

Probabilistic Linear Solvers: A Unifying View

Simon Bartels*, Jon Cockayne*, Ilse C. F. Ipsen, Mark Girolami and Philipp Hennig

Received: date / Accepted: date

Abstract Several recent works have developed a new, probabilistic interpretation for numerical algorithms solving linear systems in which the solution is inferred in a Bayesian framework, either directly or by inferring the unknown action of the matrix inverse. These approaches have typically focused on replicating the behavior of the conjugate gradient method as a prototypical iterative method. In this work surprisingly general conditions for equivalence of these disparate methods are presented. We also describe connections between probabilistic linear solvers and projection methods for linear systems,

providing a probabilistic interpretation of a far more general class of iterative methods. In particular, this provides such an interpretation of the generalised minimum residual method. A probabilistic view of preconditioning is also introduced. These developments unify the literature on probabilistic linear solvers, and provide foundational connections to the literature on iterative solvers for linear systems.

Keywords Probabilistic linear solvers · Projection methods · Iterative methods · Preconditioning

Simon Bartels
University of Tübingen
and Max Planck Institute for Intelligent Systems
72076 Tübingen, Germany
E-mail: sbartels@tue.mpg.de

Jon Cockayne
Department of Statistics
University of Warwick
Coventry, CV4 7AL, UK
E-mail: j.cockayne@warwick.ac.uk

Ilse C. F. Ipsen
Department of Mathematics
North Carolina State University
Raleigh, NC 27695-8205, USA
E-mail: ipsen@ncsu.edu

Mark Girolami
Department of Mathematics
Imperial College London
London, SW7 2AZ, UK
E-mail: m.girolami@imperial.ac.uk

Philipp Hennig
University of Tübingen
and Max Planck Institute for Intelligent Systems
72076 Tübingen, Germany
E-mail: ph@tue.mpg.de

* Authors contributed equally.

1 Introduction

Consider the linear system

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

where $A \in \mathbb{R}^{d \times d}$ is an invertible matrix, $\mathbf{b} \in \mathbb{R}^d$ is a given vector and $\mathbf{x}^* \in \mathbb{R}^d$ is an unknown to be determined. Recent work [Hennig, 2015, Cockayne et al., 2018] has constructed iterative solvers for this problem which output *probability measures*, constructed to quantify uncertainty due to terminating the algorithm before the solution has been identified completely. On the surface the approaches in these two works appear different: In the matrix-based inference (MBI) approach of Hennig [2015], a posterior is constructed on the matrix A^{-1} , while in the solution-based inference (SBI) method of Cockayne et al. [2018] a posterior is constructed on the solution vector \mathbf{x}^* .

These algorithms are instances of *probabilistic numerical methods* (PNM) in the sense of Hennig et al. [2015] and Cockayne et al. [2017]. PNM are numerical methods which output posterior distributions that quantify uncertainty due to discretisation error. An interesting property of PNM is that they often result in

a posterior distributions whose mean element coincides with the solution given by a classical numerical method for the problem at hand. The relationship between PNM and classical solvers has been explored for integration [e.g. Karvonen and Sarkka, 2017], ODE-solvers [Schober et al., 2014, 2018, Kersting et al., 2018] and PDE solvers [Cockayne et al., 2016] in some generality. For linear solvers, attention has thus far been restricted to the conjugate gradient (CG) method. Since CG is but a single member of a larger class of iterative solvers, and applicable only if the matrix A is symmetric and positive-definite, extending the probabilistic interpretation is still an interesting endeavour. Probabilistic interpretations provide an alternative perspective on numerical algorithms, and can also provide extensions such as the ability to exploit noisy or corrupted observations. The probabilistic view has also been used to the develop new numerical methods [Xi et al., 2018], and *Bayesian* PNM can be incorporated rigorously into pipelines of computation [Cockayne et al., 2017].

Preconditioning—mapping Eq. (1) to a better conditioned system which the same solution—is key to the fast convergence of iterative linear solvers, particularly those based upon Krylov methods [Liesen and Strakos, 2012]. The design of preconditioners has been referred to as “a combination of art and science” [Saad, 2003, p. 283]. In this work we also provide a new, probabilistic interpretation of preconditioning as a form of prior information.

1.1 Contribution

This text contributes three primary insights:

1. It is shown that, for particular choices of the generative model, matrix-based inference (MBI) and solution-based inference (SBI) can be equivalent (Section 2).
2. A general probabilistic interpretation of projection methods [Saad, 2003] is described (Section 3.1), leading to a probabilistic interpretation of the generalised minimum residual method (GMRES; Saad and Schultz [1986], Section 6). The connection to CG is expanded and made more concise in Section 5.
3. A probabilistic interpretation of preconditioning is presented in Section 4.

Most of the proofs are presented inline; lengthier proofs are deferred to Appendix B. While an important consideration, the predominantly theoretical contributions of this paper will not consider the impact of finite numerical precision.

1.2 Notation

For a symmetric positive-definite matrix $M \in \mathbb{R}^{d \times d}$ and two vectors $\mathbf{v}, \mathbf{w} \in \mathbb{R}^d$, we write $\langle \mathbf{v}, \mathbf{w} \rangle_M = \mathbf{v}^\top M \mathbf{w}$ for the inner product induced by M , and $\|\mathbf{v}\|_M^2 = \langle \mathbf{v}, \mathbf{v} \rangle_M$ for the corresponding norm.

A set of vectors $\mathbf{s}_1, \dots, \mathbf{s}_m$ is called *M-orthogonal* or *M-conjugate* if $\langle \mathbf{s}_i, \mathbf{s}_j \rangle_M = 0$ for $i \neq j$, and *M-orthonormal* if, in addition, $\|\mathbf{s}_i\|_M = 1$ for $1 \leq i \leq m$.

For a square matrix $A = [\mathbf{a}_1 \dots \mathbf{a}_d]^\top \in \mathbb{R}^{d \times d}$, the *vectorisation operator* $\text{vec} : \mathbb{R}^{d \times d} \rightarrow \mathbb{R}^{d^2}$ stacks the rows¹ of A into one long vector:

$$\vec{A} \equiv \text{vec}(A) = \begin{bmatrix} \mathbf{a}_1 \\ \vdots \\ \mathbf{a}_d \end{bmatrix}, \quad \text{with} \quad \left[\vec{A} \right]_{(ij)} = [A]_{ij}.$$

The *Kronecker product* of two matrices $A, B \in \mathbb{R}^{d \times d}$ is $A \otimes B$ with $[A \otimes B]_{(ij),(k\ell)} = [A]_{ik}[B]_{j\ell}$. A list of its properties is provided in Appendix A.

The Krylov space of order m generated by the matrix $A \in \mathbb{R}^{d \times d}$ and the vector $\mathbf{b} \in \mathbb{R}^d$ is

$$K_m(A, \mathbf{b}) = \text{span}(\mathbf{b}, A\mathbf{b}, A^2\mathbf{b}, \dots, A^{m-1}\mathbf{b}).$$

We will slightly abuse notation to describe shifted and scaled subspaces of \mathbb{R}^d : Let \mathbb{S} be an m -dimensional linear subspace of \mathbb{R}^d with basis $\{\mathbf{s}_1, \dots, \mathbf{s}_m\}$. Then for a vector $\mathbf{v} \in \mathbb{R}^d$ and a matrix $M \in \mathbb{R}^{d \times d}$, let

$$\mathbf{v} + M\mathbb{S} = \text{span}(\mathbf{v} + M\mathbf{s}_1, \dots, \mathbf{v} + M\mathbf{s}_m).$$

2 Probabilistic Linear Solvers

Several probabilistic framework describing the solution of Eq. (1) have been constructed in recent years. They primarily differ in the subject of inference: SBI approaches such as Cockayne et al. [2018], of which *BayesCG* is an example, place a prior distribution on the solution \mathbf{x}^* of Eq. (1). Conversely, the MBI approach of Hennig [2015] and Bartels and Hennig [2016] places a prior on A^{-1} , treating the action of the inverse operator as an unknown to be inferred². This section reviews each approach and adds some new insights. In particular, SBI can be viewed as strict special case of MBI (Section 2.4). Throughout this section, we will assume that the search

¹ Stacking the columns is equivalently possible and common. It is associated with a permutation in the definition of the Kronecker product, but the resulting inferences are equivalent.

² Hennig [2015] also discusses inference over A . This model class will not be discussed further in the present work. It has the disadvantage that the associated marginal on \mathbf{x}^* is non-analytic, but more easily lends itself to situations with noisy or otherwise perturbed matrix-vector products as observations.

directions S_m in $S_m^\top A \mathbf{x} = S_m^\top \mathbf{b}$ are given *a-priori*; Section 5 examines algorithms which iteratively generate search directions adapted to the problem at hand.

2.1 Background on Gaussian conditioning

The propositions in this section follow from the following two classic properties of Gaussian distributions.

Lemma 1 *Let $\mathbf{x} \in \mathbb{R}^d$ be Gaussian distributed with density $p(\mathbf{x}) = \mathcal{N}(\mathbf{x}; \mathbf{x}_0, \Sigma)$ for $\mathbf{x}_0 \in \mathbb{R}^d$ and $\Sigma \in \mathbb{R}^{d \times d}$ a positive semi-definite matrix. Let $M \in \mathbb{R}^{n \times d}$ and $\mathbf{z} \in \mathbb{R}^n$. Then $\mathbf{v} = M\mathbf{x} + \mathbf{z}$ is also Gaussian, with*

$$p(\mathbf{v}) = \mathcal{N}(\mathbf{v}; M\mathbf{x}_0 + \mathbf{z}, M\Sigma M^\top).$$

Lemma 2 *Let $\mathbf{x} \in \mathbb{R}^d$ be distributed as in Lemma 1, and let observations $\mathbf{y} \in \mathbb{R}^n$ be generated from the conditional density*

$$p(\mathbf{y} | \mathbf{x}) = \mathcal{N}(\mathbf{y}; M\mathbf{x} + \mathbf{z}, \Lambda)$$

with $M \in \mathbb{R}^{n \times d}$, $\mathbf{z} \in \mathbb{R}^n$, and $\Lambda \in \mathbb{R}^{n \times n}$ again positive-semidefinite. Then the associated conditional distribution on \mathbf{x} after observing \mathbf{y} is again Gaussian, with

$$p(\mathbf{x} | \mathbf{y}) = \mathcal{N}(\mathbf{x}; \bar{\mathbf{x}}, \bar{\Sigma}) \quad \text{where}$$

$$\bar{\mathbf{x}} = \mathbf{x}_0 + \Sigma M^\top (M\Sigma M^\top + \Lambda)^{-1} (\mathbf{y} - M\mathbf{x}_0 - \mathbf{z})$$

$$\bar{\Sigma} = \Sigma - \Sigma M^\top (M\Sigma M^\top + \Lambda)^{-1} M\Sigma.$$

This formula also applies if $\Lambda = 0$, i.e. observations are made without noise, with the caveat that if $M\Sigma M^\top$ is singular, the inverse should be interpreted as a pseudo-inverse.

