SUPPLEMENTARY MATERIALS: BAYESCG AS AN UNCERTAINTY AWARE VERSION OF CG*

TIM W. REID†, ILSE C. F. IPSEN†, JON COCKAYNE‡, AND CHRIS J. OATES§

SM1. Outline of Supplementary Materials. We present the proof of Theorem 2.1 (section SM2), discuss more theoretical properties of BayesCG (section SM3), and examine the performance of the Krylov posterior as a CG error estimate (section SM4).

SM2. Proof of Theorem 2.1. We present an example of search directions that satisfy the assumptions of Theorem 2.1 (Example SM2.1); review the conjugacy and stability of Gaussian distributions (Lemmas SM2.2 and SM2.3); present the proof of Theorem 2.1; and discuss the relation between the solution $\mathbf{x}_*$ and the random variable $X \sim \mathcal{N}(\mathbf{x}_0, \mathbf{\Sigma}_0)$ (Remark SM2.4).

Existence of search directions satisfying the assumptions of Theorem 2.1. The example below illustrates a non-recursive way to select search directions $\mathbf{S}_m$ so that $\mathbf{\Lambda}_m = \mathbf{S}_m^T\mathbf{A}\mathbf{\Sigma}_0\mathbf{A}\mathbf{S}_m$ is nonsingular. The purpose of this example is to show that at for all $m \leq \operatorname{rank}(\mathbf{\Sigma}_0)$, least one set of search directions $\mathbf{S}_m$ exists that satisfies the assumptions of Theorem 2.1.

EXAMPLE SM2.1. *Let $\mathbf{\Sigma}_0 = \mathbf{U}\mathbf{D}\mathbf{U}^T$ be a singular value decomposition of the prior covariance $\mathbf{\Sigma}_0$, and let $m \leq \operatorname{rank}(\mathbf{\Sigma}_0)$. Distinguish the leading m columns of $\mathbf{U}$, and the leading nonsingular $m \times m$ principal submatrix of $\mathbf{D}$*

$$\mathbf{U}_{1:m} \equiv \begin{bmatrix}\mathbf{u}_1 & \mathbf{u}_2 & \cdots & \mathbf{u}_m\end{bmatrix} \qquad \textit{and} \qquad \mathbf{D}_{1:m} \equiv \operatorname{diag}\begin{pmatrix}d_1 & d_2 & \cdots & d_m\end{pmatrix},$$

and define the search directions $\mathbf{S}_m \equiv \mathbf{A}^{-1}\mathbf{U}_m$. Then the equality

$$\mathbf{\Lambda}_m = \mathbf{S}_m^T\mathbf{A}\mathbf{\Sigma}_0\mathbf{A}\mathbf{S}_m = \mathbf{U}_m^T\mathbf{A}^{-1}\mathbf{A}\mathbf{\Sigma}_0\mathbf{A}\mathbf{A}^{-1}\mathbf{U}_m = \mathbf{U}_m^T\mathbf{\Sigma}_0\mathbf{U}_m = \mathbf{D}_m$$

and $m \leq \operatorname{rank}(\mathbf{\Sigma}_m)$ imply that $\mathbf{D}_m$, hence $\mathbf{\Lambda}_m$, is nonsingular.

This example shows that at least one set of search directions exists that satisfying the assumptions of Theorem 2.1. This example is necessary because Theorem 2.11 only shows that the recursively computed search directions from Theorem 2.8 satisfy the assumptions of Theorem 2.1 for $m \leq \mathrm{K} \leq \operatorname{rank}(\mathbf{\Sigma}_0)$, where K is the grade of $\mathbf{r}_0$ with respect to $\mathbf{A}\mathbf{\Sigma}_0\mathbf{A}$. In practice, it is best to compute the BayesCG posterior with the recursively computed search directions, even if $\mathrm{K} < \operatorname{rank}(\mathbf{\Sigma}_m)$. There is no reason to compute more than K of these search directions because they cause the posterior mean at K iterations to be $\mathbf{x}_\mathrm{K} = \mathbf{x}_*$.

