbf{f} \in \mathbb{R}^d$. It follows that $P^m(\mathbf{x}_0) = \mathbf{G}^m \mathbf{x}_0 + \sum_{i=0}^{m-1} \mathbf{G}^i \mathbf{f}$, where $\mathbf{G}^0 = \mathbf{I}$. We now recall some standard results for linear stationary iterative methods of first degree which will later be useful. The following is based on Young [1971, Section 3.3.2 and 3.3.5] and Saad [2003, Section 4.2].

Proposition 3.1. *Let \mathbf{A} be nonsingular, suppose that $\mathbf{G} \in \mathbb{R}^{d \times d}$ is such that $\|\mathbf{G}\| < 1$ and*

$$\mathbf{f} = (\mathbf{I} - \mathbf{G})\mathbf{A}^{-1}\mathbf{b} = (\mathbf{I} - \mathbf{G})\mathbf{x}_*. \tag{5}$$

Then the iterative method

$$\mathbf{x}_{m+1} = \mathbf{G}\mathbf{x}_m + \mathbf{f} \quad m \geq 1$$

converges to \mathbf{x}_* for all $\mathbf{x}_* \in \mathbb{R}^d$. Furthermore the error in \mathbf{x}_m is controlled by the bound

$$\|\mathbf{x}_* - \mathbf{x}_m\| \leq \|\mathbf{G}\|^m \|\mathbf{x}_* - \mathbf{x}_0\|.$$

Now we consider lifting linear stationary iterative methods of first degree into $\mathcal{P}(\mathbb{R}^d)$. Our main observation is that if μ_0 is Gaussian, then the distribution μ_m can be computed in closed form using standard formulae for linear transforms of Gaussian distributions.

Proposition 3.2. *Let \mathcal{I} be a linear, stationary, first degree iterative method. Let $\mu_0 = \mathcal{N}(\mathbf{x}_0, \Sigma_0)$. Then $\mu_m = \mathcal{N}(\mathbf{x}_m, \Sigma_m)$, where $\mathbf{x}_m = P^m(\mathbf{x}_0)$ coincides with the classical iterate, and*

$$\Sigma_m = \mathbf{G}^m \Sigma_0 (\mathbf{G}^m)^\top.$$

Furthermore we have the following bounds:

$$\begin{aligned} \|\mathbf{x}_* - \mathbf{x}_m\| &\leq \|\mathbf{G}\|^m \|\mathbf{x}_* - \mathbf{x}_0\| \\ \|\Sigma_m\| &\leq \|\mathbf{G}\|^m \|\mathbf{G}^\top\|^m \|\Sigma_0\|. \end{aligned}$$

Proof. From elementary properties of Gaussian distributions [Tong, 1990, Theorem 3.3.3] we have that $\mu_1 = \mathcal{N}(\mathbf{x}_1, \Sigma_1)$ where $\mathbf{x}_1 = \mathbf{G}\mathbf{x}_0 + \mathbf{f}$ and $\Sigma_1 = \mathbf{G}\Sigma_0\mathbf{G}^\top$. This can be continued inductively to achieve the form stated in the proposition for all $m \geq 1$. The bound on $\|\mathbf{x}_* - \mathbf{x}_m\|$ is a consequence of \mathbf{x}_m coinciding with the classical iterate and Proposition 3.1. The bound on Σ_m is direct by applying submultiplicativity of the norm $\|\cdot\|$ to $\|\mathbf{G}^m \Sigma_0 (\mathbf{G}^m)^\top\|$. \square

Remark 3.3. The bound on Σ_m in Proposition 3.2 does not require that $\|\cdot\|$ be the induced norm, only that it is submultiplicative. As a result, this applies to other matrix norms such as the Frobenius norm, which is submultiplicative but not induced.

3.2. Evaluation of Uncertainty Quantification

The crucial point that must be addressed in order for probabilistic iterative methods to be useful is whether the covariance matrix Σ_m relates meaningfully to the error $\mathbf{e}_m = \mathbf{x}_* - \mathbf{x}_m$. It is not possible to provide a satisfactory answer to this question by considering just one linear system; this would be akin to asking whether the number 3 is meaningfully related to the distribution $\mathcal{N}(0, 1)$. Therefore a collection of linear systems is required so that average-case properties can be discussed.

In Cockayne et al. [2019a] a criterion for meaningful UQ was introduced based on the *Z-statistic*, which is given by

$$Z_m(X) = \|X - \mathbf{x}_m\|_{\Sigma_m^\dagger}^2. \quad (6)$$

In that work, the calibration of the PNM was assessed by randomising the right hand side \mathbf{b} according to $B = \mathbf{A}X$, $X \sim \mu_0$. The PNM was said to be “well-calibrated” if the statistic $Z_m(X)$ is distributed as χ_{d-r}^2 , where $r = \text{rank}(\mathbf{\Sigma}_m)$. Note that when X is randomised in this way both the mean \mathbf{x}_m and covariance $\mathbf{\Sigma}_m$ may² also be random, as a consequence of the dependence of \mathcal{I} on \mathbf{b} . Loosely speaking, this definition implies that the “width” of the Gaussian $\mathcal{N}(\mathbf{x}_m, \mathbf{\Sigma}_m)$ is commensurate with the $\|\mathbf{x}_* - \mathbf{x}_m\|_2$ when averaged over realisations of X . Empirical studies in Cockayne et al. [2019a] illustrate that the PNM proposed in that work (called *BayesCG*) fails to be well-calibrated, though we note that Reid et al. [2020] proposes a particular prior under which BayesCG is well-calibrated.

The recently proposed notion of calibration in Graham et al. [2020] generalises and formalises the criterion introduced above. That work defines both *weak* and *strong* calibration for learning procedures such as PNMs. In this section we will show that linear, stationary, first-degree probabilistic iterative methods are *strongly* calibrated, when a Gaussian μ_0 is used. Initially we assume that $\mathbf{\Sigma}_m$ is nonsingular, which implies that \mathbf{G} must also be nonsingular.

Definition 3.4 (Strongly Calibrated, Nonsingular Case). Let $\mu_0 = \mathcal{N}(\mathbf{x}_0, \mathbf{\Sigma}_0)$, and suppose that a probabilistic iterative method $\mathcal{I}_\#$ produces iterates of the form $\mu_m = \mathcal{N}(\mathbf{x}_m, \mathbf{\Sigma}_m)$. Suppose that $\mathbf{\Sigma}_0$ and $\mathbf{\Sigma}_m$ are each symmetric positive-definite. Then $\mathcal{I}_\#$ is said to be *strongly calibrated* for (μ_0, \mathbf{A}) if, when $\mathcal{I}_\#$ is applied to solve a random linear system defined by \mathbf{A} and $B = \mathbf{A}X$, $X \sim \mu_0$, it holds for all $m > 0$ that

$$\mathbf{\Sigma}_m^{-\frac{1}{2}}(X - \mathbf{x}_m) \sim \mathcal{N}(\mathbf{0}, \mathbf{I}).$$

A similar notion of strong calibration has been exploited for verifying the correctness of algorithms for Bayesian computation in Cook et al. [2006], Talts et al. [2018]; see Graham et al. [2020] for full details. Note that the definitions in Graham et al. [2020] do not require that μ_m be the output of a linear stationary probabilistic iterative method; any PNM for linear systems may be assessed similarly to using Definition 3.4 provided it yields a Gaussian output. This means that existing probabilistic linear solvers, such as Hennig [2015] and Cockayne et al. [2019a], may also be assessed against this definition.

The criterion presented in Definition 3.4 is stronger than that from Eq. (6) due to the absence of the norm; clearly if a method is strongly calibrated then $Z_m(X)$ will follow the appropriate χ_d^2 distribution, though the converse does not necessarily hold. Nevertheless the definition may be motivated from the same intuition as used in Cockayne et al. [2019a] to justify Eq. (6): the probabilistic iterative method is strongly calibrated if the true solution looks like a sample from μ_m , on average with respect to randomness in X .

The next proposition proves that when \mathbf{G} is nonsingular, probabilistic iterative methods are strongly calibrated.

²Note that if \mathcal{I} is a linear, stationary, first-degree iterative method then $\mathbf{\Sigma}_m$ depends only on \mathbf{G} , and for such methods \mathbf{G} is often independent of \mathbf{b} . In this case $\mathbf{\Sigma}_m$ is not random when X is randomised, though in general this should not be expected to be the case.