2.2 Solution-Based Inference

To phrase the solution of Eq. (1) as a form of probabilistic inference, Cockayne et al. [2018] consider a Gaussian prior over the solution \mathbf{x}^* , and condition on observations provided by a set of search directions $\mathbf{s}_1, \dots, \mathbf{s}_m$, $m < d$. Let $S_m \in \mathbb{R}^{d \times m}$ be given by $S_m = [\mathbf{s}_1, \dots, \mathbf{s}_m]$, and let information be given by $\mathbf{y}_m := S_m^\top A \mathbf{x}^* = S_m^\top \mathbf{b}$. Since the information is clearly a linear projection of \mathbf{x}^* , the posterior distribution is a Gaussian distribution on \mathbf{x}^* :

Lemma 3 (Cockayne et al. [2018]) *Assume that the columns of S_m are linearly independent. Consider the prior*

$$p(\mathbf{x}) = \mathcal{N}(\mathbf{x}; \mathbf{x}_0, \Sigma).$$

The posterior from SBI is then given by

$$p(\mathbf{x} | \mathbf{y}_m) = \mathcal{N}(\mathbf{x}; \mathbf{x}_m, \Sigma_m)$$

where

$$\mathbf{x}_m = \mathbf{x}_0 + \Sigma_0 A^\top S_m (S_m^\top A \Sigma_0 A^\top S_m)^{-1} S_m^\top \mathbf{r}_0 \quad (2)$$

$$\Sigma_m = \Sigma_0 - \Sigma_0 A^\top S_m (S_m^\top A \Sigma_0 A^\top S_m)^{-1} S_m^\top \Sigma_0,$$

and $\mathbf{r}_0 = \mathbf{b} - A\mathbf{x}_0$.

The following proposition establishes an optimality property of the posterior mean \mathbf{x}_m . This is a relatively well-known property of Gaussian inference, but has not appeared before in the literature on these methods and will prove useful in subsequent sections.

Proposition 4 *If $S_m = \text{range}(S_m)$, then the posterior mean in Lemma 3 satisfies the optimality property*

$$\mathbf{x}_m = \arg \min_{\mathbf{x} \in \mathbf{x}_0 + \Sigma_0 A^\top S_m} \|\mathbf{x} - \mathbf{x}^*\|_{\Sigma_0^{-1}}.$$

Proof With the abbreviations $X = \Sigma_0 A^\top S_m$ and $\mathbf{y} = \mathbf{x}^* - \mathbf{x}_0$ the mean in Lemma 3 can be written as

$$\mathbf{x}_m = \mathbf{x}_0 + X\mathbf{c}_m,$$

where

$$\mathbf{c}_m = (X^\top \Sigma_0^{-1} X)^{-1} X^\top \Sigma_0^{-1} \mathbf{y}$$

is the solution of the weighted least squares problem [Golub and Van Loan, 2013, Section 6.1]

$$\mathbf{c}_m = \arg \min_{\mathbf{c} \in \mathbb{R}^m} \|X\mathbf{c} - \mathbf{y}\|_{\Sigma_0^{-1}}$$

$$= \arg \min_{\mathbf{c} \in \mathbb{R}^m} \|\mathbf{x}_0 + \Sigma_0 A^\top S_m \mathbf{c} - \mathbf{x}^*\|_{\Sigma_0^{-1}}.$$

This is equivalent to the desired statement. \square

2.3 Matrix-Based Inference

In contrast to SBI, the MBI approach of Hennig [2015] treats the matrix inverse A^{-1} as the unknown in the inference procedure. As in the previous section, search directions S_m yield matrix-vector products $Y_m \in \mathbb{R}^{d \times m}$. In Hennig [2015] these arise from *right-multiplying*³ A with S_m , i.e. $Y_m = AS_m$. Note that

$$S_m = A^{-1} Y_m, \text{ or, equivalently } \overrightarrow{S_m} = (I \otimes Y_m^\top) \overrightarrow{A^{-1}}. \quad (3)$$

Thus S_m is a linear transformation of A^{-1} and Lemma 2 can again be applied:

³ This work also considers a model class that explicitly encodes *symmetry* of A , such that the distinction between left- and right- multiplication vanishes. See Section 5.2 and Prop. 16 for more.

Lemma 5 (Lemma 2.1 in Hennig [2015]⁴) Consider the prior

$$p\left(\overrightarrow{A^{-1}}\right) = \mathcal{N}\left(\overrightarrow{A_0^{-1}}, \Sigma_0 \otimes W_0\right).$$

Then the posterior given the observations $\overrightarrow{S_m} = A^{-1}Y_m$ is given by

$$p\left(\overrightarrow{A^{-1}} \mid \overrightarrow{S_m}\right) = \mathcal{N}\left(\overrightarrow{A_m^{-1}}, \Sigma_0 \otimes W_m\right)$$

with

$$A_m^{-1} = A_0^{-1} + (S_m - A_0^{-1}Y_m)(Y_m^\top W_0 Y_m)^{-1} Y_m^\top W_0$$

$$W_m = W_0 - W_0 Y_m (Y_m^\top W_0 Y_m)^{-1} Y_m^\top W_0.$$

For linear solvers, the object of interest is $\mathbf{x}^* = A^{-1}\mathbf{b}$. Writing $A^{-1}\mathbf{b} = (I \otimes \mathbf{b}^\top) \overrightarrow{A^{-1}}$, and again using Lemma 1, we see that the associated marginal is also Gaussian, and given by

$$p(\mathbf{x} \mid S, Y) = \mathcal{N}(\mathbf{x}; A_m^{-1}\mathbf{b}, \mathbf{b}^\top W_m \mathbf{b} \cdot \Sigma_0). \quad (4)$$

In the Kronecker product specification for the prior covariance on A^{-1} , the first matrix, here Σ_0 , describes the dependence between the columns of A^{-1} . The second matrix, W_0 , captures the dependency between the rows of A^{-1} . Note that in Theorem 5, the posterior covariance has the form $\Sigma_0 \otimes W_m$. When compared to the prior covariance, $\Sigma_0 \otimes W_0$, it is clear that the observations have conveyed no new information to the first term of the Kronecker product covariance.

2.4 Equivalence of MBI and SBI

In practise Hennig [2015] notes that inference on A^{-1} should be performed only implicitly, avoiding the d^2 storage cost and the mathematical complexity of the operations involved in Lemma 5. This raises the question of when MBI is equivalent to SBI. Although, based on Lemma 1, one might suspect SBI and MBI to be equivalent, in fact the posterior from Lemma 5 is structurally different to the posterior in Lemma 3: After projecting into solution space, the posterior covariance in Lemma 5 is a scalar multiple of the matrix Σ_0 , which is not the case in general in Lemma 3.

However, the implied posterior over the solution vector can be made to coincide with the posterior from SBI if one considers observations in MBI as

$$S_m^\top = Y_m^\top A^{-1}. \quad (5)$$

That is, as *left*-multiplications of A . We will refer to the observation model of Eq. (3) as *right-multiplied information*, and to Eq. (5) as *left-multiplied information*.

⁴ This corrects a printing error in Hennig [2015]. The notation has been adapted to fit the context.

Proposition 6 Consider a Gaussian MBI prior

$$p(A^{-1}) = \mathcal{N}(A^{-1}; \overrightarrow{A_0^{-1}}, \Sigma_0 \otimes W_0),$$

conditioned on the left-multiplied information of Eq. (5). The associated marginal on \mathbf{x} is identical to the posterior on \mathbf{x} arising in Lemma 3 from $p(\mathbf{x}) = \mathcal{N}(\mathbf{x}; \mathbf{x}_0, \Sigma_0)$ under the conditions

$$A_0^{-1}\mathbf{b} = \mathbf{x}_0 \quad \text{and} \quad \mathbf{b}^\top W_0 \mathbf{b} = 1.$$

Proof See Appendix B. \square

The first of the two conditions requires that the prior mean on the matrix inverse be consistent with the prior mean on the solution, which is natural. The second condition demands that, after projection into solution space, the relationship between the rows of A^{-1} modelled by W_0 does not inflate the covariance Σ_0 . Note that this condition is trivial to enforce for an arbitrary covariance \bar{W}_0 by setting $W_0 = (\mathbf{b}^\top \bar{W}_0 \mathbf{b})^{-1} \bar{W}_0$.

2.5 Remarks

The result in Proposition 6 shows that any result proven for SBI applies immediately to MBI with left-multiplied observations. Though MBI has more model parameters than SBI, there are situations in which this point of view is more appropriate. Unlike in SBI, the information obtained in MBI need not be specific to a particular solution vector \mathbf{x}^* and thus can be propagated and recycled over several linear problems, similar to the notion of subspace recycling [Soodhalter et al., 2014]. Secondly, MBI is able to utilise both left- and right-multiplied information, while SBI is restricted to left-multiplied information. This additional generality may prove useful in some applications.

3 Projection Methods as Inference

This section discusses a connection between probabilistic numerical methods for linear systems and the classic framework of projection methods for the iterative solution of linear problems. Section 3.1 reviews this established class of solvers, while Section 3.2 presents the novel results.

3.1 Background

Many iterative methods for linear systems, including CG and GMRES, belong to the class of projection methods [Saad, 2003, p. 130f.]. Saad describes a projection

method as an iterative scheme in which, at each iteration, a solution vector \mathbf{x}_m is constructed by projecting \mathbf{x}^* into a solution space $\mathbb{X}_m \subset \mathbb{R}^d$, subject to the restriction that the residual $\mathbf{r}_m = \mathbf{b} - A\mathbf{x}_m$ is orthogonal to a constraint space $\mathbb{U}_m \subset \mathbb{R}^d$.

More formally, each iteration of a projection method is defined by two matrices $X_m, U_m \in \mathbb{R}^{d \times m}$, and by a starting point \mathbf{x}_0 . The matrices X_m and U_m each encode the solution and constraint spaces as $\mathbb{X}_m = \text{range}(X_m)$ and $\mathbb{U}_m = \text{range}(U_m)$. The projection method then constructs \mathbf{x}_m as $\mathbf{x}_m = \mathbf{x}_0 + X_m\boldsymbol{\alpha}_m$ with $\boldsymbol{\alpha}_m \in \mathbb{R}^m$ determined by the constraint $U_m^\top \mathbf{r}_m = \mathbf{0}$. This is possible only if $U_m^\top AX_m$ is nonsingular, in which case one obtains

$$\boldsymbol{\alpha}_m = (U_m^\top AX_m)^{-1} U_m^\top \mathbf{r}_0, \text{ and thus} \quad (6)$$

$$\mathbf{x}_m = \mathbf{x}_0 + X_m (U_m^\top AX_m)^{-1} U_m^\top \mathbf{r}_0. \quad (7)$$

From this perspective CG and GMRES perform only a single step with the number of iterations m fixed and determined in advance. For CG the spaces are $\mathbb{U}_m = \mathbb{X}_m = K_m(A, \mathbf{b})$, while for GMRES they are $\mathbb{X}_m = K_m(A, \mathbf{b})$ and $\mathbb{U}_m = AK_m(A, \mathbf{b})$ [Saad, 2003, Proposition 5.1].