---

*Submitted to the editors

Funding: The work was supported in part by NSF grant DMS-1745654 (TWR, ICFI), NSF grant DMS-1760374 and DOE grant DE-SC0022085 (ICFI), and the Lloyd's Register Foundation Programme on Data Centric Engineering at the Alan Turing Institute (CJO).

†Department of Mathematics, North Carolina State University, Raleigh, NC 27695-8205, USA, (twreid@alumni.ncsu.edu, ipsen@ncsu.edu)

‡Department of Mathematical Sciences, University of Southampton, Southampton SO17 1BJ, UK (jon.cockayne@soton.ac.uk)

§School of Mathematics, Statistics and Physics, Newcastle University, Newcastle upon Tyne NE1 7RU, UK (chris.oates@ncl.ac.uk)

Review of stability and conjugacy of Gaussian distributions. The proof of Theorem 2.1 relies on the *stability* and *conjugacy* of Gaussian distributions.

LEMMA SM2.2 (Stability of Gaussian distributions [SM14, Section 1.2]). *Let $X \sim \mathcal{N}(\mathbf{x}, \mathbf{\Sigma}) \in \mathbb{R}^n$ be a Gaussian random variable with mean $\mathbf{x} \in \mathbb{R}^n$ and covariance $\mathbf{\Sigma} \in \mathbb{R}^{n\times n}$. If $\mathbf{y} \in \mathbb{R}^n$ is a vector and $\mathbf{F} \in \mathbb{R}^{n\times n}$ is a matrix, then $Z = \mathbf{y} + \mathbf{F}X$ is again a Gaussian random variable distibuted as*

$$Z \sim \mathcal{N}(\mathbf{y} + \mathbf{F}\mathbf{x}, \mathbf{F}\mathbf{\Sigma}\mathbf{F}^T).$$

LEMMA SM2.3 (Conjugacy of Gaussian distributions [SM15, Section 6.1], [SM19, Corollary 6.21]). *Let $X \sim \mathcal{N}(\mathbf{x}, \mathbf{\Sigma}_x)$ and $Y \sim \mathcal{N}(\mathbf{y}, \mathbf{\Sigma}_y)$ be Gaussian random variables. The jointly Gaussian random variable $\begin{bmatrix} X^T & Y^T \end{bmatrix}^T$ has the distribution*

$$\begin{bmatrix} X \\ Y \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} \mathbf{x} \\ \mathbf{y} \end{bmatrix}, \begin{bmatrix} \mathbf{\Sigma}_x & \mathbf{\Sigma}_{xy} \\ \mathbf{\Sigma}_{xy}^T & \mathbf{\Sigma}_y \end{bmatrix} \right),$$

where $\mathbf{\Sigma}_{xy} \equiv \mathrm{Cov}(X, Y) = \mathbb{E}[(X - \mathbf{x})(Y - \mathbf{y})^T]$ and the conditional distribution of X given Y is

$$(X \mid Y) \sim \mathcal{N}(\overbrace{\mathbf{x} + \mathbf{\Sigma}_{xy}\mathbf{\Sigma}_y^\dagger(Y - \mathbf{y})}^{mean}, \overbrace{\mathbf{\Sigma}_x - \mathbf{\Sigma}_{xy}\mathbf{\Sigma}_y^\dagger\mathbf{\Sigma}_{xy}^T}^{covariance}).$$

Proof of Theorem 2.1. Since $m \le \mathrm{rank}(\mathbf{\Sigma}_0)$, we can choose search directions $\mathbf{S}_m$ with linearly independent columns so that $\mathbf{\Lambda}_m$ is nonsingular, see Example SM2.1. Then the proof reduces to that of [SM5, Proof of Proposition 1].