Proposition 3.5. *Let the assumptions of Definition 3.4 hold. Additionally assume that \mathcal{I} is a linear first degree stationary iterative method with nonsingular \mathbf{G} , and such that Eq. (5) holds with probability one when \mathcal{I} is applied to solve a system defined by the right hand side $B = \mathbf{A}X$, $X \sim \mu_0$. Then $\mathcal{I}_\#$ is strongly calibrated for (μ_0, \mathbf{A}) .*

Proof. We begin by examining $\Sigma_m^{-\frac{1}{2}}(\mathbf{x}_* - \mathbf{x}_m)$ for a fixed true solution, then we complete the proof by randomising \mathbf{x}_* to obtain the result. First note that Σ_m is nonsingular since \mathbf{G} and Σ_0 are nonsingular. Now, for each fixed \mathbf{x}_* and all $m > 0$ we have:

$$\begin{aligned}\Sigma_m^{-\frac{1}{2}}(\mathbf{x}_* - \mathbf{x}_m) &= \Sigma_{m-1}^{-\frac{1}{2}}(\mathbf{G}^{-1}(\mathbf{x}_* - \mathbf{G}\mathbf{x}_{m-1} - \mathbf{f})) \\ &= \Sigma_{m-1}^{-\frac{1}{2}}(\mathbf{G}^{-1}(\mathbf{x}_* - \mathbf{f}) - \mathbf{x}_{m-1}).\end{aligned}$$

Now we have $\mathbf{G}^{-1}(\mathbf{x}_* - \mathbf{f}) = \mathbf{x}_*$, from nonsingularity of \mathbf{G} and Eq. (5). It follows inductively over m that

$$\begin{aligned}\Sigma_m^{-\frac{1}{2}}(\mathbf{x}_* - \mathbf{x}_m) &= \Sigma_{m-1}^{-\frac{1}{2}}(\mathbf{x}_* - \mathbf{x}_{m-1}) \\ &= \Sigma_0^{-\frac{1}{2}}(\mathbf{x}_* - \mathbf{x}_0).\end{aligned}$$

Thus, if we now randomise \mathbf{x}_* according to $X \sim \mu_0 = \mathcal{N}(\mathbf{x}_0, \Sigma_0)$, we obtain $\Sigma_m^{-\frac{1}{2}}(X - \mathbf{x}_m) \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$, completing the proof. \square

Remark 3.6. Note that the only demand Proposition 3.5 makes of \mathcal{I} is that Eq. (5) holds, loosely speaking, for all the linear systems on which μ_0 places probability mass. This is distinct from the *convergence* of the method, for which we require that $\|\mathbf{G}\| < 1$. Thus calibratedness does not imply a requirement that μ_m contracts to the truth, only that the width of μ_m should reflect the error. If \mathcal{I} diverges for some \mathbf{x}_0 it is natural that μ_m should tend to a distribution with infinite variance as $m \rightarrow \infty$.

The assumption of nonsingular \mathbf{G} permits a straightforward proof for Proposition 3.5, but unfortunately \mathbf{G} may be singular even for such elementary methods as the Jacobi iterations. The next definition adapts Definition 3.4 to the case where \mathbf{G} , and therefore also Σ_m , are singular. It simplifies the subsequent presentation to focus on the case where Σ_m does not depend on \mathbf{b} . To the best of our knowledge this is the case for the majority of stationary iterative methods.

Definition 3.7 (Strongly Calibrated, Singular Case). Let $\mu_0 = \mathcal{N}(\mathbf{x}_0, \Sigma_0)$, and suppose the probabilistic iterative method $\mathcal{I}_\#$ produces iterates of the form $\mu_m = \mathcal{N}(\mathbf{x}_m, \Sigma_m)$, where Σ_m does not depend on the right hand side \mathbf{b} . Suppose that Σ_0 is symmetric positive-definite and Σ_m is positive semidefinite with rank $0 < r < d$. Let $\mathbf{R} \in \mathbb{R}^{d \times r}$, $\mathbf{N} \in \mathbb{R}^{d \times (d-r)}$ be matrices such that $\text{range}(\mathbf{R}) = \text{range}(\Sigma_m)$ and $\text{range}(\mathbf{N}) = \text{ker}(\Sigma_m)$. Then μ_m is said to be *strongly calibrated* for (μ_0, \mathbf{A}) if, when $\mathcal{I}_\#$ is applied to solve a random linear system defined by \mathbf{A} and $B = \mathbf{A}X$, $X \sim \mu_0$, the following two conditions are satisfied:

1. $(\mathbf{R}^\top \Sigma_m \mathbf{R})^{-\frac{1}{2}} \mathbf{R}^\top (X - \mathbf{x}_m) \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_r)$.
2. $\mathbf{N}^\top (X - \mathbf{x}_m) = \mathbf{0}$.

Remark 3.8. Both Definitions 3.4 and 3.7 provide *sufficient* conditions for a method to be strongly calibrated. However dependence of \mathbf{G} on \mathbf{b} does not preclude strong calibration. Thus the work in Reid et al. [2020] may yield a strongly calibrated PNM, but a more general definition than Definition 3.7 would be required to verify this.

Remark 3.9. Since Σ_m does not depend on \mathbf{b} , both \mathbf{R} and \mathbf{N} are fixed matrices independent of X . Furthermore while \mathbf{R} and \mathbf{G} must be bases of the range and kernel of Σ_m respectively, neither Definition 3.7 nor Proposition 3.10 depend on any particular choice of basis.

The above definition is an intuitive extension of Definition 3.4 to the case of singular Σ_m ; it demands that in any subspace of \mathbb{R}^d in which Σ_m is nonzero, $\mathcal{I}_\#$ is strongly calibrated as in Definition 3.4, and in any subspace in which it is zero and thus no uncertainty remains, \mathbf{x}_m is identically equal to the true solution X .

We then have the following result, the proof of which is provided in Appendix Appendix A. The intuition behind the proof in the singular case is the same as in the nonsingular case, but additional technical effort is required to project into the null space of Σ_m .

Proposition 3.10. *Let \mathcal{I} be a linear first degree stationary iterative method such that Eq. (5) holds with probability one when \mathcal{I} is applied to solve a system defined by the right hand side $\mathbf{b} = \mathbf{A}X$, $X \sim \mu_0$. Suppose that \mathbf{G} is independent of \mathbf{b} , and that \mathbf{G} is diagonalisable with rank $0 < r \leq d$. Then $\mathcal{I}_\#$ is strongly calibrated for (μ_0, \mathbf{A}) .*

Proof. See Appendix Appendix A. □

Proposition 3.10, together with Proposition 3.5, provides a clear and defensible sense in which the output μ_m from a probabilistic iterative method $\mathcal{I}_\#$, arising from a linear first degree stationary iterative method \mathcal{I} , can be considered to be meaningful. Specifically, one has a guarantee that the unknown solution \mathbf{x}_* is indistinguishable, in a statistical sense, from samples drawn from μ_m . Thus one may interpret μ_m as quantifying uncertainty with respect to the unknown true value of \mathbf{x}_* .

3.3. Generalisations

Here we discuss generalisations to both non-stationary and higher degree iterative methods, while remaining in the linear framework.

3.3.1. Non-Stationary Methods

In a non-stationary linear iterative method of first degree [Young, 1971, Chapter 9], the iteration is of the form:

$$\mathbf{x}_m = \mathbf{G}_m \mathbf{x}_{m-1} + \mathbf{f}_m \tag{7}$$

where $\mathbf{f}_m \in \mathbb{R}^d$ and $\mathbf{G}_m \in \mathbb{R}^{d \times d}$ for all $m \geq 0$. The map P^m is then of the form:

$$P^m(\mathbf{x}_0) = \hat{\mathbf{G}}_m \mathbf{x}_0 + \hat{\mathbf{f}}_m$$

$$\hat{\mathbf{G}}_m = \prod_{i=1}^m \mathbf{G}_i \quad \hat{\mathbf{f}}_m = \mathbf{f}_m + \sum_{i=1}^m \left(\prod_{j=i+1}^m \mathbf{G}_j \right) \mathbf{f}_i.$$

From this it follows by an identical argument to Proposition 3.2 that $\mu_m = \mathcal{N}(\mathbf{x}_m, \boldsymbol{\Sigma}_m)$, with $\mathbf{x}_m = \hat{\mathbf{G}}_m \mathbf{x}_0 + \hat{\mathbf{f}}_m$ and $\boldsymbol{\Sigma}_m = \hat{\mathbf{G}}_m \boldsymbol{\Sigma}_0 \hat{\mathbf{G}}_m^\top$.

Considering the consistency of the implied probabilistic iterative method, as in the stationary setting, \mathbf{x}_m coincides with the classical iterate. Furthermore, Young [1971] notes that the iteration from Eq. (7) converges to \mathbf{x}_* if and only if $\hat{\mathbf{G}}_m \rightarrow \mathbf{0}$. In this event clearly $\boldsymbol{\Sigma}_m \rightarrow \mathbf{0}$, and so provided the underlying iterative converges, μ_m converges to a Dirac mass on \mathbf{x}_* as $m \rightarrow \infty$.

From the perspective of calibration of UQ, the proofs in Section 3.2 do not apply to non-stationary methods \mathcal{I} since those proofs exploit that $\boldsymbol{\Sigma}_m = \mathbf{G}^m \boldsymbol{\Sigma}_0 (\mathbf{G}^m)^\top$, which no longer holds in the non-stationary setting. However if one instead directly assumes $\hat{\mathbf{G}}_m$ to be diagonalisable for each m , the proof of Proposition 3.5 would need only minor modifications to establish strong calibration of $\mathcal{I}_\#$ in the non-stationary setting.

3.3.2. Higher Degree Methods

Modifying Definition 2.4 to allow methods of higher degree requires changing the space on which μ is defined, and the domain of P_m (and by extension $(P_m)_\#$), to a cartesian product of s instances of \mathbb{R}^d .

In terms of such methods, when $s = 2$ [Young, 1971, Chapter 16] the iteration takes the form

$$\mathbf{x}_m = \mathbf{G}\mathbf{x}_{m-1} + \mathbf{H}\mathbf{x}_{m-2} + \mathbf{k} \quad (8)$$

where $\mathbf{G}, \mathbf{H} \in \mathbb{R}^{d \times d}$ and $\mathbf{k} \in \mathbb{R}^d$. While second degree methods are seldom used in practice, higher order methods can accelerate convergence and raise some interesting statistical questions. These methods are analysed by augmenting the space as follows, to obtain a first degree linear stationary iterative method on \mathbb{R}^{2d} :

$$\begin{pmatrix} \mathbf{x}_{m-1} \\ \mathbf{x}_m \end{pmatrix} = \begin{pmatrix} \mathbf{0} & \mathbf{I} \\ \mathbf{H} & \mathbf{G} \end{pmatrix} \begin{pmatrix} \mathbf{x}_{m-1} \\ \mathbf{x}_m \end{pmatrix} + \begin{pmatrix} \mathbf{0} \\ \mathbf{k} \end{pmatrix} = \tilde{\mathbf{G}} \begin{pmatrix} \mathbf{x}_{m-1} \\ \mathbf{x}_m \end{pmatrix} + \tilde{\mathbf{k}}.$$

Convergence of the iterate, and hence the covariance in Proposition 3.2, then require $\rho(\tilde{\mathbf{G}}) < 1$. Similarly, provided $\tilde{\mathbf{G}}$ satisfies the assumptions of one of the theorems in Section 3.2, μ_m will provide useful uncertainty quantification according to Definitions 3.4 and 3.7.