3.2 Probabilistic Perspectives

In this section we first show, in Proposition 7, that the conditional mean from SBI after m steps corresponds to some projection method. Then, in Proposition 8 we prove the converse: that each projection method is also the posterior mean of a probabilistic method, for some prior covariance and choice of information.

Proposition 7 *Let the columns of S_m be linearly independent. Consider SBI under the prior*

$$p(\mathbf{x}) = \mathcal{N}(\mathbf{x}_0, \Sigma_0),$$

and with observations $\mathbf{y}_m = S_m^\top \mathbf{b}$. Then the posterior mean \mathbf{x}_m in Lemma 3 is identical to the iterate from a projection method defined by the matrices $U_m = S_m$ and $X_m = \Sigma_0 A^\top S_m$, and the starting vector \mathbf{x}_0 .

Proof Substituting $U_m = S_m$ and $X_m = \Sigma_0 A^\top S_m$ into Lemma 3 gives Eq. (7), as required. \square

The converse to this also holds:

Proposition 8 *Consider a projection method defined by the matrices $X_m, U_m \in \mathbb{R}^{d \times m}$, each with linearly independent columns, and the starting vector $\mathbf{x}_0 \in \mathbb{R}^d$. Then the iterate \mathbf{x}_m in Eq. (7) is identical to the SBI posterior mean in Lemma 3 under the prior*

$$p(\mathbf{x}) = \mathcal{N}(\mathbf{x}; \mathbf{x}_0, X_m X_m^\top) \quad (8)$$

when search directions $S_m = U_m$ are used.

Proof Abbreviate $Z = X_m^\top A^\top U_m$ and write the projection method iterate from Eq. (7) as

$$\mathbf{x}_m = \mathbf{x}_0 + X_m Z^{-T} U_m^\top \mathbf{r}_0.$$

Multiply the middle matrix by the identity,

$$\begin{aligned} Z^{-T} &= Z Z^{-1} Z^{-T} = Z (Z^\top Z)^{-1} \\ &= X_m^\top A^\top U_m (U_m^\top A \Sigma_0 A^\top U_m)^{-1}, \end{aligned}$$

and insert this into the expression for \mathbf{x}_0 ,

$$\mathbf{x}_m = \mathbf{x}_0 + \Sigma_0 A^\top U_m (U_m^\top A \Sigma_0 A^\top U_m)^{-1} U_m^\top \mathbf{r}_0.$$

Setting $U_m = S_m$ gives the mean in Lemma 3. \square

Including a basis of the solution space in the prior may seem problematic. A direct way to enforce the posterior occupying the solution space is by placing a prior on the coefficients $\boldsymbol{\alpha}$ in $\mathbf{x} = \mathbf{x}_0 + X_m \boldsymbol{\alpha}$. Under a unit Gaussian prior $\boldsymbol{\alpha} \sim \mathcal{N}(\mathbf{0}, I)$, the implied prior on \mathbf{x} naturally has the form of Eq. (8). However, this prior is nevertheless unsatisfying both since it requires the solution space to be specified *a-priori*, precluding adaptivity in the algorithm, and, perhaps more worryingly, because the posterior uncertainty over the solution is a matrix of zeros even though the solution is not fully identified. Again taking $Z = X_m^\top A^\top U_m$:

$$\begin{aligned} \Sigma_m &= \Sigma_0 - \Sigma_0 A^\top U_m (U_m^\top A \Sigma_0 A^\top U_m)^{-1} U_m^\top A \Sigma_0 \\ &= X_m X_m^\top - X_m Z (Z^\top Z)^{-1} Z^\top X_m^\top \\ &= X_m X_m^\top - X_m X_m^\top \\ &= 0. \end{aligned}$$

[Hennig, 2015] and [Bartels and Hennig, 2016] each proposed to address this issue by adding additional uncertainty in the null space of X_m . This empirical uncertainty calibration step has not yet been analysed in detail. Such analysis is left for future work. Nevertheless, the proposition provides a probabilistic view for *arbitrary* projection methods and does not require knowledge of A^{-1} , unlike some of the results presented in [Hennig, 2015, Cockayne et al., 2017] and in the following propositions.

This prior is not unique. The next proposition establishes more restrictive conditions under which a projection method may have a probabilistic interpretation and still result in a nonzero posterior uncertainty.

Proposition 9 *Consider a projection method defined by $X_m, U_m \in \mathbb{R}^{d \times m}$ and the starting vector \mathbf{x}_0 . Further suppose that $U_m = R X_m$ for some invertible $R \in \mathbb{R}^{d \times d}$, and that $A^\top R$ is symmetric positive-definite. Then under the prior*

$$p(\mathbf{x}) = \mathcal{N}(\mathbf{x}; \mathbf{x}_0, (A^\top R)^{-1})$$

and the search directions $S_m = U_m = RX_m$, the iterate in the projection method is identical to the posterior mean in Lemma 3.

Proof First substitute $X_m = R^{-1}U_m$ into Eq. (7) to obtain

$$\begin{aligned} \mathbf{x}_m &= \mathbf{x}_0 + R^{-1}U_m(U_m^\top AR^{-1}U_m)^{-1}U_m^\top \mathbf{r}_0 \\ &= \mathbf{x}_0 + R^{-1}A^{-\top}A^\top U_m(U_m^\top AR^{-1}A^{-\top}A^\top U_m)^{-1}U_m^\top \mathbf{r}_0 \\ &= \mathbf{x}_0 + \Sigma_0 A^\top U_m(U_m^\top A \Sigma_0 A^\top U_m)^{-1}U_m^\top \mathbf{r}_0. \end{aligned}$$

The third line uses $\Sigma_0 = (A^\top R)^{-1} = R^{-1}A^{-\top}$. This is equivalent to the posterior mean in Eq. (2) with $S_m = U_m$. \square

A corollary which provides further insight arises when one considers the polar decomposition of A . Recall that an invertible matrix A has a unique polar decomposition $A = PH$, where $P \in \mathbb{R}^{d \times d}$ is orthogonal and $H \in \mathbb{R}^{d \times d}$ is symmetric positive-definite.

Corollary 10 Consider a projection method defined by $X_m, U_m \in \mathbb{R}^{d \times m}$ and the starting vector \mathbf{x}_0 , and suppose that $U_m = PX_m$, where P arises from the polar decomposition $A = PH$. Then under the prior

$$p(\mathbf{x}) = \mathcal{N}(\mathbf{x}; \mathbf{x}_0, H^{-1})$$

and the search directions $S_m = U_m = PX_m$, the iterate in the projection method is identical to the posterior mean in Lemma 3.

Proof This follows from Proposition 9. Setting $R = P$ aligns the search directions in Corollary 10 with those in Proposition 9. Since P is orthogonal, $P^{-1} = P^\top$, and since H is symmetric positive-definite, $A^\top P = P^\top A = H$ by definition of the polar decomposition, which gives the prior covariance required for Proposition 9. \square

This is an intuitive analogue of similar results in Hennig [2015] and Cockayne et al. [2017] which show that CG is recovered under certain conditions involving a prior $\Sigma_0 = A^{-1}$. When A is not symmetric and positive definite it cannot be used as a prior covariance. This corollary suggests a natural way to select a prior covariance still linked to the linear system, though this choice is still not computationally convenient. Furthermore, in the case that A is symmetric positive-definite, this recovers the prior which replicates CG described in Cockayne et al. [2018]. Note that each of H and P can be stated explicitly as $H = (A^\top A)^{\frac{1}{2}}$ and $P = A(A^\top A)^{-\frac{1}{2}}$. Thus in the case of symmetric positive-definite A we have that $H = A$ and $P = I$, so that the prior covariance $\Sigma_0 = A^{-1}$ arises naturally from this interpretation.

4 Preconditioning

This section discusses probabilistic views on preconditioning. Preconditioning is a widely-used technique accelerating the convergence of iterative methods [Saad, 2003, Sections 9 and 10]. A preconditioner P is a nonsingular matrix satisfying two requirements:

1. Linear systems $Pz = c$ can be solved at low computational cost (i.e. ‘‘analytically’’)
2. P is ‘‘close’’ to A in some sense.

In this sense, solving systems based upon a preconditioner can be viewed as approximately inverting A , and indeed many preconditioners are constructed based upon this intuition. One distinguishes between *right preconditioners* P_r and *left preconditioners* P_l , depending on whether they act on A from the left or the right. Two-sided preconditioning with nonsingular matrices P_l and P_r transforms implicitly Eq. (1) into a new linear problem

$$P_l A P_r \mathbf{z}^* = P_l \mathbf{b}, \quad \text{with } \mathbf{x}^* = P_r \mathbf{z}^*. \quad (9)$$

The preconditioned system can then be solved using arbitrary projection methods as described in Section 3.1, from the starting point \mathbf{z}_0 defined by $\mathbf{x}_0 = P_r \mathbf{z}_0$. The probabilistic view can be used to create a nuanced description of preconditioning as a form of prior information. In the SBI framework, Proposition 11 below shows that solving a right-preconditioned system is equivalent to modifying the prior, while in Proposition 12 shows that left-preconditioning is equivalent to making a different choice of observations.