Let the random variable $X_0 \sim \mathcal{N}(\mathbf{x}_0, \mathbf{\Sigma}_0)$ represent the prior belief about the unknown solution $\mathbf{x}_*$, and let the random variable $Y_m \equiv \mathbf{S}_m^T\mathbf{A}X_0$ represent the implied prior belief about the unknown values $\mathbf{S}_m^T\mathbf{A}\mathbf{x}_*$ before they are computed. The posterior is the conditional distribution $(X_0 \mid Y_m = \mathbf{S}_m^T\mathbf{A}\mathbf{x}_*)$ [SM3, Proposition 1]. Thus, we first determine the conditional distribution $(X_0 \mid Y_m)$ and then substitute $Y_m = \mathbf{S}_m^T\mathbf{A}\mathbf{x}_*$.

The joint distribution of X_0 and Y_m is

$$\begin{bmatrix} X_0 \\ Y_m \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} \mathbf{x}_0 \\ \mathbb{E}[Y_m] \end{bmatrix}, \begin{bmatrix} \mathbf{\Sigma}_0 & \mathrm{Cov}(X_0, Y_m) \\ \mathrm{Cov}(X_0, Y_m)^T & \mathrm{Cov}(Y_m, Y_m) \end{bmatrix} \right). \tag{SM2.1}$$

The mean and covariance of Y_m follow from Lemma SM2.2,

$$\mathbb{E}[Y_m] = \mathbf{S}_m^T\mathbf{A}\mathbf{x}_0 \qquad \text{and} \qquad \mathrm{Cov}(Y_m, Y_m) = \mathbf{S}_m^T\mathbf{A}\mathbf{\Sigma}_0\mathbf{A}\mathbf{S}_m = \mathbf{\Lambda}_m,$$

while the linearity of the expectation implies for the covariance that

$$\mathrm{Cov}(X_0, Y_m) = \mathbb{E}[(X_0 - \mathbf{x}_0)(\mathbf{S}_m^T\mathbf{A}(X_0 - \mathbf{x}_0))^T] = \mathbf{\Sigma}_0\mathbf{A}\mathbf{S}_m.$$

Substituting all of the above into (SM2.1) gives

$$\begin{bmatrix} X_0 \\ Y_m \end{bmatrix} \sim \mathcal{N}\left(\begin{bmatrix} \mathbf{x}_0 \\ \mathbf{S}_m^T\mathbf{A}\mathbf{x}_0 \end{bmatrix}, \begin{bmatrix} \mathbf{\Sigma}_0 & \mathbf{\Sigma}_0\mathbf{A}\mathbf{S}_m \\ (\mathbf{\Sigma}_0\mathbf{A}\mathbf{S}_m)^T & \mathbf{\Lambda}_m \end{bmatrix} \right).$$

Thus we can invoke Lemma SM2.3 to conclude that the conditional distribution for $(X_0 \mid Y_m)$ is a Gaussian $\mathcal{N}(\mathbf{x}_m, \mathbf{\Sigma}_m)$ with mean and covariance

$$\mathbf{x}_m = \mathbf{x}_0 + \mathbf{\Sigma}_0\mathbf{A}\mathbf{S}_m\mathbf{\Lambda}_m^{-1}(Y_m - \mathbf{S}_m^T\mathbf{A}\mathbf{x}_0)$$
$$\mathbf{\Sigma}_m = \mathbf{\Sigma}_0 - \mathbf{\Sigma}_0\mathbf{A}\mathbf{S}_m\mathbf{\Lambda}_m^{-1}\mathbf{S}_m^T\mathbf{A}\mathbf{\Sigma}_0.$$

At last, substitute $Y_m = \mathbf{S}_m^T \mathbf{A}\mathbf{x}_* = \mathbf{S}_m^T \mathbf{b}$ to obtain $(X_0 \mid Y_m = \mathbf{S}_m^T \mathbf{A}\mathbf{x}_*)$. ∎

Below we discuss the relation between the solution vector $\mathbf{x}_*$ and the random variable X from the proof of Theorem 2.1