An interesting technicality for higher degree methods is that, whereas in first degree methods only an initial iterate \mathbf{x}_0 must be supplied, in second degree methods both the iterates \mathbf{x}_0 and \mathbf{x}_1 are required. This raises a challenge in the probabilistic framework because it is not clear how one should specify an initial distribution jointly over \mathbf{x}_0 and \mathbf{x}_1 . While expert knowledge may be exploited to build a distribution over \mathbf{x}_0 , the same is not true of \mathbf{x}_1 . Several possible approaches are considered experimentally in Section 5.

4. Beyond Linearity

In the non-Gaussian and non-linear setting it is significantly more difficult to prove that output from a probabilistic iterative method is strongly calibrated. Instead, in this section we adopt an alternative notion from Graham et al. [2020] called *weak* calibratedness, which is more straightforward to empirically test. In Section 4.1 we present that definition and in Section 4.2 discuss statistical tests for calibration which will be applied in Section 5 to test for calibration when nonlinear iterative methods are used.

4.1. Weakly Calibrated Probabilistic Iterative Methods

The chief issue with Definitions 3.4 and 3.7 is that in order to prove calibration we require that both μ_0 and μ_m are Gaussian. This is problematic because Gaussian distributions are unable to express all initial beliefs about components of \mathbf{x}_* , and because the linear iterative methods which result in a Gaussian μ_m are less widely-used compared to non-linear iterative methods, such as CG. The definition of strong calibration is difficult to verify in such general cases, and so we turn to an alternative, weaker sense in which the output μ_m from a (possibly nonlinear) probabilistic iterative method can be considered to be meaningful.

To define what it means for a probabilistic iterative method $\mathcal{I}_\#$ to be weakly calibrated we proceed as follows: In the same setting as Section 3.2, we fix \mathbf{A} and randomly generate a right hand side $\mathbf{b} = \mathbf{A}X$, $X \sim \mu_0$. Then, conditional on X and for each $m > 0$, we introduce a second random variable $Y^{(m)}|X \sim \mu_m$ that is sampled from the output μ_m from the probabilistic iterative method applied to solve the linear system defined by \mathbf{A} and \mathbf{b} . Let $Y^{(m)}$ denote the random variable obtained by marginalising $Y^{(m)}|X$ over realisations of X . Note that realisations of $Y^{(m)}$ may be accessed trivially by sampling \mathbf{x}_0 and \mathbf{x}_* independently from μ_0 and applying the iterative method $\mathcal{I}(\mathbf{A}, \mathbf{A}\mathbf{x}_*)$, initialised at \mathbf{x}_0 .

Definition 4.1 (Weakly Calibrated). A probabilistic iterative method $\mathcal{I}_\#$ is said to be *weakly calibrated* to (μ_0, \mathbf{A}) if, when $\mathcal{I}_\#$ is applied to solve a random linear system defined by \mathbf{A} and $\mathbf{b} = \mathbf{A}X$, $X \sim \mu_0$, and when $Y^{(m)}|X \sim \mu_m$, it holds for all $m > 0$ that

$$Y^{(m)} \sim \mu_0. \tag{9}$$

Thus the sense in which Definition 4.1 is weaker than Definitions 3.4 and 3.7 is that from a simulation perspective, in *strong* calibration we:

1. Draw $\mathbf{x}_* \sim \mu_0$,
2. Compute μ_m according to $\mathcal{I}_\#(\mathbf{A}, \mathbf{A}\mathbf{x}_*)$ and draw $\mathbf{x} \sim \mu_m$,
3. Compare \mathbf{x}_* to \mathbf{x} .

Conversely in *weak* calibration we:

1. Draw $\mathbf{x}_* \sim \mu_0$ and $\mathbf{x}'_* \sim \mu_0$ *independently*,

2. Compute μ_m according to $\mathcal{I}_{\#}(\mathbf{A}, \mathbf{A}\mathbf{x}_*)$ and draw $\mathbf{x} \sim \mu_m$,
3. Compare \mathbf{x}'_* to \mathbf{x} .

In the former \mathbf{x} is dependent on the sample from μ_m it is compared to, whereas in the latter only a marginal comparison is performed.

Eq. (9) is sometimes called the *self-consistency property* and, as with strong calibration, the notion of weak calibration has previously been exploited to verify the correctness of algorithms for Bayesian computation [Geweke, 2004]. Graham et al. [2020, Lemma 2.19] establishes that strong calibration implies weak calibration. Although a weaker property, when the initial and output distributions are non-Gaussian Definition 4.1 allows for statistical tests of distributional equality to be used to assess the quality of the uncertainty quantification provided.

4.2. Testing for Weak Calibration

We now present a statistical test to determine whether a probabilistic iterative method is weakly calibrated or not. For convenience we let ν_m denote the distribution of $Y^{(m)}$, so that we aim to test whether $\nu_m = \mu_0$. Since ν_m does not necessarily have a closed form and, in particular, it may only be possible to access samples $Y_1^{(m)}, \dots, Y_N^{(m)}$ from ν_m , we aim to perform a goodness-of-fit test to determine whether these samples are consistent with being drawn from μ_0 or not. To achieve this we adopt a general purpose goodness-of-fit test based on *maximum mean discrepancy* (MMD), due to Gretton et al. [2012], which we briefly describe next.

Definition 4.2 (Maximum Mean Discrepancy). Let $\mu, \nu \in \mathcal{P}(\mathbb{R}^d)$ and let \mathcal{F} be a set of real-valued, μ and ν -integrable functions on \mathbb{R}^d . Then the MMD between μ and ν , based on \mathcal{F} , is given by

$$\text{MMD}_{\mathcal{F}}(\mu, \nu) := \sup_{f \in \mathcal{F}} \left(\int f(\mathbf{x}) \mu(d\mathbf{x}) - \int f(\mathbf{x}) \nu(d\mathbf{x}) \right).$$

Gretton et al. [2012] considered taking \mathcal{F} to be a unit ball in a reproducing kernel Hilbert space (RKHS), showing that when the RKHS is chosen judiciously then MMD is a metric on $\mathcal{P}(\mathbb{R}^d)$. Moreover, this choice ensures that an unbiased estimator for MMD can be constructed, as will now be explained. Recall that an RKHS is associated with a symmetric positive definite *kernel* $k : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$; we emphasise this using the notation $\mathcal{F} \equiv \mathcal{F}_k = \{f \in \mathcal{H}_k : \|f\|_{\mathcal{H}_k} \leq 1\}$ where \mathcal{H}_k is the unique RKHS with kernel k . Define the *kernel mean embedding* of μ in \mathcal{H}_k as $\mu[k]$ where $\mu[k](\mathbf{x}) := \int k(\mathbf{x}, \mathbf{x}') \mu(d\mathbf{x}')$. Then Gretton et al. [2012, Lemma 4] asserts that $\text{MMD}_{\mathcal{F}_k}(\mu, \nu)$ can be expressed as a difference between the kernel mean embeddings of μ and ν :

$$\text{MMD}_{\mathcal{F}_k}(\mu, \nu) := \|\mu[k] - \nu[k]\|_{\mathcal{H}_k}. \quad (10)$$

For convenient choices of k and μ it may be possible to compute $\mu[k]$ in closed-form, but in general one must resort to approximating Eq. (10) based on samples from one or

both of μ and ν . Given independent samples $X_1, \dots, X_N \sim \mu_0$ and $Y_1^{(m)}, \dots, Y_N^{(m)} \sim \nu_m$, we define an estimator

$$\widehat{\text{MMD}}_{\mathcal{F}_k}^2 := \frac{1}{N(N-1)} \sum_{\substack{i,j=1 \\ i \neq j}}^N k(X_i, X_j) + k(Y_i^{(m)}, Y_j^{(m)}) - k(X_i, Y_i^{(m)}) - k(Y_i^{(m)}, X_j), \quad (11)$$

which can be verified to be an unbiased estimator of $\text{MMD}_{\mathcal{F}_k}(\mu, \nu_m)^2$ provided that, in addition to having the stated distribution, the samples $Y_1^{(m)}, \dots, Y_N^{(m)}$ are generated independently from the samples X_1, \dots, X_N .

The statistic in Eq. (11) enables a goodness-of-fit test to be performed, and the distribution of this test statistic under the null hypothesis $\nu_m = \mu_0$ may be estimated using a standard bootstrap procedure as described in Gretton et al. [2012, Section 5]. Having obtained M approximate samples from the distribution of Eq. (11) using the bootstrap, we determine a threshold for a prescribed power level $\alpha \in (0, 1)$ by computing a $(1 - \alpha)$ -quantile of this empirical distribution.

5. Empirical Assessment

The aim of this section is to gain insight into the uncertainty being quantified by a probabilistic iterative method. We accomplish this by analysing the performance of these methods when applied to the problem of inverting a linear system that arises when building a kernel interpolant. Our aim is not to address the problem of computing kernel interpolants *per se*, as many powerful methods exist for this task, but this problem serves as a convenient test-bed in which probabilistic iterative methods can be examined.