Proposition 11 (Right preconditioning) Consider the right-preconditioned system

$$A P_r \mathbf{z}^* = \mathbf{b} \quad \text{where } \mathbf{x}^* = P_r \mathbf{z}^*. \quad (10)$$

SBI on Eq. (10) under the prior

$$\mathbf{z} \sim \mathcal{N}(\mathbf{z}; \mathbf{z}_0, \Sigma_0) \quad (11)$$

is equivalent to solving Eq. (1) under the prior

$$\mathbf{x} \sim \mathcal{N}(\mathbf{x}; P_r \mathbf{z}_0, P_r \Sigma_0 P_r^\top).$$

Proof Let $p(\mathbf{x}) = \mathcal{N}(\mathbf{x}; \mathbf{x}_0, \Sigma_r)$. Lemma 3 implies that after observing information from search directions S_m , the posterior mean equals

$$\mathbf{x}_m = \mathbf{x}_0 + \Sigma_r A^\top S_m (S_m^\top A \Sigma_r A^\top S_m)^{-1} S_m^\top \mathbf{r}_0$$

where $\mathbf{r}_0 = \mathbf{b} - A\mathbf{x}_0$. Setting $\mathbf{x}_0 = P_r \mathbf{z}_0$ and letting $\Sigma_r = P_r \Sigma_0 P_r^\top$ gives

$$\mathbf{x}_m = P_r \mathbf{z}_0 + P_r \Sigma_0 B^\top S_m (S_m^\top B \Sigma_0 B^\top S_m)^{-1} S_m^\top \hat{\mathbf{r}}_0$$

where $B := AP_r$ and $\hat{\mathbf{r}}_0 = \mathbf{b} - B\mathbf{z}_0$. Left multiplying by P_r^{-1} shows that this is equivalent to

$$\begin{aligned} \mathbf{z}_m &:= P_r^{-1} \mathbf{x}_m \\ &= \mathbf{z}_0 + \Sigma_0 B^\top S_m (S_m^\top B \Sigma_0 B^\top S_m)^{-1} S_m^\top \hat{\mathbf{r}}_0. \end{aligned}$$

Thus \mathbf{z}_m is the posterior mean of the system $B\mathbf{z}^* = \mathbf{b}$ with prior Eq. (11) after observing search directions S_m . \square

Proposition 12 (Left preconditioning) *Consider the left-preconditioned system*

$$P_l A \mathbf{x}^* = P_l \mathbf{b} \quad (12)$$

And the SBI prior

$$p(\mathbf{x}) = \mathcal{N}(\mathbf{x}; \mathbf{x}_0, \Sigma_0).$$

Then the posterior from SBI on Eq. (12) under search directions S_m is equivalent to the posterior from SBI applied to the system Eq. (1) under search directions $P_l^\top S_m$.

Proof Lemma 3 implies that after observing search directions T_m , the posterior mean over the solution of Eq. (1) equals

$$\mathbf{x}_m = \mathbf{x}_0 + \Sigma_0 A^\top T_m (T_m^\top A \Sigma_0 A^\top T_m)^{-1} T_m^\top \mathbf{r}_0$$

where $\mathbf{r}_0 = \mathbf{b} - A\mathbf{x}_0$. Setting $T_m = P_l^\top S_m$ gives

$$\mathbf{x}_m = \mathbf{x}_0 + \Sigma_0 B^\top S_m (S_m^\top B \Sigma_0 B^\top S_m)^{-1} S_m^\top P_l \hat{\mathbf{r}}_0$$

where $B := P_l A$ and $\hat{\mathbf{r}}_0 = P_l \mathbf{b} - P_l A \mathbf{x}_0$. Thus, \mathbf{x}_m is the posterior mean of the system $B\mathbf{x}^* = P_l \mathbf{b}$ after observing search directions S_m . \square

If a probabilistic linear solver has a posterior mean which coincides with a projection method (as discussed in Section 3.1), the Propositions 11 and 12 show how to obtain a probabilistic interpretation of the preconditioned version of that algorithm. Furthermore, the equivalence demonstrated in Section 2.4 shows that the reasoning from Propositions 11 and 12 carries over to MBI based on left-multiplied observations: right-preconditioning corresponds to a change in prior belief, while left-preconditioning corresponds to a change in observations.

We do not claim that this probabilistic interpretation of preconditioning is unique. For example, when using MBI with right-multiplied observations, the same line of reasoning can be used to show the converse: right-preconditioning corresponds to a change in the observations and left-preconditioning to a change in the prior.

5 Conjugate Gradients

Conjugate gradients has been studied from a probabilistic point of view before by Hennig [2015] and Cockayne et al. [2018]. This section generalizes the results of Hennig [2015] and leverages Proposition 6 for new insights on BayesCG. For this Section (but not thereafter) assume that A is a symmetric and positive definite matrix.

5.1 Left-multiplied view

The BayesCG algorithm proposed by Cockayne et al. [2018] encompasses conjugate gradients as a special case. BayesCG uses left-multiplied observations and was derived in the solution-based perspective.

The posterior in Lemma 3 does not immediately result in a practical algorithm as it involves the solution of a linear system based on the matrix $S_m^\top A \Sigma_0 A^\top S_m \in \mathbb{R}^{m \times m}$, which requires $\mathcal{O}(m^3)$ arithmetic operations. BayesCG avoids this cost by constructing search directions that are $A \Sigma_0 A^\top$ -orthonormal, as shown below, see [Cockayne et al., 2018, Proposition 7].

Proposition 13 (Proposition 7 of Cockayne et al. [2018] (BayesCG)) *Let $\tilde{\mathbf{s}}_1 = \mathbf{b} - A\mathbf{x}_0$, and let $\mathbf{s}_1 = \tilde{\mathbf{s}}_1 / \|\tilde{\mathbf{s}}_1\|$. For $j = 2, \dots, m$ let*

$$\begin{aligned} \tilde{\mathbf{s}}_j &= \mathbf{b} - A\mathbf{x}_{j-1} - \langle \mathbf{b} - A\mathbf{x}_{j-1}, \mathbf{s}_{j-1} \rangle_{A \Sigma_0 A^\top} \mathbf{s}_{j-1} \\ \mathbf{s}_j &= \tilde{\mathbf{s}}_j / \|\tilde{\mathbf{s}}_j\|_{A \Sigma_0 A^\top}. \end{aligned}$$

Then the set $\{\mathbf{s}_1, \dots, \mathbf{s}_m\}$ is $A \Sigma_0 A^\top$ -orthonormal, and consequently $S_m^\top A \Sigma_0 A^\top S_m = I$.

With these search directions constructed, BayesCG becomes an iterative method:

Proposition 14 (Proposition 6 of Cockayne et al. [2018]) *Using the search directions from Proposition 13, the posterior from Lemma 3 reduces to:*

$$\begin{aligned} \mathbf{x}_m &= \mathbf{x}_{m-1} + \Sigma_0 A^\top \mathbf{s}_m (\mathbf{s}_m^\top (\mathbf{b} - A\mathbf{x}_{m-1})) \\ \Sigma_m &= \Sigma_{m-1} - \Sigma_0 A^\top \mathbf{s}_m \mathbf{s}_m^\top A \Sigma_0 \end{aligned}$$

In Proposition 4 of Cockayne et al. [2018] it was shown that the BayesCG posterior mean corresponds to the CG solution estimate when the prior covariance is taken to be $\Sigma_0 = A^{-1}$, though this is not a practical choice of prior covariance as it requires access to the unavailable A^{-1} . Furthermore, in Proposition 9 it was shown that when using the search directions from Proposition 13, the posterior mean from BCG has the following optimality property:

$$\mathbf{x}_m = \arg \min_{\mathbf{x} \in K_m(\Sigma_0 A^\top A, \Sigma_0 A^\top \mathbf{b})} \|\mathbf{x} - \mathbf{x}^*\|_{\Sigma_0^{-1}}$$

Note that this is now a trivial special case of Proposition 4.

The following proposition leverages these results along with Proposition 6 to show that there exists an MBI method which, under a particular choice of prior and with a particular methodology for the generation of search directions, is consistent with CG.

Proposition 15 *Consider the MBI prior*

$$p(\overrightarrow{A^{-1}}) = \mathcal{N}(\overrightarrow{A^{-1}}; \overrightarrow{A_0^{-1}}, A^{-1} \otimes W_0)$$

where $W_0 \in \mathbb{R}^d$ is symmetric positive-definite and so that $\mathbf{b}^\top W_0 \mathbf{b} = 1$. Suppose left-multiplied information is used, and that the search directions are generated sequentially according to:

$$\tilde{\mathbf{s}}_1 = (I - AA_0^{-1})\mathbf{b}$$

$$\mathbf{s}_1 = \frac{\tilde{\mathbf{s}}_1}{\|\tilde{\mathbf{s}}_1\|_A}$$

and for $j = 2, \dots, m$

$$\tilde{\mathbf{s}}_j = (I - AA_{j-1}^{-1})\mathbf{b} - \mathbf{b}^\top (I - AA_{j-1}^{-1})^\top A \mathbf{s}_{j-1} \cdot \mathbf{s}_{j-1}$$

$$\mathbf{s}_j = \frac{\tilde{\mathbf{s}}_j}{\|\tilde{\mathbf{s}}_j\|_A}.$$

Then it holds that the implied posterior mean on solution space, given by $A_m^{-1}\mathbf{b}$, corresponds to the CG solution estimate after m iterations, with starting point $\mathbf{x}_0 = A_0^{-1}\mathbf{b}$.

Proof First note that, by Proposition 6, since left-multiplied observations are used and since $\mathbf{b}^\top W_0 \mathbf{b} = 1$, the implied posterior distribution on solution space from MBI is identical to the posterior distribution from SBI under the prior

$$p(\mathbf{x}) = \mathcal{N}(\mathbf{x}; A_0^{-1}\mathbf{b}, A^{-1}).$$

It thus remains to show that the sequence of search directions generated is identical to those in Proposition 13 for this prior. For $\tilde{\mathbf{s}}_1$:

$$\tilde{\mathbf{s}}_1 = (I - AA_0^{-1})\mathbf{b} = \mathbf{b} - A\mathbf{x}_0$$

as required. For $\tilde{\mathbf{s}}_j$:

$$\begin{aligned} \tilde{\mathbf{s}}_j &= (I - AA_{j-1}^{-1})\mathbf{b} - \mathbf{b}^\top (I - AA_{j-1}^{-1})^\top A \mathbf{s}_{j-1} \cdot \mathbf{s}_{j-1} \\ &= \mathbf{b} - A\mathbf{x}_{m-1} - (\mathbf{b} - A\mathbf{x}_{j-1})^\top A \mathbf{s}_{j-1} \cdot \mathbf{s}_{j-1} \\ &= \mathbf{b} - A\mathbf{x}_{m-1} - \langle \mathbf{b} - A\mathbf{x}_{j-1}, \mathbf{s}_{j-1} \rangle_A \cdot \mathbf{s}_{j-1} \end{aligned}$$

where the second line uses that $A_{j-1}^{-1}\mathbf{b} = \mathbf{x}_{j-1}$. Thus, the search directions coincide with those in Proposition 13. It therefore holds that the implied posterior mean on solution space, $A_m^{-1}\mathbf{b}$, coincides with the solution estimate produced by CG. \square

5.2 Right-multiplied view

Interpretations of CG (and general projection methods) that use right-multiplied observations seems to require more care than those based on left-multiplied observations. Nevertheless, Hennig [2015] provided an interpretation for CG in this framework, essentially showing⁵ that Algorithm 1 reproduces both the search directions and solution estimates from CG under the prior

$$p(A^{-1}) = \mathcal{N}(\overrightarrow{A^{-1}}; \overrightarrow{\alpha I}, \beta A^{-1} \otimes A^{-1}).$$

where $\alpha \in \mathbb{R} \setminus \{0\}$, $\beta \in \mathbb{R}^+$ and \otimes denotes the symmetric Kronecker product (see Section A.1). The posterior under such a prior is described in Lemma 2.2 of Hennig [2015] (see Lemma 21), though we note that the sense in which the solution estimate \mathbf{x}_m output by this algorithm is related to the posterior over A^{-1} differs from that in the previous section, in the sense that $A_m^{-1}\mathbf{b} \neq \mathbf{x}_m$. (More precisely, $\mathbf{x}_m = A_m^{-1}(\mathbf{b} - A\mathbf{x}_0) - \mathbf{x}_0 - (1 - \alpha_m)\mathbf{d}_m$, as the CG estimate is corrected by the step size computed in line 6. Fixing this rank-1 discrepancy would complicate the exposition of Algorithm 1 and yield a more cumbersome algorithm). The following proposition generalizes this result.