Remark SM2.4. *The solution vector $\mathbf{x}_*$ is a deterministic value, but we do not know its true value. The prior distribution $\mathcal{N}(\mathbf{x}_0, \mathbf{\Sigma}_0)$ models the initial epistemic uncertainty in $\mathbf{x}_*$, that is, the uncertainty in our knowledge of the true value of $\mathbf{x}_*$. The random variable $X \sim \mathcal{N}(\mathbf{x}_0, \mathbf{\Sigma}_0)$ in the proof of Theorem 2.1 is a surrogate for $\mathbf{x}_*$.*

When we compute $\mathbf{x}_*$ with an iterative linear solver, we gain more information about the true value of $\mathbf{x}_*$. Since we gain information about $\mathbf{x}_*$, we can update the surrogate for $\mathbf{x}_*$ by conditioning X on the new information. In BayesCG, the information we gain is that $Y \equiv \mathbf{S}_m^T \mathbf{A} X$ takes the value $\mathbf{S}_m^T \mathbf{b}$. Therefore, our updated surrogate is $X \mid Y = \mathbf{S}_m^T \mathbf{b}$, and it is distributed according to the posterior distribution $\mathcal{N}(\mathbf{x}_m, \mathbf{\Sigma}_m)$. The posterior distribution models the uncertainty remaining $\mathbf{x}_*$ after we obtained the additional information about it.

SM3. Additional Theoretical Properties of BayesCG.

We discuss the relationship between BayesCG and CG (section SM3.1), present an alternative proof of Theorem 3.3 (section SM3.2), and present an alternative definition of $\mathbf{\Phi}$ that has the same convergence properties as in section 3.3 (section SM3.3).

SM3.1. Relationship Between BayesCG and CG.

We discuss the relationship between BayesCG and CG.

The posterior mean from Algorithm 2.1 is closely related to the approximate solution from CG. For *nonsingular* $\mathbf{\Sigma}_0$, BayesCG can be interpreted as CG applied to a right-preconditioned linear system. Specifically, [SM8] showed the posterior means $\mathbf{x}_i$ in Algorithm 2.1 are equal to the iterates of Algorithm 2.2 applied to the right preconditioned system

$$\mathbf{A}\,(\mathbf{\Sigma}_0 \mathbf{A})\mathbf{w}_* = \mathbf{b} \quad \text{where} \quad \mathbf{w}_* = (\mathbf{\Sigma}_0 \mathbf{A})^{-1}\,\mathbf{x}_*. \tag{SM3.1}$$

It can be seen in (SM3.1) that if $\mathbf{\Sigma}_0 = \mathbf{A}^{-1}$, then the BayesCG posterior mean is equal to the approximate solution computed by CG. This was originally shown in [SM3, Section 2.3] and can also be seen by comparing Algorithms 2.1 and 2.2. Additionally, if $\mathbf{\Sigma}_0 = \mathbf{A}^{-1}$, then the search directions in Algorithms 2.1 and 2.2 are equal as well.

Similarly to CG, the termination criterion in Algorithm 2.1 can be the usual relative residual norm, or it can be statistically motivated [SM2, Section 2], [SM4, Section 1.3].

The similarity of BayesCG (Algorithm 2.1) and CG (Algorithm 2.2) strongly suggests that both algorithms have similar finite precision behavior. The search directions in Algorithm 2.1 lose orthogonality through the course of the iteration, thereby slowing down the convergence of the posterior means [SM3, Section 6.1], similar to what happens in CG [SM9, Section 5.8], [SM11, Section 5]. In addition, loss of orthogonality causes loss of semi-definiteness in the posterior covariances $\mathbf{\Sigma}_m$, prohibiting their interpretation as covariance matrices since covariance matrices must be positive semi-definite [SM3, Section 6.1]. The remedy recommended in [SM3, Section 6.1] is reorthogonalization of the search directions.

SM3.2. Alternative Version of Theorem 3.3.

We present an alternative version of Theorem 3.3, the theorem that shows the Krylov posterior means are equal to CG iterates. This version additionally shows the search directions computed in Algorithm 2.1 under the Krylov prior are equal to the search directions in Algorithm 2.2.