5.1. Problem Definition

Consider a dataset consisting of pairs (z_i, y_i) , $i = 1, \dots, d$, $d \in \mathbb{N}$, where the $z_i \in [0, 1]$ are distinct locations at which observations $y_i \in \mathbb{R}$ of some physical phenomenon were obtained. The aim is to compute an interpolant of this dataset, that is, a function $g : [0, 1] \rightarrow \mathbb{R}$ which is such that $g(z_i) = y_i$ for all $i = 1, \dots, d$. For a given symmetric positive definite kernel $c : [0, 1] \times [0, 1] \rightarrow \mathbb{R}$, we consider an interpolant of the form

$$g(z) := \sum_{i=1}^d x_i c(z, z_i) \quad (12)$$

and note that there is a unique set of weights $x_i \in \mathbb{R}$ such that the interpolation equations

$$g(z_i) = y_i, \quad i = 1, \dots, d$$

are satisfied. The vector $\mathbf{x}_* = [x_1, \dots, x_d]^\top$ of such weights satisfies the d -dimensional linear system in Eq. (1) with $A_{i,j} = c(z_i, z_j)$ and $\mathbf{b} = (y_1, \dots, y_d)^\top$.

This linear system is representative of linear systems that are widely encountered in statistics and machine learning, and naturally a variety of methods have been proposed to

circumvent the need to solve them; for example, based on reducing the degrees of freedom of the parametric function g so that the dataset is only approximately interpolated. Our aim is to use a finite number of iterations, m , of a probabilistic iterative method on the full problem in Eq. (1) and to lift the distribution μ_m over the unknown solution vector \mathbf{x}_* into the function space spanned by functions of the form in Eq. (12). This enables uncertainty due to limited computation to be interpreted in the domain on which the interpolation problem was defined.

The condition number of \mathbf{A} depends on the spectrum of the kernel k and the closeness of the elements in $\{z_1, \dots, z_d\}$. For kernels with rapidly decaying spectrum, such as the squared exponential kernel

$$c(x, y) = \exp\left(-\frac{\|x - y\|_2^2}{2\ell^2}\right) \quad (13)$$

with length-scale parameter $\ell > 0$, it is common for \mathbf{A} to be badly conditioned. Thus even when d is small, direct solution of Eq. (1) can be difficult and careful numerical analysis is required.

A dataset of size $d = 440$ was generated, with $(z_i)_{i=1, \dots, d}$ consisting of 20 evenly spaced points in $[0, 0.1]$, 400 evenly spaced points in $[0.2, 0.8]$ and 20 evenly spaced points in $[0.9, 1]$, and $y_i = f(z_i)$ where $f(z) = 1_{z < 0.5} \sin(2\pi z) + 1_{z \geq 0.5} \sin(4\pi z)$. The parameter $\ell = 0.0012$ was used, which produces a system for which a direct solver can be used, so that a ground-truth is accessible, but which is not entirely trivial.

5.2. Choice of μ_0

For the initial distribution μ_0 several candidates were considered. Firstly a `DEFAULT` choice given by $\mu_0 = \mathcal{N}(\mathbf{0}, \mathbf{I})$ which can be interpreted as a lack of *a priori* insight. Secondly the `NATURAL` choice $\mu_0 = \mathcal{N}(\mathbf{0}, \mathbf{A}^{-1})$ which incorporates the structure of \mathbf{A} into the initial distribution, and has been noted in the past to produce posterior means that replicate CG in works such as Cockayne et al. [2019a], Hennig [2015]. We note that it is not a practical choice in general as it requires computation of \mathbf{A}^{-1} .

The third initial distribution we consider is based on an attempt to adapt to the scale of the problem by maximising the probability of observing a small number of *ansatz* solutions to the linear system, perhaps obtained by expert knowledge of the system at hand. Let $\mathbf{x}_{*,i}$, $i = 1, \dots, N$ be these solutions, and suppose that we wish to use these to estimate the scaling parameter ν^2 for the prior $\mu_0 = \mathcal{N}(\mathbf{0}, \nu^2 \Sigma_0)$ where Σ_0 . Standard results yield that the estimate

$$\nu_{\text{opt}}^2 := \frac{1}{Nd} \sum_{i=1}^N \|\mathbf{x}_{*,i}\|_{\Sigma_0^{-1}}^2.$$

In the experiments below we assume that $\Sigma_0 = \mathbf{I}$ and term this initial distribution `OPT`. We used $N = 5$ *ansatz* solutions, and these were obtained by sampling 5 right-hand-sides $\mathbf{b}_1, \dots, \mathbf{b}_5 \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and computing $\mathbf{x}_{*,i} = \mathbf{A}^{-1} \mathbf{b}_i$. We note that this does not result in an entirely fair comparison since 5 exact solutions to the linear system are

used to construct the prior. However since the focus of this paper is not primarily on the problem of selecting μ_0 , we have opted for this simple approach. Nevertheless, owing to the importance of this choice the OPT initial distribution is included to highlight that objective selection procedures may be considered, and we expect these to be a subject of future research.

5.3. Results in Function Space

In this section we examine the resulting distributions μ_m from application of a number of probabilistic iterative methods to the problem above, for each choice of initial distribution from Section 5.2.

Stationary Iterative Methods We first consider Richardson’s iteration with a constant step size. Since this method is stationary and linear, the theoretical results obtained in Section 3 apply. The step size ω was set to either the optimal value, $\omega = 2/(\lambda_{\min}(\mathbf{A}) + \lambda_{\max}(\mathbf{A}))$, that minimises the spectral radius of \mathbf{G} , or a default value $\omega = 2/3$. Jacobi’s method was also considered, but in our simulations the results were virtually identical for this problem, so they are not presented.

Fig. 1 displays samples (grey curves) from each of the probabilistic iterative methods that we considered and the blue curve represents the exact kernel interpolant. For each probabilistic iterative method considered, the output was seen to contract around the exact solution as the number m of iterations is increased. Interestingly, very little variation is observed in the intervals $[0, 0.1]$ and $[0.9, 1]$, which accords with the fact that the interpolant is being approximated well in these regions and suggests that the distributional output can act as a local error indicator.

Non-Stationary and Higher-Order Methods We now consider non-stationary and higher-order schemes. As discussed in Section 3.3, these schemes are expected to be well-calibrated as they are still linear, though calibration has not been rigorously proven. For the non-stationary scheme we considered Richardson iteration again but with the step-size chosen adaptively, with $\omega_m = \mathbf{r}_m^\top \mathbf{A} \mathbf{r}_m / \|\mathbf{A} \mathbf{r}_m\|_2^2$ minimising the Euclidean norm of the residual $\mathbf{r}_{m+1} = \mathbf{b} - \mathbf{A} \mathbf{x}_{m+1}$. Results for the non-stationary scheme are presented in Fig. 2, with qualitative behaviour appearing to be similar to with the default step size from Fig. 1a. Since the non-stationary scheme is better able to adapt to the problem at hand, this seems a more prudent choice than an arbitrary $\omega = 2/3$, though we note that the calibration of this method remains to be assessed empirically; this will be considered in Section 5.4.

For the higher order scheme we opted for a second-degree version of Richardson iteration presented in Young [1972]. In this method the iteration was of the form

$$\mathbf{x}_m = \gamma\sigma \left(\frac{2}{\beta - \alpha} \mathbf{G} - \frac{\beta + \alpha}{\beta - \alpha} \right) \mathbf{x}_{m-1} + (1 - \gamma) \mathbf{I} \mathbf{x}_{m-2} + \frac{2\gamma\sigma}{\beta - \alpha} \mathbf{f}$$

where \mathbf{G} and \mathbf{f} are as given in the classical first-order Richardson iteration from Example 2.1, with optimal step size $\omega = 2/(\lambda_{\min}(\mathbf{A}) + \lambda_{\max}(\mathbf{A}))$, while $\alpha = \lambda_{\min}(\mathbf{G})$, $\beta =$

$\lambda_{\max}(\mathbf{G})$ and

$$\sigma = \frac{\beta - \alpha}{2 - (\beta - \alpha)} \quad \gamma = \frac{2}{1 + \sqrt{1 - \sigma^2}}.$$

Recall that for a second degree probabilistic iterative method, a joint initial distribution must be specified for \mathbf{x}_0 and \mathbf{x}_1 . The distribution assigned to \mathbf{x}_0 was fixed to the OPT, since this appeared to provide better uncertainty quantification across different choices of ω . Three choices were considered for initial distributions for \mathbf{x}_1 : IID, in which \mathbf{x}_1 is an independent copy of \mathbf{x}_0 , CORR, in which \mathbf{x}_1 is identical to \mathbf{x}_0 and RICH, in which \mathbf{x}_1 is obtained from \mathbf{x}_0 by performing one iteration of Richardson iteration with optimal step size. Note that both IID and CORR yield the same marginal distribution for \mathbf{x}_1 , but the joint distribution differs

Fig. 3 displays samples from the output of the probabilistic iterative methods just described. Qualitatively, the results appear to be similar to those from Fig. 1b with initial distribution OPT, as one would expect given that μ_0 in all three rows is that same distribution. Of the three choices for μ_1 , RICH appears to contract marginally faster, though in all three methods the improvement over the first order method from Fig. 1b appears to be negligible.

Nonlinear methods We now turn to the conjugate gradient method, which is the most practically relevant of the probabilistic iterative methods we consider, but for which our theoretical results on strong calibration do not hold. Results are displayed in Fig. 4. Convergence is clearly seen to be faster than in the other methods considered, though qualitatively the samples obtained otherwise seem to be similar. This hints at the results from the next section, in which we will see that CG is weakly calibrated for this problem and for the initial distributions that we considered.