Proposition 16 *Consider the prior*

$$p(A^{-1}) = \mathcal{N}(A^{-1}; \alpha I, (\beta I + \gamma A^{-1}) \otimes (\beta I + \gamma A^{-1})).$$

For all choices $\alpha \in \mathbb{R} \setminus \{0\}$ and $\beta, \gamma \in \mathbb{R}_{+,0}$ with $\beta + \gamma > 0$, Algorithm 1 is equivalent to CG, in the sense that it produces the exact same sequence of estimates \mathbf{x}_i and scaled search directions \mathbf{s}_i .

Proof The proof is extensive and has been moved to Appendix B. \square

Algorithm 1 The algorithm referred to by Proposition 16, which reproduces the search directions and solution estimates from CG.

```

1  $\mathbf{x}_0 \leftarrow A_0^{-1}\mathbf{b}$  // initial guess
2  $\mathbf{r}_0 \leftarrow A\mathbf{x}_0 - \mathbf{b}$ 
3 for  $i = 1, \dots, m$  do
4    $\mathbf{d}_i \leftarrow -A_{i-1}^{-1}\mathbf{r}_{i-1}$  // compute optimization direction
5    $\mathbf{z}_i \leftarrow A\mathbf{d}_i^\top \mathbf{r}_{i-1}$  // observe
6    $\alpha_i \leftarrow -\frac{\mathbf{d}_i^\top \mathbf{r}_{i-1}}{\mathbf{d}_i^\top \mathbf{z}_i}$  // optimal step-size
7    $\mathbf{s}_i \leftarrow \alpha_i \mathbf{d}_i$  // re-scale step
8    $\mathbf{y}_i \leftarrow \alpha_i \mathbf{z}_i$  // re-scale observation
9    $\mathbf{x}_i \leftarrow \mathbf{x}_{i-1} + \mathbf{s}_i$  // update estimate for  $\mathbf{x}$ 
10   $\mathbf{r}_i \leftarrow \mathbf{r}_{i-1} + \mathbf{y}_i$  // new gradient at  $\mathbf{x}_i$ 
11   $A_i^{-1} \leftarrow \mathbb{E}_{p(A^{-1}|\mathcal{S}, Y)} A^{-1}$  // estimate  $A^{-1}$ 
12 end for
13 return  $\mathbf{x}_m$ 

```

⁵ Algorithm 1 is not included in this form in the op.cit.

Note that, unlike previous propositions, Proposition 16 proposes a prior that does not involve A^{-1} for the case when $\gamma = 0$.

6 GMRES

The *Generalised Minimal Residual Method* [Saad, 2003, Section 6.5] applies to general nonsingular matrices A . At iteration m , GMRES minimises the residual over the affine space $\mathbf{x}_0 + K_m(A, \mathbf{r}_0)$. That is, $\mathbf{r}_m = \mathbf{r}_0 - A\mathbf{x}_m$ satisfies

$$\begin{aligned} \|\mathbf{r}_m\|_2 &= \min_{\mathbf{x} \in K_m(A, \mathbf{r}_0)} \|A\mathbf{x} - \mathbf{r}_0\|_2 \\ &= \min_{\mathbf{x} \in \mathbf{x}_0 + K_m(A, \mathbf{r}_0)} \|A\mathbf{x} - \mathbf{b}\|_2. \end{aligned} \quad (13)$$

Since $A\mathbf{x} - \mathbf{b} = A(\mathbf{x} - \mathbf{x}^*)$, this corresponds to minimizing the error in the $A^\top A$ norm.

We present a brief development of GMRES, starting with Arnoldi's method (Section 6.1) and the GMRES algorithm (Section 6.2), before presenting our Bayesian interpretation (Section 6.3).

6.1 Arnoldi's Method

GMRES uses Arnoldi's method [Saad, 2003, Section 6.3] to construct orthonormal bases for Krylov spaces of general, nonsingular matrices A . Starting with $\mathbf{q}_1 = \mathbf{r}_0 / \|\mathbf{r}_0\|_2$, Arnoldi's method recursively computes the orthonormal basis

$$Q_m = [\mathbf{q}_1 \dots \mathbf{q}_m] \in \mathbb{R}^{d \times m}$$

for $K_m(A, \mathbf{r}_0)$. The basis vectors satisfy the relations

$$AQ_m = Q_{m+1} \tilde{H}_m = Q_m H_m + h_{m+1,m} \mathbf{q}_{m+1} \mathbf{e}_m^\top \quad (14)$$

and $Q_m^\top A Q_m = H_m$, where the *upper Hessenberg* matrix H_m is defined as

$$H_m = \begin{bmatrix} h_{11} & h_{12} & h_{13} & \dots & h_{1,m-1} & h_{1m} \\ h_{21} & h_{22} & h_{23} & \dots & h_{2,m-1} & h_{2m} \\ 0 & h_{32} & h_{33} & \dots & h_{3,m-1} & h_{3m} \\ \vdots & 0 & h_{43} & \dots & h_{4,m-1} & h_{4m} \\ \vdots & & \ddots & \ddots & \vdots & \vdots \\ 0 & \dots & \dots & 0 & h_{m,m-1} & h_{mm} \end{bmatrix} \in \mathbb{R}^{m \times m}$$

and

$$\tilde{H}_m = \begin{bmatrix} H_m \\ h_{m+1,m} \mathbf{e}_m^\top \end{bmatrix} \in \mathbb{R}^{(m+1) \times m}.$$

6.2 GMRES

GMRES computes the iterate

$$\mathbf{x}_m = \mathbf{x}_0 + Q_m \mathbf{c}_m$$

based on the optimality condition in Eq. (13), which can equivalently be expressed as

$$\begin{aligned} \mathbf{c}_m &= \arg \min_{\mathbf{c} \in \mathbb{R}^m} \|AQ_m \mathbf{c} - \mathbf{r}_0\|_2 \\ &= ((AQ_m)^\top (AQ_m))^{-1} (AQ_m)^\top \mathbf{r}_0. \end{aligned} \quad (15)$$

Thus

$$\mathbf{x}_m = \mathbf{x}_0 + Q_m (Q_m^\top A^\top A Q_m)^{-1} Q_m^\top A^\top \mathbf{r}_0, \quad (16)$$

confirming that GMRES is a projection method with $X_m = Q_m$ and $U_m = AQ_m$.

GMRES solves the least squares problem in Eq. (15) efficiently by projecting it to a lower dimensional space via Arnoldi's method. To this end, express the starting vector in the Krylov basis,

$$\mathbf{r}_0 = \|\mathbf{r}_0\|_2 \mathbf{q}_1 = \|\mathbf{r}_0\|_2 Q_{m+1} \mathbf{e}_1,$$

and exploit the Arnoldi recursion from Eq. (14),

$$AQ_m \mathbf{c} - \mathbf{r}_0 = Q_{m+1} (\tilde{H}_{m+1} \mathbf{c} - \|\mathbf{r}_0\|_2 \mathbf{e}_1),$$

followed by the unitary invariance of the two-norm,

$$\|AQ_m \mathbf{c} - \mathbf{r}_0\|_2 = \|\tilde{H}_m \mathbf{c} - \|\mathbf{r}_0\|_2 \mathbf{e}_1\|_2.$$

Thus, instead of solving the least squares problem Equation (15) with d rows, GMRES solves instead a problem with only $m+1$ rows,

$$\mathbf{c}_m = \arg \min_{\mathbf{c} \in \mathbb{R}^m} \|\tilde{H}_m \mathbf{c} - \|\mathbf{r}_0\|_2 \mathbf{e}_1\|_2. \quad (17)$$

The computations are summarized in Algorithm 2.

Algorithm 2 GMRES [Saad, 2003, Algorithm 6.9]

```

1  $\mathbf{r}_0 \leftarrow \mathbf{b} - A\mathbf{x}_0$ ,  $\beta \leftarrow \|\mathbf{r}_0\|_2$ ,  $\mathbf{q}_1 \leftarrow \mathbf{r}_0/\beta$ 
2 for  $j = 1, \dots, m$  do
3    $\mathbf{w}_j \leftarrow A\mathbf{q}_j$ 
4   for  $i = 1, \dots, j$  do
5      $h_{ij} \leftarrow \langle \mathbf{w}_j, A\mathbf{q}_i \rangle$ 
6      $\mathbf{w}_j \leftarrow \mathbf{w}_j - h_{ij} \mathbf{q}_i$ 
7   end for
8    $h_{j+1,j} \leftarrow \|\mathbf{w}_j\|_2$ 
9   if  $h_{j+1,j} = 0$  then
10     $m \leftarrow j$ , go to 14
11  end if
12   $\mathbf{q}_{j+1} \leftarrow \mathbf{w}_j/h_{j+1,j}$ 
13 end for
14 Define  $\tilde{H}_m \in \mathbb{R}^{(m+1) \times m}$  with elements  $h_{ij}$ 
15  $\mathbf{c}_m \leftarrow \arg \min_{\mathbf{c}} \|\tilde{H}_m \mathbf{c} - \beta \mathbf{e}_1\|_2$ 
16  $\mathbf{x}_m \leftarrow \mathbf{x}_0 + Q_m \mathbf{c}_m$ 

```

6.3 Bayesian Interpretation of GMRES

We now present probabilistic linear solvers with posterior means that coincide with the solution estimate from GMRES.

6.3.1 Left-multiplied view

Proposition 17 *Under the SBI prior*

$$p(\mathbf{x}) = \mathcal{N}(\mathbf{x}; \mathbf{x}_0, \Sigma_0) \quad \text{where} \quad \Sigma_0 = (A^\top A)^{-1}$$

and the search directions $U_m = AQ_m$, the posterior mean is identical to the GMRES iterate \mathbf{x}_m in Eq. (16).