The alternative version of Theorem 3.3 also verifies the claim in Remark 3.9 that the approximate Krylov posterior (3.20) can be viewed as as the posterior from the rank-$(m+d)$, $1 \leq d \leq \mathrm{K}-m$, approximation of the prior $\mathcal{N}(\mathbf{x}_0, \widehat{\mathbf{\Gamma}}_0)$ with

$$\widehat{\mathbf{\Gamma}}_0 = \mathbf{V}_{1:m+d}\mathbf{\Phi}_{1:m+d}(\mathbf{V}_{1:m+d})^T. \tag{SM3.2}$$

Similarly to Theorem 3.3, the alternative version of the theorem relies on (3.4). Equation (3.4) remains true for the approximate posterior:

$$\widehat{\mathbf{\Gamma}}_0\mathbf{A}\tilde{\mathbf{v}}_i = \phi_i\tilde{\mathbf{v}}_i, \qquad 1 \leq i \leq m+d. \tag{SM3.3}$$

THEOREM SM3.1. *Let $\mathbf{s}_i$ and $\mathbf{x}_i$, $1 \leq i \leq m$ be the search directions and posterior means computed in m iterations of Algorithm 2.1 starting from the prior $\mathcal{N}(\mathbf{x}_0, \widehat{\mathbf{\Gamma}}_0)$. Similarly, let $\mathbf{v}_i$ and $\mathbf{z}_i$, $1 \leq i \leq m$ be the search directions and solution iterates computed in m iterations of Algorithm 2.2 staring at initial guess $\mathbf{z}_0$. If $\mathbf{z}_0 = \mathbf{x}_0$, then*

$$\mathbf{s}_i = \mathbf{v}_i \quad \textit{and} \quad \mathbf{x}_i = \mathbf{z}_i, \qquad 1 \leq i \leq m. \tag{SM3.4}$$

Proof. We give an induction proof to establish the equality of iterates and search directions. In this proof we denote by

$$\mathbf{q}_i = \mathbf{b} - \mathbf{A}\mathbf{z}_i, \qquad 0 \leq i \leq m,$$

the residuals in Algorithm 2.2.

Induction base: The equality of the initial iterates follows from the assumption that $\mathbf{z}_0 = \mathbf{x}_0$. This, in turn, implies the equality of the corresponding residuals and search directions,

$$\mathbf{s}_1 = \mathbf{r}_0 = \mathbf{b} - \mathbf{A}\mathbf{x}_0 = \mathbf{b} - \mathbf{A}\mathbf{z}_0 = \mathbf{q}_0 = \mathbf{v}_1.$$

Induction hypothesis: Assume equality of the first m search directions and iterates,

$$\mathbf{x}_i = \mathbf{z}_i, \quad 0 \leq i \leq m-1, \qquad \text{and} \qquad \mathbf{s}_i = \mathbf{v}_i, \quad 1 \leq i \leq m. \tag{SM3.5}$$

The equality of the iterates implies the equality of the residuals

$$\mathbf{r}_i = \mathbf{b} - \mathbf{A}\mathbf{x}_i = \mathbf{b} - \mathbf{A}\mathbf{z}_i = \mathbf{q}_i, \qquad 0 \leq i \leq m-1. \tag{SM3.6}$$

Induction step: Show $\mathbf{x}_m = \mathbf{z}_m$ and $\mathbf{s}_{m+1} = \mathbf{v}_{m+1}$ via the recursions from Algorithms 2.1 and 2.2.

Iterates. Apply $\mathbf{z}_{m-1} = \mathbf{x}_{m-1}$ from (SM3.5) and $\mathbf{q}_{m-1} = \mathbf{r}_{m-1}$ from (SM3.6) to the iterate from Algorithm 2.2,

$$\mathbf{z}_m = \mathbf{z}_{m-1} + \frac{\mathbf{q}_{m-1}^T\mathbf{q}_{m-1}}{\mathbf{v}_m^T\mathbf{A}\mathbf{v}_m}\mathbf{v}_m = \mathbf{x}_{m-1} + \frac{\mathbf{r}_{m-1}^T\mathbf{r}_{m-1}}{\mathbf{v}_m^T\mathbf{A}\mathbf{v}_m}\mathbf{v}_m.$$