5.4. Testing Calibration

We now test for evidence against weak calibration for all of the probabilistic iterative methods and initial distributions considered. Recall that, according to the results in Section 3.2, stationary Richardson iterations give rise to probabilistic iterative methods that are strongly calibrated when ω is fixed (irrespective of whether the optimal stepsize or a default choice is used). Non-stationary Richardson iteration with adaptive stepsize ω_m are expected to give rise to probabilistic iterative methods that are strongly calibrated but no proof has been obtained, as is the higher order method described above. It is unknown whether probabilistic iterative methods based on CG are strongly or weakly calibrated. In addition to probabilistic iterative methods, we also include BayesCG from Cockayne et al. [2019a], which is not a probabilistic iterative method in the sense of this paper and is not expected to be strongly calibrated owing to results based on the Z -statistic in Eq. (6) presented in Cockayne et al. [2019a] and in Reid et al. [2020]. It was hitherto unknown whether BayesCG is weakly calibrated.

We apply the test described in Section 4.2. For each initial distribution and each method we generated $N = 100$ independent samples of μ_0 and ν_m from which the test statistic Eq. (11) was computed. Significance was assessed using the bootstrap method

		Rich. (default)	Rich. (optimal)	Rich. (adaptive)	Rich. (2o)	CG	BayesCG
DEFAULT	MMD	1.90e-04	-3.11e-05	9.76e-06	5.36e-05	-2.80e-05	1.14e-03
	q	0.34	0.52	0.45	0.43	0.49	0.03
NATURAL	MMD	-1.72e-04	-2.71e-04	-2.44e-04	-3.20e-04	-2.98e-04	4.18e-03
	q	0.60	0.64	0.64	0.68	0.68	0.00
OPT	MMD	3.59e-05	1.00e-05	4.30e-06	-6.62e-06	3.57e-05	6.57e-03
	q	0.48	0.47	0.48	0.49	0.47	0.00

Table 1.: Results from applying the maximum mean discrepancy (MMD) based test from Section 4.2 to the methods described in Section 5. The abbreviation “Rich.” refers to Richardson iteration. “2o” refers to the second order method. The test does not reject the null that each of the methods assessed is weakly calibrated, with the exception of BayesCG where the null is rejected. Results that are statistically significant at the 5% level, indicating that the method is not weakly calibrated, are highlighted in bold.

with $M = 1000$. The kernel used was the squared exponential kernel from Eq. (13), with the length-scale set using the *median heuristic* as recommended in Gretton et al. [2012]. For each method $m = 10$ iterations were performed. For the second order method we opted to use the RICH initial distribution for \mathbf{x}_1 .

Table 1 shows the statistics obtained for each of these methods by applying the test for weak calibration described in Section 4.2 for each choice of initial distribution from Section 5.2. Reported are the value of Eq. (11) (as MMD in Table 1). Note that while strictly speaking MMD ought to be positive, due to sampling error it may be negative. This was also observed in Gretton et al. [2012]. Also reported is the statistic q which is analogous to a \mathcal{I} -value in a classical statistical test; if q' is the empirical quantile of MMD within the empirical distribution based on M bootstrap samples of Eq. (11), then $q = 1 - q'$. Thus, a small q represents evidence in support of rejection of the null hypothesis. We used the value 0.05, representing a 5% significance level, as a threshold in Table 1; thus, if a value of q below 0.05 was determined this constitutes evidence that the method is *not* weakly calibrated. Note that owing to the fact that q is based on a sample from the bootstrapped distribution, it is possible to obtain $q = 0$; we would expect the true \mathcal{I} -value to be small but positive.

Examining the results, Richardson iteration with both default and optimal step sizes is seen to be weakly calibrated. This provides support for our testing methodology, since from Graham et al. [2020, Lemma 2.19] any strongly calibrated method must be weakly calibrated. Similarly the second order method is weakly calibrated, which is to be expected since the proof of strong calibration for this method would require only a small extension of the proof for first order methods. Richardson iteration with the adaptive step size appears to be weakly calibrated for all initial distributions considered, suggesting that the non-stationarity implied by the adaptive step size does not affect the weak calibration of the method.

Perhaps more surprisingly owing to its high degree of nonlinearity, CG also appears to be weakly calibrated. This hints at the possibility of a more fundamental result regarding the calibration of probabilistic iterative methods in the general setting, though we leave study of this fact to future work.

Concerning BayesCG (which we emphasise again is *not* a probabilistic iterative method in the same sense as the other methods considered), the results show that BayesCG is *not* weakly calibrated for either the NATURAL or OPT initial distributions μ_0 even when the prior distribution, required in BayesCG, is set equal to μ_0 itself. This is to be expected, considering that this method is widely known *not* to produce well-calibrated posteriors apart from in special cases [e.g. Reid et al., 2020]. One other noteworthy point is that for the DEFAULT initial distribution the MMD obtained for BayesCG has a slightly higher value of $q = 0.03$. This is perhaps due to the fact that with such an uninformative prior BayesCG is known to converge quite slowly. Thus, the posterior after 10 iterations may not have deviated far from the prior.

5.5. Spectral Behaviour

Lastly we examine the spectral behaviour of one of the methods above by performing a principal component analysis, to illustrate how the output of a probabilistic iterative method can provide a richer description of error than classical error bounds. In this section we fixed the distribution μ_0 to NATURAL.

Here we consider principal components (leading eigenvectors) of the covariance matrix $\mathbf{A}\Sigma_m\mathbf{A}^\top$, which describes covariance in the domain of the function Eq. (12). The six leading principal components for the probabilistic iterative method based on Richardson iteration with default step size $\omega = 2/3$ are displayed in Figure 5. At each of the values of m considered, the low frequency variation over the interval $[0.2, 0.8]$ is seen to be the dominant principal component (more so as m is increased), which accords with the result of Figure 1a in that the error of NATURAL is mainly manifest in a low-frequency vertical shift between the exact interpolant and the sampled output. At $m = 100$ the first six components account for over 60% of the variability in the distributional output, with the remaining variability dedicated to higher-frequency aspects of the solution.

The detailed nature of these error indicators may be useful to shed light on the aspects of the exact solution \mathbf{x}_* that we are most uncertain about, having run a probabilistic iterative method. This rich description of numerical uncertainty can trivially be propagated through subsequent computation $F(\mathbf{x}_*)$, e.g. by sampling from μ_m and then applying F , in order to probabilistically assess the impact of numerical uncertainty on any subsequent computational output.

6. Conclusion

In this paper we have introduced probabilistic iterative methods, a new class of probabilistic numerical methods for solving linear systems. We have provided theoretical results concerning the convergence and calibration of these methods in the stationary and linear setting, and examined their empirical performance using a synthetic test-bed.

Finally, we alluded to how the output of a probabilistic iterative method could be used to represent *numerical uncertainty* and how such a representation could be propagated through subsequent computational output.

Several interesting avenues for future related work are now highlighted:

6.1. Generalisation to Nonlinear Methods

The generalisation of this work to nonlinear iterative methods, such as the method of conjugate gradients Hestenes and Stiefel [1952] and other Krylov methods is of interest. These methods are more widely used than stationary iterative methods in modern applications, owing both to their faster convergence and that they only require access to the action of \mathbf{A} , rather than needing to interrogate and modify the elements of \mathbf{A} .

The definition that we proposed for probabilistic iterative methods in Definition 2.4, and the sampling algorithm for accessing the output of a probabilistic iterative method described in Section 3, do not require the generating iterative method to be linear. However, with the exception of Proposition 2.5, the results presented in this paper depend strongly on linearity. Nevertheless the experimental results in Section 5.4 indicate that CG, a prototypical nonlinear iterative method, may be weakly calibrated. Thus, we consider theoretical properties of nonlinear probabilistic iterative methods to be an important line of future work.


6.2. Gradient Flow Interpretation

Recent work in the numerical analysis community highlights that iterative methods for linear systems may be interpreted as the discrete-time solution of an underlying dynamical system [Chu, 2008]. Insight may then be gained by studying the original dynamical system. This connects with recent work in the statistics and machine learning communities that has provided gradient flow interpretations of various sampling and variational inference algorithms [e.g. Arbel et al., 2019, Liu et al., 2019]. An interesting avenue for future work would be to consider whether the methods presented in this paper may be interpreted in a similar way, and whether insight can be gained by performing analysis of the original probability flow.


6.3. Wider Applications

In this paper we have focussed on iterative methods for solving linear systems. However, the assumption that \mathcal{I} was an iterative method for solving such systems was not essential to Definition 2.4. Provided an initial distribution μ_0 can be constructed in the domain of $\mathcal{I}_\#$, probabilistic iterative methods could be applied to *any* classical problem for which iterative methods are used, such as solvers for eigenproblems, numerical optimisation problems or even solvers for nonlinear differential equations. Proposition 2.5 also applies to this general case, provided a suitable bound of the form Eq. (3) can be derived in a norm adapted to the problem and, when the iteration is an affine map, we expect that the proof techniques from Section 3.2 could be applied to establish strong calibration.

Acknowledgements JC was supported by Wave 1 of The UKRI Strategic Priorities Fund under the EPSRC Grant EP/T001569/1, particularly the “Digital Twins for Complex Engineering Systems” theme within that grant, and The Alan Turing Institute. The work of ICFI was supported in part by National Science Foundation grants DMS-1760374 and DMS-1745654. CJO was supported by the Lloyd’s Register Foundation programme on data-centric engineering at the Alan Turing Institute, UK. The work of TWR was supported in part by National Science Foundation grant DMS-1745654.



(a) Default Step Size



(b) Optimal Step Size

Figure 1.: Samples from the distributional output of a probabilistic iterative method based on Richardson iteration, visualised in the physical domain in which the interpolant is defined. The rows in each figure represent the three choices of initial distribution described in Section 5.2. In each panel we present 50 samples (grey curves) from the output of the probabilistic iterative method after m iterations have been performed. The interpolant, corresponding to the exact solution of the linear system, is also shown in blue.