Proof Substitute $R = A$ and $U_m = AQ_m$ into Proposition 9. \square

Proposition 17 is intuitive in the context of Proposition 4: Setting $\Sigma_0 = (A^\top A)^{-1}$ ensures that the norm being minimised coincides with that of GMRES, as does the solution space $X_m = AQ_m$. This interpretation exhibits an interesting duality with CG for which $\Sigma_0 = A^{-1}$.

Another probabilistic interpretation follows from Proposition 8.

Corollary 18 *Under the prior*

$$p(\mathbf{x}) = \mathcal{N}(\mathbf{x}; \mathbf{x}_0, \Sigma_0) \quad \text{where} \quad \Sigma_0 = Q_m Q_m^\top, \quad (18)$$

and with observations $\mathbf{y}_m = Q_m^\top \mathbf{b}$, the posterior mean from SBI is identical to the GMRES iterate \mathbf{x}_m in Eq. (16).

Note that Proposition 17 has a posterior covariance which is not practical, as it involves A^{-1} . [Cockayne et al., 2017] proposed replacing A^{-1} in the prior covariance with a preconditioner to address this, which does yield a practically computable posterior, but this extension was not explored here. Furthermore, that approach yields poorly calibrated posterior uncertainty, as described in that work. Corollary 18 does not have this drawback, but the posterior covariance is a matrix of zeroes.

6.3.2 Right-multiplied view

As for CG in Section 5.2, finding interpretations of GMRES that use right-multiplied observations appears to be more difficult.

Proposition 19 *Under the prior*

$$p(A^{-1}) = \mathcal{N}(0, \Sigma \otimes I) \quad (19)$$

and given $Y_m = AQ_m$, the implied posterior mean on the solution space given by $A_m^{-1} \mathbf{b}$ is equivalent to the GMRES solution. This correspondence breaks when $\mathbf{x}_0 \neq \mathbf{0}$.

Proof Under this prior, \mathbf{b} applied to the posterior mean is

$$\begin{aligned} A_m^{-1} \mathbf{b} &= A_0^{-1} \mathbf{b} + (Q_m - A_m^{-1} Y_m)(Y_m^\top Y_m)^{-1} Y_m^\top \mathbf{b} \\ &= Q_m (Y_m^\top Y_m)^{-1} Y_m^\top \mathbf{b} \\ &= Q_m (Q_m^\top A^\top A Q_m)^{-1} Q_m^\top A^\top \mathbf{b} \end{aligned}$$

which is the GMRES projection step if $\mathbf{x}_0 = \mathbf{0}$. \square

6.4 Simulation Study

In this section the simulation study of Cockayne et al. [2018] will be replicated to demonstrate that the uncertainty produced from GMRES in Proposition 17 is similarly poorly calibrated, owing to the dependence of Q_m on \mathbf{x}^* by way of its dependence on \mathbf{b} . Throughout the size of the test problems is set to $d = 100$. The eigenvalues of A were drawn from an exponential distribution with parameter $\gamma = 10$, and eigenvectors uniformly from the Haar-measure over rotation-matrices (see Diaconis and Shahshahani [1987]). In contrast to Cockayne et al. [2018] the entries of \mathbf{b} are drawn from a standard Gaussian distribution, rather than \mathbf{x}_* . By Lemma 1, the prior is then perfectly calibrated for this scenario, providing justification for the expectation that the posterior should be equally well-calibrated for $m \geq 1$.

Figure 6.4 shows on the left the convergence of GMRES and on the right the convergence rate of the trace of the posterior covariance.

Figure 6.4 repeats the uncertainty quantification study of Cockayne et al. [2018]. Cockayne et al. [2018] argue that if the uncertainty is well-calibrated then \mathbf{x}^* can be considered as a draw from the posterior. Under this assumption, i.e. $\Sigma_m^{-1/2}(\mathbf{x}^* - \mathbf{x}_m) \sim \mathcal{N}(\mathbf{0}, I)$ they derive the test statistic:

$$Z(\mathbf{x}^*) := \|\Sigma_m^{-1/2}(\mathbf{x}^* - \mathbf{x}_m)\| \sim \chi_{d-m}^2.$$

It can be seen that the same poor uncertainty quantification occurs in BayesGMRES; even after just 10 iterations, the empirical distribution of the test statistic exhibits a profound left-shift, indicating an overly conservative posterior distribution. Producing well-calibrated posteriors remains an open issue in the field of probabilistic linear solvers.

7 Discussion

We have established many new connections between probabilistic linear solvers and a broad class of iterative methods. Matrix-based and solution-based inference were shown to be equivalent in a particular regime, showing that results from SBI transfer to MBI with




Fig. 1 Convergence of posterior mean and variance of the probabilistic interpretation of GMRES from Proposition 17.




Fig. 2 Assessment of the uncertainty quantification. Plotted are kernel density estimates for the statistic Z based on 500 randomly sampled test problems for steps $m = \{1, 3, 5, 8, 10\}$. These are compared with the theoretical distribution of Z when the posterior distribution is well-calibrated.

left-multiplied observations. Since SBI is a special case of MBI, future research will establish what additional benefits the increased generality of MBI can provide.

We also established a connection between the wide class of projection methods and probabilistic linear solvers. The common practise of preconditioning has an intuitive probabilistic interpretation, and all probabilistic linear solvers can be interpreted as projection methods. While the converse was shown to hold, the conditions under which generic projection methods can be reproduced are somewhat restrictive; however, GMRES and CG, which are among the most commonly used projection methods, have a well-defined probabilistic interpretation. Probabilistic interpretations of other widely used iterative methods can, we anticipate, be established from the results presented in this work.

Posterior uncertainty remains a challenge for probabilistic linear solvers. Direct probabilistic interpretations of CG and GMRES yield posterior covariance matrices which are not always computable, and even when the posterior can be computed the uncertainty remains poorly calibrated. This is owed to the dependence of the search directions in Krylov methods on $A\mathbf{x}^* = \mathbf{b}$, resulting in an algorithm which is not strictly Bayesian. Mitigating this issue without sacrificing the fast rate of convergence provided by Krylov methods remains an important focus for future work.

Acknowledgements

Ilse Ipsen was supported in part by NSF grant DMS-1760374. Mark Girolami was supported by EPSRC grants [EP/R034710/1, EP/R018413/1, EP/R004889/1, EP/P020720/1], an EPSRC Established Career Fellowship EP/J016934/3, a Royal Academy of Engineering Research Chair, and The Lloyds Register Foundation Programme on Data Centric Engineering. Philipp Hennig was supported by an ERC grant [757275/PANAMA].

Appendix A Properties of Kronecker Products

The following identities about Kronecker products and the vectorization operator are easily derived, but recalled here for the convenience of the reader:

$$(A \otimes B)\vec{C} = \overrightarrow{ACB^T} \quad (\text{K1})$$

$$(A \otimes B)(C \otimes D) = (AC) \otimes (BD) \quad (\text{K2})$$

$$(A \otimes B)^{-1} = A^{-1} \otimes B^{-1} \quad (\text{K3})$$

$$(A \otimes B)^T = A^T \otimes B^T \quad (\text{K4})$$

$$(A + B) \otimes C = A \otimes C + B \otimes C \quad (\text{K5})$$

A.1 The Symmetric Kronecker Product

Definition 20 (symmetric Kronecker-product)

The *symmetric* Kronecker-product for two square matrices $A, B \in \mathbb{R}^{N \times N}$ of equal size is defined as

$$A \otimes B := \Gamma(A \otimes B)\Gamma$$

where $[\Gamma]_{ij,kl} := 1/2\delta_{ik}\delta_{jl} + 1/2\delta_{il}\delta_{jk}$ satisfies

$$\Gamma\vec{C} = 1/2\vec{C} + 1/2\vec{C}^T$$

for all square-matrices $C \in \mathbb{R}^{N \times N}$.

Proposition 21 (Theorem 2.3 in Hennig [2015])

Let $W \in \mathbb{R}^{d \times d}$ be symmetric and positive definite. Assume a Gaussian prior of symmetric mean A_0^{-1} and covariance $W \otimes W$ on the elements of a symmetric matrix A^{-1} . After m linearly independent noise-free observations of the form $S = A^{-1}Y$, $Y \in \mathbb{R}^{d \times m}$, $\text{rk}(Y) = m$, the posterior belief over A^{-1} is a Gaussian with mean

$$\begin{aligned} A_m^{-1} &= A_0^{-1} + (S - A_0^{-1}Y)GY^T W \\ &\quad + WYG(S - A_0^{-1}Y)^T \\ &\quad + WYGY^T(S - A_0^{-1}Y)GY^T W \end{aligned} \quad (20)$$

and posterior covariance

$$V_m = (W - WYGY^T W) \otimes (W - WYGY^T W) \quad (21)$$

where $G := (Y^T W Y)^{-1}$.

Remark 22 Since A_0^{-1} is symmetric and the symmetric prior places mass only on symmetric matrices, the posterior mean A_m^{-1} is also symmetric.

Appendix B Proofs

B.1 Proposition 6

Proof (Proof of Proposition 6)

Let $H = A^{-1}$ and let $A_0^{-1} = H_0$. First note that by right-multiplying the information in Eq. (5) by H :

$$\begin{aligned} Y_m^T H &= S_m^T \\ \implies \overrightarrow{Y_m^T H} &= \overrightarrow{S_m^T} \\ \implies (Y_m \otimes I)\vec{H} &= \vec{S}_m^T \quad (\text{from K1}) \end{aligned}$$

Now the implied posterior on \vec{H} can be computed using the standard laws of Gaussian conditioning:

$$\begin{aligned} \vec{H} &\sim \mathcal{N}(\vec{H}_0, \Sigma \otimes W) \\ \implies \vec{H} | \mathcal{D} &\sim \mathcal{N}(\vec{H}_m, \Omega_m). \end{aligned}$$

Let $\Omega_0 = \Sigma_0 \otimes W_0$ and let $P = Y_m^\top \otimes I$. Then

$$\begin{aligned} \vec{H}_m &= \vec{H}_0 + [P\Omega_0]^\top [P\Omega_0 P^\top]^{-1} (\vec{S}_m^\top - \overline{Y_m^\top H_0}) \\ \Omega_m &= \Omega_0 - [P\Omega_0]^\top [P\Omega_0 P^\top]^{-1} (P\Omega_0) \end{aligned}$$