Figure 2.: As in Fig. 1, but with step size chosen adaptively.




Figure 3.: A probabilistic iterative method based on a second degree version of Richardson iteration, as described in Section 5. Each row uses OPT as the initial distribution for \mathbf{x}_0 and a different initial distribution for \mathbf{x}_1 , as described in the main text.




Figure 4.: Samples from the distributional output of the probabilistic iterative method implied by using the conjugate gradient method as the underlying iterative method.

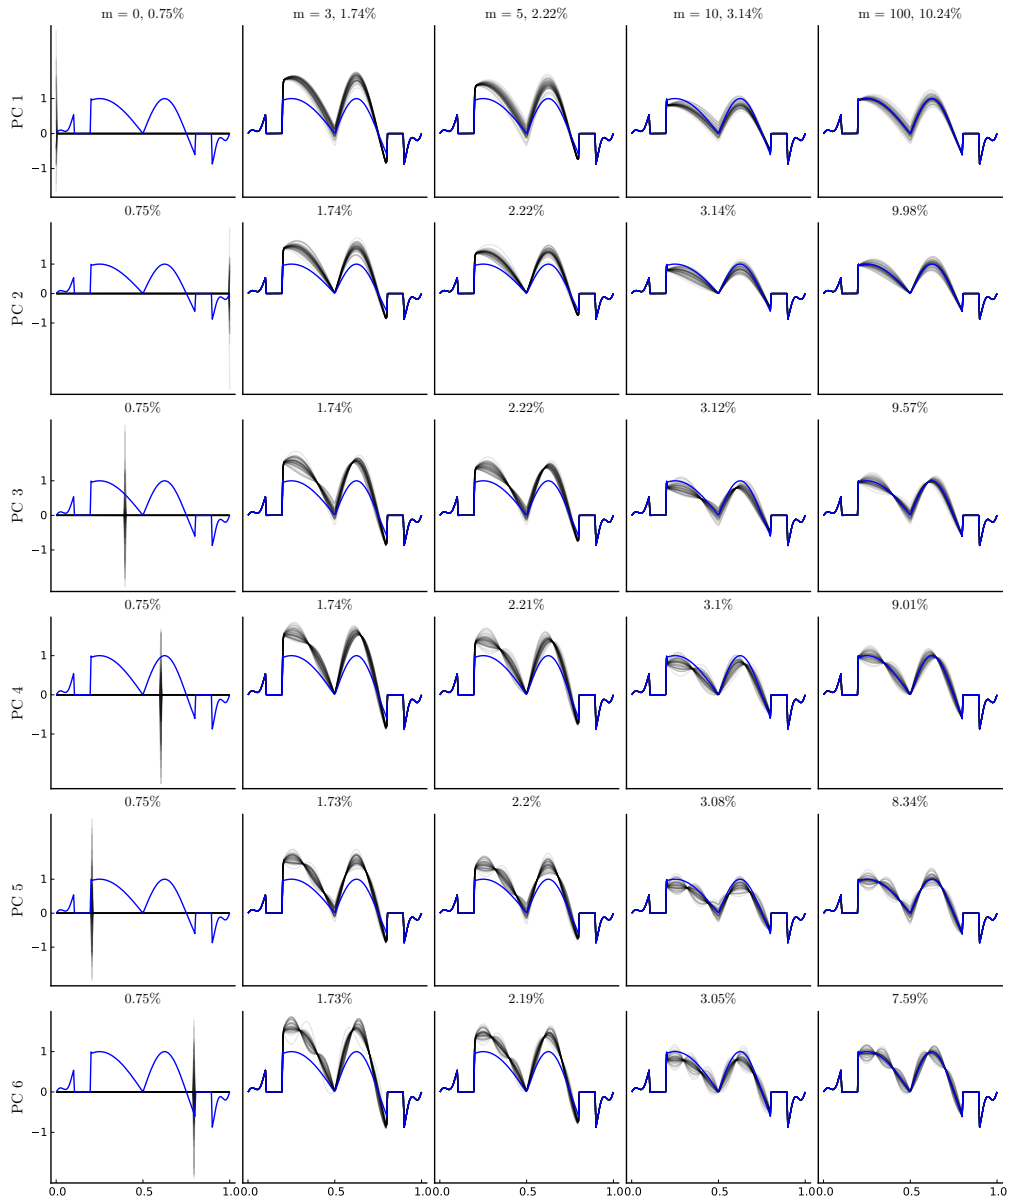


Figure 5.: A closer look at the distributional output: principal components from a probabilistic iterative method based on Richardson iteration with the default step size and initial distribution `NATURAL`. Here the first 6 principal components (PC) are displayed for the same values of m used in Figure 1. The percentages indicate the percentage of the total variation that is explained by that component. Each grey line is constructed as the mean of μ_m , plus a sample in the direction of the relevant principal component, re-scaled to improve visualisation, with 50 samples shown in total.

Appendix A. Proof of Proposition 3.10

In order to prove Proposition 3.10, we need several results from linear algebra about the range and kernel of products of matrices, as well as decomposition of a diagonalizable matrix.

Lemma Appendix A.1 (Ipsen [2009, Fact 6.3]). *Let $\mathbf{Y}, \mathbf{W} \in \mathbb{R}^{d \times d}$. If \mathbf{Y} is non-singular, then $\ker(\mathbf{Y}\mathbf{W}) = \ker(\mathbf{W})$.*

Lemma Appendix A.2 (Ipsen [2009, Facts 6.3 and 6.4]). *Let $\mathbf{Y}, \mathbf{\Omega}, \mathbf{W} \in \mathbb{R}^{d \times d}$ where \mathbf{Y} and \mathbf{W} are non-singular and $\mathbf{\Omega}$ is diagonal. If \mathbf{Y} , $\mathbf{\Omega}$, and \mathbf{W} have the partitions*

$$\mathbf{Y} = [\mathbf{Y}_1 \quad \mathbf{Y}_2], \quad \mathbf{\Omega} = \begin{bmatrix} \mathbf{\Omega}_{11} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix}, \quad \text{and} \quad \mathbf{W} = \begin{bmatrix} \mathbf{W}_1^\top \\ \mathbf{W}_2^\top \end{bmatrix},$$

with $\mathbf{Y}_1, \mathbf{W}_1 \in \mathbb{R}^{d \times r}$, $\mathbf{Y}_2, \mathbf{W}_2 \in \mathbb{R}^{d \times (d-r)}$, and $\mathbf{\Omega} \in \mathbb{R}^{r \times r}$, then

$$\text{range}(\mathbf{Y}\mathbf{\Omega}\mathbf{W}) = \text{range}(\mathbf{Y}_1) \quad \text{and} \quad \ker(\mathbf{Y}\mathbf{\Omega}\mathbf{W}) = \text{range}(\mathbf{W}_2).$$

Lemma Appendix A.3. *Horn and Johnson [2009, Lemma 3.4.1.10] Let $\mathbf{G} \in \mathbb{R}^{d \times d}$ be diagonalisable and of rank $r < d$. Then \mathbf{G} may be represented in its real Jordan canonical form as*

$$\mathbf{G} = \mathbf{Y}\mathbf{\Omega}\mathbf{Y}^{-1}$$

where $\mathbf{Y} \in \mathbb{R}^{d \times d}$ is invertible, while $\mathbf{\Omega} \in \mathbb{R}^{d \times d}$ is of the form

$$\mathbf{\Omega} = \begin{pmatrix} \mathbf{\Omega}_{11} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{pmatrix}.$$

Here $\mathbf{\Omega}_{11} \in \mathbb{R}^{r \times r}$ is nonsingular and block-diagonal, with ℓ 2×2 blocks and s 1×1 blocks, where ℓ is the number of nonzero conjugate pairs of complex eigenvalues of \mathbf{G} and s is the number of nonzero real eigenvalues of \mathbf{G} , so that $r = 2\ell + s$.

With these results stated we proceed to the main proof.

Proof of Proposition 3.10. First note that if $\text{rank}(\mathbf{G}) = d$ then \mathbf{G} is invertible, so the probabilistic iterative method is strongly calibrated as a result of Proposition 3.5. Thus we focus on the case that $\text{rank}(\mathbf{G}) < d$.

We complete this proof in multiple steps:

Step 1 We express the range and kernel of $\mathbf{\Sigma}_m$ in terms of the matrices forming the real Jordan canonical form of \mathbf{G} , thus identifying the matrices \mathbf{R} and \mathbf{N} from Proposition 3.10.

Step 2 We compute $(\mathbf{R}^\top \mathbf{\Sigma}_m \mathbf{R})^{1/2}$, $(\mathbf{R}^\top \mathbf{\Sigma}_m \mathbf{R})^{1/2} \mathbf{R}^\top (\mathbf{x}_m - \mathbf{x}_*)$ and $\mathbf{N}^\top (\mathbf{x}_m - \mathbf{x}_*)$.

Step 3 We combine these results to show that stationary iterative methods are strongly calibrated when \mathbf{G} is diagonalisable.

Step 1 We first compute the range and kernel of Σ_m . The posterior covariance is defined as

$$\Sigma_m = \mathbf{G}^m \Sigma_0 (\mathbf{G}^m)^\top.$$

From Lemma Appendix A.3 we have that

$$\mathbf{G}^i = \mathbf{Y} \Omega^i \mathbf{Y}^{-1}, \quad 0 \leq i \leq m.$$

We partition the diagonalization of \mathbf{G} as

$$\mathbf{Y} = [\mathbf{Y}_1 \quad \mathbf{Y}_2], \quad \Omega = \begin{bmatrix} \Omega_{11} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} \end{bmatrix}, \quad \text{and} \quad \mathbf{Y}^{-1} = \begin{bmatrix} \mathbf{W}_1^\top \\ \mathbf{W}_2^\top \end{bmatrix},$$

where $\mathbf{Y}_1, \mathbf{W}_1 \in \mathbb{R}^{d \times r}$, $\mathbf{Y}_2, \mathbf{W}_2 \in \mathbb{R}^{d \times (d-r)}$, and $\Omega_{11} \in \mathbb{R}^{r \times r}$. With this partitioning and Lemma Appendix A.2 we have

$$\text{range}(\mathbf{G}^i) = \text{range}(\mathbf{Y}_1) \quad \text{and} \quad \ker((\mathbf{G}^i)^\top) = \text{range}(\mathbf{W}_2), \quad 0 \leq i \leq m. \quad (14)$$

We now express the range and kernel of Σ_m in terms of \mathbf{Y}_1 and \mathbf{W}_2 . Express Σ_m as the product $\Sigma_m = \mathbf{Q} \mathbf{Q}^\top$, where $\mathbf{Q} = \mathbf{G}^m \Sigma_0^{1/2}$. For any $\mathbf{x} \in \ker(\Sigma_m)$ we have

$$\Sigma_m \mathbf{x} = \mathbf{0} \iff \mathbf{x}^\top \Sigma_m \mathbf{x} = 0 \iff (\mathbf{Q}^\top \mathbf{x})^\top \mathbf{Q}^\top \mathbf{x} = 0 \iff \mathbf{Q}^\top \mathbf{x} = \mathbf{0}.$$

Thus $\ker(\Sigma_m) = \ker(\mathbf{Q}^\top)$. Because $\Sigma_0^{1/2}$ is the non-singular square root of the non-singular matrix Σ_0 , we can apply Lemma Appendix A.1 to $\mathbf{Q}^\top = \Sigma_0^{1/2} (\mathbf{G}^m)^\top$ to obtain

$$\ker(\Sigma_m) = \ker(\underbrace{\Sigma_0^{1/2} (\mathbf{G}^m)^\top}_{\mathbf{Q}^\top}) = \ker((\mathbf{G}^m)^\top). \quad (15)$$

By the fundamental theorem of linear algebra, $\ker((\mathbf{G}^m)^\top)$ is the orthogonal complement of $\text{range}(\mathbf{G}^m)$ and $\ker(\Sigma_m)$ is the orthogonal complement of $\text{range}(\Sigma_m) = \text{range}(\Sigma_m)$. This combined with Eq. (15) implies

$$\text{range}(\Sigma_m) = \text{range}(\mathbf{G}^m). \quad (16)$$

Applying Lemma Appendix A.2 with $\mathbf{W} = \mathbf{Y}^{-1}$ gives

$$\text{range}(\Sigma_m) = \text{range}(\mathbf{Y}_1) \quad \text{and} \quad \ker(\Sigma_m) = \text{range}(\mathbf{W}_2). \quad (17)$$

Therefore, referring to Proposition 3.10, we have that $\mathbf{R} = \mathbf{Y}_1$ and $\mathbf{N} = \mathbf{W}_2$.

Step 2 We begin by computing $(\mathbf{Y}_1^\top \Sigma_m \mathbf{Y}_1)^{1/2}$. We have that

$$\begin{aligned} \mathbf{Y}_1^\top \Sigma_m \mathbf{Y}_1 &= \mathbf{Y}_1^\top \mathbf{G}^m \Sigma_0 (\mathbf{G}^m)^\top \mathbf{Y}_1 \\ &= \mathbf{Y}_1^\top \mathbf{Y} \Omega^m \mathbf{Y}^{-1} \Sigma_0 \mathbf{Y}^{-\top} (\Omega^m)^\top \mathbf{Y}^\top \mathbf{Y}_1 \\ &= \mathbf{Y}_1^\top \mathbf{Y}_1 \Omega_{11}^m \mathbf{W}_1^\top \Sigma_0 \mathbf{W}_1 (\Omega_{11}^m)^\top \mathbf{Y}_1^\top \mathbf{Y}_1. \end{aligned}$$

The product $\mathbf{Y}_1^\top \mathbf{Y}_1$ is Hermitian positive definite because \mathbf{Y}_1 is full rank. Additionally, $\mathbf{Y}_1 \mathbf{W}_1^\top = \mathbf{I}_r$ because $\mathbf{Y} \mathbf{Y}^{-1} = \mathbf{I}_n$. Therefore the inverse square root³ is,

$$(\mathbf{Y}_1^\top \boldsymbol{\Sigma}_m \mathbf{Y}_1)^{-1/2} = \mathbf{B} \boldsymbol{\Omega}_{11}^{-m} (\mathbf{Y}_1^\top \mathbf{Y}_1)^{-1}, \quad (18)$$

where $\mathbf{B} = [\mathbf{W}_1^\top \boldsymbol{\Sigma}_0 \mathbf{W}_1]^{-1/2} \in \mathbb{R}^{r \times r}$.

Next, we compute $(\mathbf{Y}_1^\top \boldsymbol{\Sigma}_m \mathbf{Y}_1)^{1/2} \mathbf{Y}_1^\top (\mathbf{x}_* - \mathbf{x}_m)$. Left-multiplying $\mathbf{x}_* - \mathbf{x}_m$ by \mathbf{Y}_1^\top yields

$$\mathbf{Y}_1^\top (\mathbf{x}_* - \mathbf{x}_m) = \mathbf{Y}_1^\top \left(\mathbf{x}_* - \mathbf{G}^m \mathbf{x}_0 - \sum_{i=0}^{m-1} \mathbf{G}^i \mathbf{f} \right) \quad (19)$$

$$= \mathbf{Y}_1^\top \left(\mathbf{x}_* - \mathbf{Y} \boldsymbol{\Omega}^m \mathbf{Y}^{-1} \mathbf{x}_0 - \mathbf{f} - \sum_{i=1}^{m-1} \mathbf{Y} \boldsymbol{\Omega}^i \mathbf{Y}^{-1} \mathbf{f} \right)$$

$$= \mathbf{Y}_1^\top \mathbf{x}_* - \mathbf{Y}_1^\top \mathbf{Y}_1 \boldsymbol{\Omega}_{11}^m \mathbf{W}_1^\top \mathbf{x}_0 - \mathbf{Y}_1^\top \mathbf{f} - \sum_{i=1}^{m-1} \mathbf{Y}_1^\top \mathbf{Y}_1 \boldsymbol{\Omega}_{11}^i \mathbf{W}_1^\top \mathbf{f}. \quad (20)$$

Now left-multiplying by Eq. (18) gives

$$\begin{aligned} & (\mathbf{Y}_1^\top \boldsymbol{\Sigma}_m \mathbf{Y}_1)^{-1/2} \mathbf{Y}_1^\top (\mathbf{x}_* - \mathbf{x}_m) \\ &= \underbrace{\mathbf{B} \boldsymbol{\Omega}_{11}^{-m} (\mathbf{Y}_1^\top \mathbf{Y}_1)^{-1} \left(\mathbf{Y}_1^\top (\mathbf{x}_* - \mathbf{f}) - \sum_{i=1}^{m-1} \mathbf{Y}_1^\top \mathbf{Y}_1 \boldsymbol{\Omega}_{11}^i \mathbf{W}_1^\top \mathbf{f} \right)}_{(*)} - \mathbf{B} \mathbf{W}_1^\top \mathbf{x}_0. \end{aligned} \quad (21)$$

We now focus on simplifying $(*)$. Left-multiplying Eq. (5) by \mathbf{Y}_1^\top gives

$$\begin{aligned} \mathbf{Y}_1^\top \mathbf{x}_* &= \mathbf{Y}_1^\top \mathbf{Y}_1 \boldsymbol{\Omega}_{11} \mathbf{W}_1^\top \mathbf{x}_* + \mathbf{Y}_1^\top \mathbf{f}. \\ \implies \mathbf{W}_1^\top \mathbf{x}_* &= \boldsymbol{\Omega}_{11}^{-1} (\mathbf{Y}_1^\top \mathbf{Y}_1)^{-1} \mathbf{Y}_1^\top (\mathbf{x}_* - \mathbf{f}) \end{aligned} \quad (22)$$

while left-multiplying by \mathbf{W}_1^\top gives

$$\begin{aligned} \mathbf{W}_1^\top \mathbf{x}_* &= \boldsymbol{\Omega}_{11} \mathbf{W}_1^\top \mathbf{x}_* + \mathbf{W}_1^\top \mathbf{f} \\ \implies \mathbf{W}_1^\top \mathbf{x}_* &= \boldsymbol{\Omega}_{11}^{-1} \mathbf{W}_1^\top (\mathbf{x}_* - \mathbf{f}). \end{aligned} \quad (23)$$

Substituting Eq. (22) into $(*)$ results in

$$\begin{aligned} (*) &= \mathbf{B} \boldsymbol{\Omega}_{11}^{-m} (\mathbf{Y}_1^\top \mathbf{Y}_1)^{-1} \left(\mathbf{Y}_1^\top (\mathbf{x}_* - \mathbf{f}) - \sum_{i=1}^{m-1} \mathbf{Y}_1^\top \mathbf{Y}_1 \boldsymbol{\Omega}_{11}^i \mathbf{W}_1^\top \mathbf{f} \right) \\ &= \mathbf{B} \boldsymbol{\Omega}_{11}^{-(m-1)} \left(\boldsymbol{\Omega}_{11}^{-1} (\mathbf{Y}_1^\top \mathbf{Y}_1)^{-1} \mathbf{Y}_1^\top (\mathbf{x}_* - \mathbf{f}) - \boldsymbol{\Omega}_{11}^{-1} (\mathbf{Y}_1^\top \mathbf{Y}_1)^{-1} \sum_{i=1}^{m-1} \mathbf{Y}_1^\top \mathbf{Y}_1 \boldsymbol{\Omega}_{11}^i \mathbf{W}_1^\top \mathbf{f} \right) \\ &= \mathbf{B} \boldsymbol{\Omega}_{11}^{-(m-1)} \left(\mathbf{W}_1^\top \mathbf{x}_* - \mathbf{W}_1^\top \mathbf{f} - \sum_{i=1}^{m-2} \boldsymbol{\Omega}_{11}^i \mathbf{W}_1^\top \mathbf{f} \right). \end{aligned}$$

³This is a square root in the sense of Section 1.4, a matrix $\mathbf{T}^{1/2}$ such that $\mathbf{T}^{1/2} (\mathbf{T}^{1/2})^\top = \mathbf{T}$.