Now note that

$$\begin{aligned} P\Omega_0 &= (Y_m^\top \otimes I)(\Sigma_0 \otimes W) \\ &= (Y_m^\top \Sigma_0) \otimes W \\ \implies (P\Omega_0)^\top &= (\Sigma_0 Y_m) \otimes W \end{aligned}$$

where the second line uses Eq. (K2) and the third uses Eq. (K4). Thus

$$\begin{aligned} P\Omega_0 P^\top &= (Y_m^\top \otimes I)(\Sigma_0 \otimes W_0)(Y_m^\top \otimes I)^\top \\ &= (Y_m^\top \Sigma_0 Y_m) \otimes W_0 \\ \implies (P\Omega_0 P^\top)^{-1} &= (Y_m^\top \Sigma_0 Y_m)^{-1} \otimes W_0^{-1} \end{aligned}$$

where the second line is again using Eq. (K2) and Eq. (K4), while the third line uses Eq. (K3). We conclude that

$$\begin{aligned} (P\Omega_0)^\top (P\Omega_0 P^\top)^{-1} &= [(\Sigma_0 Y_m) \otimes W][(Y_m^\top \Sigma_0 Y_m)^{-1} \otimes W^{-1}] \\ &= (\Sigma_0 Y_m (Y_m^\top \Sigma_0 Y_m)^{-1}) \otimes I \\ \implies (P\Omega_0)^\top (P\Omega_0 P^\top)^{-1} (P\Omega_0) &= (\Sigma_0 Y_m (Y_m^\top \Sigma_0 Y_m)^{-1} Y_m^\top \Sigma_0) \otimes W_0. \end{aligned}$$

From these expressions it is straightforward to simplify the expressions for \vec{H}_m :

$$\begin{aligned} \vec{H}_m &= \vec{H}_0 + (\Sigma_0 Y_m (Y_m^\top \Sigma_0 Y_m)^{-1} \otimes I) (\vec{S}_m^\top - \overline{Y_m^\top H_0}) \\ &= \text{vec} (H_0 + \Sigma_0 Y_m (Y_m^\top \Sigma_0 Y_m)^{-1} (S_m^\top - Y_m^\top H_0)) \end{aligned}$$

where the last line follows from K1. For Ω_m :

$$\begin{aligned} \Omega_m &= \Sigma_0 \otimes W - (\Sigma_0 Y_m (Y_m^\top \Sigma_0 Y_m)^{-1} Y_m^\top \Sigma_0) \otimes W_0 \\ &= (\Sigma_0 - \Sigma_0 Y_m (Y_m^\top \Sigma_0 Y_m)^{-1} Y_m^\top \Sigma_0) \otimes W_0 \end{aligned}$$

where the last line is from application of K5.

It remains to project the posterior into \mathbb{R}^d by performing the matrix-vector product $H\mathbf{b}$.

$$\mathbf{x} = \vec{H}\mathbf{b} = (I \otimes \mathbf{b}^\top)H. \quad (\text{from K1})$$

Thus, the implied posterior is $\mathbf{x} \sim \mathcal{N}(\bar{\mathbf{x}}_m, \bar{\Sigma}_m)$, with

$$\begin{aligned} \bar{\mathbf{x}}_m &= (I \otimes \mathbf{b}^\top) \text{vec} (\\ &\quad H_0 + \Sigma_0 Y_m (Y_m^\top \Sigma_0 Y_m)^{-1} (S_m^\top - Y_m^\top H_0) \\ &= \text{vec} (H_0 \mathbf{b} + \Sigma_0 Y_m (Y_m^\top \Sigma_0 Y_m)^{-1} (S_m^\top \mathbf{b} - Y_m^\top H_0 \mathbf{b})) \\ &= \mathbf{x}_0 + \Sigma_0 A^\top S_m (S_m^\top A \Sigma_0 A^\top S_m)^{-1} S_m^\top (\mathbf{b} - A \mathbf{x}_0) \end{aligned}$$

where in the last line we have used that $H_0 \mathbf{b} = \mathbf{x}_0$ and that $Y_m = A^\top S_m$. Furthermore

$$\begin{aligned} \bar{\Sigma}_m &= (I \otimes \mathbf{b}^\top) \\ &\quad \cdot [(\Sigma_0 - \Sigma_0 Y_m (Y_m^\top \Sigma_0 Y_m)^{-1} Y_m^\top \Sigma_0) \otimes W_0] \\ &\quad \cdot (I \otimes \mathbf{b}^\top)^\top \\ &= (\Sigma_0 - \Sigma_0 Y_m (Y_m^\top \Sigma_0 Y_m)^{-1} Y_m^\top \Sigma_0) \times \mathbf{b}^\top W_0 \mathbf{b} \\ &= \Sigma_0 - \Sigma_0 A^\top S_m (S_m^\top A \Sigma_0 A^\top S_m)^{-1} S_m^\top A \Sigma_0 \end{aligned}$$

where in the second line we have used K2 and the fact that $\mathbf{b}^\top W_0 \mathbf{b}$ is a scalar, while in the third line we have used that $\mathbf{b}^\top W_0 \mathbf{b} = 1$ and that $Y_m = A^\top S_m$.

Note that $\mathbf{x}_m = \bar{\mathbf{x}}_m$ and $\Sigma_m = \bar{\Sigma}_m$, as defined in Cockayne et al. [2018]. Thus, the proof is complete. \square

B.2 Theorem 16

Proof (Proof of Theorem 16.) Denote by \mathbf{x}_i^{CG} the conjugate gradient estimate in iteration i and with \mathbf{p}_i the search direction in that iteration. From one iteration to the next, the update to the solution can be written as [Nocedal and Wright, 1999, p. 108]

$$\mathbf{x}_{i+1}^{CG} = \mathbf{x}_i^{CG} + \frac{\mathbf{r}_i^\top \mathbf{p}_i}{\mathbf{p}_i^\top A \mathbf{p}_i} \mathbf{p}_i. \quad (22)$$

Comparing this update to lines 7 to 10 in Algorithm 1 it is sufficient to show that $\mathbf{d}_i \propto \mathbf{p}_i$ which follows from Lemma 23. \square

Lemma 23 *Assume that CG does not terminate before d iterations. Using the prior of Theorem 16 in Algorithm 1, the directions \mathbf{d}_i are scaled conjugate gradients search directions, i.e.*

$$\mathbf{d}_i = \gamma_i \mathbf{p}_i^{CG}$$

where \mathbf{p}_i^{CG} is the CG search direction in iteration i and $\gamma_i \in \mathbb{R} \setminus \{0\}$.

Proof The proof proceeds by induction. Throughout we will suppress the superscript CG on the CG search directions, i.e. $\mathbf{p}_i^{CG} = \mathbf{p}_i$. For $i = 1$, $A_{i-1}^{-1} = \alpha I$ by assumption and therefore $\mathbf{d}_i = \alpha \mathbf{r}_0$ which is the first CG search direction scaled by $\gamma_1 = \alpha \neq 0$.

For the inductive step, suppose that the search directions $\mathbf{s}_1, \dots, \mathbf{s}_{i-1}$ are scaled CG directions and that

the vectors $\mathbf{x}_1, \dots, \mathbf{x}_{i-1}$ are the same as the first $i-1$ solution estimates produced by CG. We will prove that \mathbf{s}_i is the i^{th} CG search direction, and that \mathbf{x}_i is the i^{th} solution estimate from CG. Lemma 25 states that \mathbf{d}_i can be written as

$$\mathbf{d}_i = A_{i-1}^{-1} \mathbf{r}_{i-1} = \sum_{j < i} \nu_j \mathbf{s}_j + \nu_i \mathbf{r}_{i-1}. \quad (23)$$

where $\nu_j \in \mathbb{R}, j = 1, \dots, i$. Under the prior, the posterior mean A_i^{-1} is always symmetric as stated in Remark 22. This allows application of Lemma 24, so that $\{\mathbf{s}_1, \dots, \mathbf{s}_{i-1}, \mathbf{d}_i\}$ is an A -conjugate set. Thus we have, for $\ell < i$:

$$\begin{aligned} 0 &= \mathbf{s}_\ell^\top A \mathbf{d}_i = \nu_\ell \mathbf{s}_\ell^\top A \mathbf{s}_\ell + \nu_i \mathbf{s}_\ell^\top A \mathbf{r}_{i-1} \\ &= \nu_\ell \mathbf{s}_\ell^\top A \mathbf{s}_\ell + \nu_i \mathbf{y}_\ell^\top \mathbf{r}_{i-1}. \end{aligned} \quad (24)$$

Now note that

$$\mathbf{y}_\ell^\top \mathbf{r}_{i-1} = (\mathbf{r}_\ell - \mathbf{r}_{\ell-1})^\top \mathbf{r}_{i-1}.$$

This follows from Line 10 of Algorithm 1, from which it is clear that $\mathbf{y}_\ell = \mathbf{r}_\ell - \mathbf{r}_{\ell-1}$. Recall that the CG residuals \mathbf{r}_j are orthogonal [Nocedal and Wright, 1999, p. 109], and that from the inductive assumption, Algorithm 1 is equivalent to CG up to iteration $i-1$). Thus, for $\ell < i-1$ we have that

$$\begin{aligned} \mathbf{y}_\ell^\top \mathbf{r}_{i-1} &= 0 \\ \implies \mathbf{s}_\ell^\top A \mathbf{d}_i &= \nu_\ell \mathbf{s}_\ell^\top A \mathbf{s}_\ell = 0 \quad \forall \ell < i-1 \end{aligned}$$

where the second line is from application of the first line in Eq. (24). However, A is positive definite and by assumption the algorithm has not converged, so $\mathbf{d}_\ell \neq \mathbf{0}$. Furthermore clearly $\mathbf{s}_\ell^\top A \mathbf{s}_\ell \neq 0$. Hence we must have that

$$\nu_\ell = 0 \quad \forall j < i-1.$$

Equation (23) thus simplifies to

$$\mathbf{d}_i = \nu_{i-1} \mathbf{s}_{i-1} + \nu_i \mathbf{r}_{i-1} = \nu_{i-1} \alpha_{i-1} \mathbf{d}_{i-1} + \nu_i \mathbf{r}_{i-1}. \quad (25)$$

Now, again by Lemma 24, \mathbf{d}_i must be conjugate to \mathbf{s}_{i-1} which implies $\nu_i \neq 0$. Pre-multiplying Eq. (25) by $\mathbf{s}_{i-1}^\top A$ gives

$$\begin{aligned} 0 &= \nu_{i-1} \alpha_{i-1} \mathbf{s}_{i-1}^\top A \mathbf{d}_{i-1} + \nu_i \mathbf{s}_{i-1}^\top A \mathbf{r}_{i-1} \\ \implies \nu_{i-1} \alpha_{i-1} &= -\nu_i \frac{\mathbf{s}_{i-1}^\top A \mathbf{r}_{i-1}}{\mathbf{s}_{i-1}^\top A \mathbf{d}_{i-1}}. \end{aligned}$$