Repeatedly substituting Eq. (23) into the previous equation gives

$$\begin{aligned}
(\star) &= \mathbf{B}\Omega_{11}^{-(m-1)} \left(\mathbf{W}_1^\top (\mathbf{x}_* - \mathbf{f}) - \sum_{i=1}^{m-2} \Omega_{11}^i \mathbf{W}_1^\top \mathbf{f} \right) \\
&= \mathbf{B}\Omega_{11}^{-(m-2)} \left(\Omega_{11}^{-1} \mathbf{W}_1^\top (\mathbf{x}_* - \mathbf{f}) - \Omega_{11}^{-1} \sum_{i=1}^{m-2} \Omega_{11}^i \mathbf{W}_1^\top \mathbf{f} \right) \\
&= \mathbf{B}\Omega_{11}^{-(m-2)} \left(\mathbf{W}_1^\top (\mathbf{x}_* - \mathbf{f}) - \sum_{i=1}^{m-3} \mathbf{W}_1^\top \mathbf{f} \right) \\
&\quad \vdots \\
&= \mathbf{B} \left(\Omega_{11}^{-1} \mathbf{W}_1^\top (\mathbf{x}_* - \mathbf{f}) \right) \\
&= \mathbf{B}\mathbf{W}_1^\top \mathbf{x}_*.
\end{aligned}$$

Finally substituting this back into Eq. (21) shows

$$(\mathbf{Y}_1^\top \Sigma_m \mathbf{Y}_1)^{-1/2} \mathbf{Y}_1^\top (\mathbf{x}_* - \mathbf{x}_m) = \mathbf{B}\mathbf{W}_1^\top (\mathbf{x}_* - \mathbf{x}_0). \quad (24)$$

Lastly we compute $\mathbf{W}_2^\top (\mathbf{x}_* - \mathbf{x}_m)$. This follows a similar argument to the above. We have

$$\mathbf{W}_2^\top (\mathbf{x}_* - \mathbf{x}_m) = \mathbf{W}_2^\top (\mathbf{x}_* - \mathbf{f}) \quad (25)$$

since $\mathbf{W}_2^\top \mathbf{G} = \mathbf{0}$. Similarly, left-multiplying the fixed-point equation Eq. (5) by \mathbf{W}_2^\top gives

$$\mathbf{W}_2^\top \mathbf{x}_* = \mathbf{W}_2^\top \mathbf{f}$$

Substituting this into Eq. (25) gives

$$\mathbf{W}_2^\top (\mathbf{x}_* - \mathbf{x}_m) = \mathbf{0}. \quad (26)$$

Completing the proof Eq. (26) validates the second requirement of Definition 3.7, since $\mathbf{N} = \mathbf{W}_2$. It remains to establish the first requirement. To accomplish this replace \mathbf{x}_* with $X \sim \mathcal{N}(\mathbf{x}_0, \Sigma_0)$ in Eq. (24). Since $\mathbf{W}_1^\top X \sim \mathcal{N}(\mathbf{W}_1^\top \mathbf{x}_0, \mathbf{W}_1^\top \Sigma_0 \mathbf{W}_1^\top)$, it follows that

$$\mathbf{B}\mathbf{W}_1^\top (X - \mathbf{x}_0) = (\mathbf{W}_1^\top \Sigma_0 \mathbf{W}_1^\top)^{-\frac{1}{2}} \mathbf{W}_1^\top (X - \mathbf{x}_0) \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_r).$$

which verifies the first requirement and completes the proof. \square

References

M. Arbel, A. Korba, A. Salim, and A. Gretton. Maximum mean discrepancy gradient flow. In H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché Buc, E. Fox, and R. Garnett, editors, *Advances in Neural Information Processing Systems 32*, pages 6484–6494. Curran Associates, Inc., 2019.

- S. Bartels and P. Hennig. Probabilistic approximate least-squares. In *Proceedings of the 19th International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 51 of *JMLR Workshop and Conference Proceedings*, pages 676–684, May 2016.
- S. Bartels, J. Cockayne, I. C. F. Ipsen, and P. Hennig. Probabilistic linear solvers: a unifying view. *Stat. Comput.*, 29(6):1249–1263, 2019.
- P. G. Bissiri, C. C. Holmes, and S. G. Walker. A general framework for updating belief distributions. *J. R. Stat. Soc. B*, 78(5):1103–1130, 2016.
- M. T. Chu. Linear algebra algorithms as dynamical systems. *Acta Numer.*, 17:1–86, 2008. ISSN 0962-4929. doi: 10.1017/S0962492906340019.
- J. Cockayne, C. J. Oates, I. C. F. Ipsen, and M. Girolami. A Bayesian conjugate gradient method (with discussion). *Bayesian Anal.*, 14(3):937–1012, 2019a.
- J. Cockayne, C. J. Oates, T. J. Sullivan, and M. Girolami. Bayesian probabilistic numerical methods. *SIAM Rev.*, 61(4):756–789, 2019b.
- S. R. Cook, A. Gelman, and D. B. Rubin. Validation of software for bayesian models using posterior quantiles. *J. Comput. Graph. Stat.*, 15(3):675–692, 2006.
- P. Diaconis. Bayesian numerical analysis. *Statistical Decision Theory and Related Topics IV*, 1:163–175, 1988.
- J. Geweke. Getting it right: Joint distribution tests of posterior simulators. *J. Am. Stat. Assoc.*, 99(467):799–804, 2004.
- G. H. Golub and C. F. Van Loan. *Matrix Computations*. The Johns Hopkins University Press, Baltimore, fourth edition, 2013.
- J. C. M. Graham, C. Oates, and T. Sullivan. Testing whether a learning procedure is calibrated, 2020.
- A. Gretton, K. M. Borgwardt, M. J. Rasch, B. Schölkopf, and A. Smola. A kernel two-sample test. *J. Mach. Learn. Res.*, 13(25):723–773, 2012.
- P. Hennig. Probabilistic interpretation of linear solvers. *SIAM J. Optim.*, 25(1):234–260, 2015. doi: 10.1137/140955501.
- P. Hennig, M. A. Osborne, and M. Girolami. Probabilistic numerics and uncertainty in computations. *J. R. Stat. Soc. A Stat.*, 471(2179):20150142, 17, 2015. ISSN 1364-5021. doi: 10.1098/rspa.2015.0142.
- M. R. Hestenes and E. Stiefel. Methods of conjugate gradients for solving linear systems. *J. Res. Nat. Bur. Stand.*, 49(6), December 1952.
- R. A. Horn and C. R. Johnson. *Matrix Analysis*. Cambridge University Press, 2009. doi: 10.1017/cbo9781139020411.

- I. C. F. Ipsen. *Numerical Matrix Analysis*. Society for Industrial and Applied Mathematics, Jan. 2009. doi: 10.1137/1.9780898717686.
- O. Kallenberg. *Foundations of Modern Probability*. Springer New York, 2002. doi: 10.1007/978-1-4757-4015-8.
- F. Larkin. Gaussian measure in Hilbert space and applications in numerical analysis. *Rocky Mt. J. Math.*, pages 379–421, 1972.
- J. Liesen and Z. Strakos. *Krylov Subspace Methods*. Oxford University Press, Oct. 2012. doi: 10.1093/acprof:oso/9780199655410.001.0001.
- C. Liu, J. Zhuo, P. Cheng, R. Zhang, and J. Zhu. Understanding and accelerating particle-based variational inference. volume 97 of *Proceedings of Machine Learning Research*, pages 4082–4092. PMLR, 2019.
- J. K. Reid. On the method of conjugate gradients for the solution of large sparse systems of linear equations. *Large Sparse Sets of Linear Equations (Proc. Conf. St. Catherine’s Coll., Oxford, 1970)*, pages 231–254, 1971.
- T. W. Reid, I. C. F. Ipsen, J. Cockayne, and C. J. Oates. A probabilistic numerical extension of the conjugate gradient method. *arXiv preprint arXiv:2008.03225*, 2020.
- Y. Saad. *Iterative Methods for Sparse Linear Systems*. Society for Industrial and Applied Mathematics, jan 2003. doi: 10.1137/1.9780898718003.
- R. C. Smith. *Uncertainty quantification*, volume 12 of *Computational Science & Engineering*. Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2014. Theory, implementation, and applications.
- S. Talts, M. Betancourt, D. Simpson, A. Vehtari, and A. Gelman. Validating Bayesian inference algorithms with simulation-based calibration. *arXiv preprint arXiv:1804.06788*, 2018.
- Y. L. Tong. *The Multivariate Normal Distribution*. Springer New York, 1990. doi: 10.1007/978-1-4613-9655-0.
- J. Wegner and P. Hennig. Probabilistic linear solvers for machine learning. *arXiv preprint arXiv:2010.09691*, 2020.
- D. M. Young. *Iterative Solution of Large Linear Systems*. Elsevier, 1971. doi: 10.1016/c2013-0-11733-3.
- D. M. Young. Second-degree iterative methods for the solution of large linear systems. *J. Approx. Theory*, 5:137–148, 1972.