Thus, \mathbf{d}_i can be written as

$$\begin{aligned} \mathbf{d}_i &= \nu_i \left(\mathbf{r}_{i-1} - \frac{\mathbf{s}_{i-1}^\top A \mathbf{r}_{i-1}}{\mathbf{s}_{i-1}^\top A \mathbf{d}_{i-1}} \mathbf{d}_{i-1} \right) \\ &= \nu_i \left(\mathbf{r}_{i-1} - \frac{\mathbf{p}_{i-1}^\top A \mathbf{r}_{i-1}}{\mathbf{p}_{i-1}^\top A \mathbf{p}_{i-1}} \mathbf{p}_{i-1} \right) \end{aligned} \quad (26)$$

where the second line again applies the inductive assumption, that \mathbf{d}_{i-1} and \mathbf{s}_{i-1} are proportional to the CG search direction \mathbf{p}_{i-1} , noting that the proportionality constants on numerator and denominator cancel. The term inside the brackets is precisely the i^{th} CG search direction. This completes the result. \square

Lemma 24 *If the belief over A_m^{-1} is symmetric for all $m = 0, \dots, d$ and A is symmetric and positive definite, then Algorithm 1 produces A -conjugate directions.*

Proof The proof is by induction. Note that the case $i = 1$ is irrelevant since a set consisting of one element is trivially A -conjugate. On many occasions the proof relies on the consistency of the MBI belief, i.e. $A_i^{-1} \mathbf{z}_k = \mathbf{d}_k$ for $k \leq i$ and by symmetry $\mathbf{z}_k^\top A_i^{-1} = \mathbf{d}_k^\top$. Thus, for the base case $i = 2$ we have:

$$\begin{aligned} \mathbf{d}_1^\top A \mathbf{d}_2 &= -\mathbf{d}_1^\top A (A_1^{-1} \mathbf{r}_1) \\ &= -\mathbf{d}_1^\top A (A_1^{-1} (\mathbf{y}_1 + \mathbf{r}_0)) \\ &= -\mathbf{d}_1^\top A (\mathbf{s}_1 + A_1^{-1} \mathbf{r}_0) \end{aligned}$$

where the second line is by Line 10 of Algorithm 1. Now recall that $\alpha_1 = -\mathbf{d}_1^\top \mathbf{r}_0 / \mathbf{d}_1^\top A \mathbf{d}_1$ to give:

$$\begin{aligned} \mathbf{d}_1^\top A \mathbf{d}_2 &= -\alpha_1 \mathbf{d}_1^\top A \mathbf{d}_1 - \mathbf{d}_1^\top A A_1^{-1} \mathbf{r}_0 \\ &= \mathbf{d}_1^\top \mathbf{r}_0 - \mathbf{d}_1^\top A A_1^{-1} \mathbf{r}_0 \\ &= \mathbf{d}_1^\top \mathbf{r}_0 - \mathbf{z}_1^\top A_1^{-1} \mathbf{r}_0 \\ &= \mathbf{d}_1^\top \mathbf{r}_0 - \mathbf{d}_1^\top \mathbf{r}_0 \\ &= 0. \end{aligned} \quad (27)$$

Here, the symmetry of the estimator A_i^{-1} is used in Eq. (27). For the inductive step, assume $\{\mathbf{d}_0, \dots, \mathbf{d}_{i-1}\}$ are pairwise A -conjugate. For any $k < i$ we have:

$$\begin{aligned} \mathbf{d}_k^\top A \mathbf{d}_i &= -\mathbf{d}_k^\top A (A_i^{-1} \mathbf{r}_i) \\ &= -\mathbf{d}_k^\top A A_i^{-1} \left(\sum_{j \leq i} \mathbf{y}_j + \mathbf{r}_0 \right) \end{aligned}$$

where the second line follows from the fact that $\mathbf{r}_i = \mathbf{r}_{i-1} + \mathbf{y}_i$. Thus, we have:

$$\begin{aligned} \mathbf{d}_k^\top A \mathbf{d}_i &= -\mathbf{d}_k^\top A \left(\sum_{j \leq i} \mathbf{s}_j + A_i^{-1} \mathbf{r}_0 \right) \\ &= -\mathbf{d}_k^\top A \left(\sum_{j \leq i} \alpha_j \mathbf{d}_j + A_i^{-1} \mathbf{r}_0 \right). \end{aligned}$$

Now, applying the conjugacy from the inductive assumption:

$$\begin{aligned} \mathbf{d}_k^\top \mathbf{A} \mathbf{d}_i &= -\alpha_k \mathbf{d}_k^\top \mathbf{A} \mathbf{d}_k - \mathbf{d}_k^\top \mathbf{A} (\mathbf{A}_i^{-1} \mathbf{r}_0) \\ &= \mathbf{d}_k^\top \mathbf{r}_{k-1} - \mathbf{d}_k^\top \mathbf{r}_0 \\ &= \mathbf{d}_k^\top \left(\sum_{j < k} \mathbf{y}_j + \mathbf{r}_0 \right) - \mathbf{d}_k^\top \mathbf{r}_0 = 0 \\ &= \sum_{j < k} \alpha_j \mathbf{d}_k^\top \mathbf{A} \mathbf{d}_j = 0. \end{aligned}$$

where the second line rearranges line 6 of the algorithm to obtain $\alpha_i \mathbf{d}_i^\top \mathbf{z}_i = -\mathbf{d}_i^\top \mathbf{r}_{i-1}$. The third line again uses that $\mathbf{r}_i = \mathbf{r}_{i-1} + \mathbf{y}_i$, while the fourth line is from the assumed conjugacy. \square

Lemma 25 *Under the prior in Theorem 16 and given scaled CG search directions $\mathbf{p}_1, \dots, \mathbf{p}_i$, it holds that $\mathbf{A}_i^{-1} \mathbf{r}_i \in \text{span}\{\mathbf{p}_1, \dots, \mathbf{p}_i, \mathbf{r}_i\}$.*

Proof Recall first that under the prior in Theorem 16, $\mathbf{A}_0^{-1} = \alpha \mathbf{I}$. Then by inspection of Eq. (20) we have $\mathbf{A}_i^{-1} \mathbf{r}_i \in \mathcal{S}$ where

$$\mathcal{S} = \text{span}\{\mathbf{r}_i, \mathbf{p}_1, \dots, \mathbf{p}_i, \mathbf{y}_1, \dots, \mathbf{y}_i, \mathbf{W} \mathbf{y}_1, \dots, \mathbf{W} \mathbf{y}_i\}$$

By choice of $\mathbf{W} = \beta \mathbf{I} + \gamma \mathbf{A}^{-1}$, $\mathcal{S} = \text{span}\{\mathbf{r}_i, \mathbf{p}_1, \dots, \mathbf{p}_i, \mathbf{y}_1, \dots, \mathbf{y}_i\}$. From line 10 of Algorithm 1 $\mathbf{y}_i = \mathbf{r}_i - \mathbf{r}_{i-1}$ and therefore $\mathcal{S} = \text{span}\{\mathbf{r}_1, \dots, \mathbf{r}_i, \mathbf{p}_1, \dots, \mathbf{p}_i\}$. By Theorem 5.3 in [Nocedal and Wright, 1999, p. 109] the span of the conjugate gradients residuals and search directions are equivalent. Therefore $\mathcal{S} \subseteq \{\mathbf{r}_i, \mathbf{p}_1, \dots, \mathbf{p}_i\}$. \square

References

- S. Bartels and P. Hennig. Probabilistic approximate least-squares. In *Proceedings of Artificial Intelligence and Statistics (AISTATS)*, 2016.
- J. Cockayne, C. Oates, T. Sullivan, and M. Girolami. Probabilistic numerical methods for partial differential equations and bayesian inverse problems, 2016.
- J. Cockayne, C. Oates, T. Sullivan, and M. Girolami. Bayesian probabilistic numerical methods, 2017.
- J. Cockayne, C. Oates, and M. Girolami. A bayesian conjugate gradient method, 2018.
- P. Diaconis and M. Shahshahani. The subgroup algorithm for generating uniform random variables. *Probability in the Engineering and Informational Sciences*, 1(01):15, jan 1987. doi: 10.1017/s0269964800000255.
- G. H. Golub and C. F. Van Loan. *Matrix computations*. Johns Hopkins Studies in the Mathematical Sciences. Johns Hopkins University Press, Baltimore, MD, fourth edition, 2013.
- P. Hennig. Probabilistic interpretation of linear solvers. *SIAM Journal on Optimization*, 25(1):234–260, jan 2015. doi: 10.1137/140955501. URL <https://doi.org/10.1137/140955501>.
- P. Hennig, M. A. Osborne, and M. Girolami. Probabilistic numerics and uncertainty in computations. *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 2015.
- T. Karvonen and S. Sarkka. Classical quadrature rules via gaussian processes. In *2017 IEEE 27th International Workshop on Machine Learning for Signal Processing (MLSP)*. IEEE, sep 2017. doi: 10.1109/mlsp.2017.8168195.
- H. Kersting, T. J. Sullivan, and P. Hennig. Convergence Rates of Gaussian ODE Filters. *ArXiv e-prints*, 1807.09737, 7 2018.
- J. Liesen and Z. Strakos. *Krylov Subspace Methods. Principles and Analysis*. Oxford University Press, Oct. 2012. doi: 10.1093/acprof:oso/9780199655410.001.0001.
- J. Nocedal and S. J. Wright. *Numerical Optimization*. Springer Verlag, 1999.
- Y. Saad. *Iterative Methods for Sparse Linear Systems*. Society for Industrial and Applied Mathematics, second edition, 2003.
- Y. Saad and M. H. Schultz. GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM Journal on Scientific and Statistical Computing*, 7(3):856–869, jul 1986. doi: 10.1137/0907058. URL <https://doi.org/10.1137/0907058>.
- M. Schober, D. Duvenaud, and P. Hennig. Probabilistic ODE solvers with runge-kutta means. In *Advances in Neural Information Processing Systems 27*, pages 739–747. Curran Associates, Inc., 2014. URL <http://papers.nips.cc/paper/5451-probabilistic-ode-solvers-with-runge-kutta-means.pdf>.
- M. Schober, S. Särkkä, and P. Hennig. A probabilistic model for the numerical solution of initial value problems. *Statistics and Computing*, 2018.
- K. M. Soodhalter, D. B. Szyld, and F. Xue. Krylov subspace recycling for sequences of shifted linear systems. *Applied Numerical Mathematics*, 81:105–118, jul 2014. doi: 10.1016/j.apnum.2014.02.006. URL <https://doi.org/10.1016/j.apnum.2014.02.006>.
- X. Xi, F.-X. Briol, and M. Girolami. Bayesian quadrature for multiple related integrals. In *Proceedings of the 35th International Conference on Machine Learning (ICML)*, 2018. arXiv:801.04153.