Apply $\mathbf{s}_m = \mathbf{v}_m$ from (SM3.5) the iterate from Algorithm 2.1 and simplify with (SM3.3),

$$\mathbf{x}_m = \mathbf{x}_{m-1} + \frac{\mathbf{r}_{m-1}^T\mathbf{r}_{m-1}}{\mathbf{s}_m^T\mathbf{A}\mathbf{\Gamma}_0\mathbf{A}\mathbf{s}_m}\mathbf{\Gamma}_0\mathbf{A}\mathbf{s}_m = \mathbf{x}_{m-1} + \frac{\phi_m}{\phi_m}\frac{\mathbf{r}_{m-1}^T\mathbf{r}_{m-1}}{\mathbf{v}_m^T\mathbf{A}\mathbf{v}_m}\mathbf{v}_m = \mathbf{z}_m,$$

which proves the equality of the iterates, and implies equality of the residuals $\mathbf{r}_m = \mathbf{q}_m$.

Search Directions. Apply $\mathbf{s}_m = \mathbf{v}_m$ from (SM3.5), and $\mathbf{r}_m = \mathbf{q}_m$ to the search direction from Algorithm 2.2,

$$\mathbf{s}_{m+1} = \mathbf{r}_m + \frac{\mathbf{r}_m^T \mathbf{r}_m}{\mathbf{r}_{m-1}^T \mathbf{r}_{m-1}} \mathbf{s}_m = \mathbf{q}_m + \frac{\mathbf{q}_m^T \mathbf{q}_m}{\mathbf{q}_{m-1}^T \mathbf{q}_{m-1}} \mathbf{v}_m = \mathbf{v}_{m+1},$$

which proves the equality of the search directions. □

Showing that the posterior covariance under the approximate Krylov prior is

$$\widehat{\mathbf{\Gamma}}_m = \mathbf{V}_{m+1:m+d} \mathbf{\Phi}_{m+1:m+d} (\mathbf{V}_{m+1:m+d})^T$$

follows the same argument as in Theorem 3.3.

Theorem SM3.1 shows that the search directions under the approximate Krylov prior are not in $\ker(\widehat{\mathbf{\Gamma}}_0 \mathbf{A})$. This is important to show because the approximate Krylov posterior does not satisfy the condition $\mathbf{x}_* - \mathbf{x}_0 \in \operatorname{range}(\widehat{\mathbf{\Gamma}}_0)$ from Theorem 2.11 which guarantees $\mathbf{s}_i \notin \ker(\widehat{\mathbf{\Gamma}}_0 \mathbf{A})$.

SM3.3. Alternative Definition of $\mathbf{\Phi}$. In the following theorem, we discuss a definition of $\mathbf{\Phi}$ that is equivalent to the definition in Theorem 3.7.

THEOREM SM3.2. *The diagonal elements of $\mathbf{\Phi}$ in Theorem 3.7 are equal to*

$$\phi_i = (\tilde{\mathbf{v}}_i^T \mathbf{r}_0)^2 = (\tilde{\mathbf{v}}_i^T \mathbf{A}(\mathbf{x}_* - \mathbf{x}_0))^2, \qquad 1 \le i \le K. \tag{SM3.7}$$

Proof. From Theorem 3.7, we have that $\phi_i = \gamma_i \|\mathbf{r}_{i-1}\|_2^2$, $1 \le i \le \mathrm{K}$. Substituting $\gamma_i = \mathbf{r}_{i-1}^T \mathbf{r}_{i-1} / \left(\mathbf{v}_i^T \mathbf{A} \mathbf{v}_i\right)$ from Algorithm 3.1 into ϕ_i results in

$$\phi_i = \frac{\|\mathbf{r}_{i-1}\|_2^4}{\mathbf{v}_i^T \mathbf{A} \mathbf{v}_i}, \qquad 1 \le i \le \mathrm{K}.$$

From the previous equation and $\mathbf{v}_i^T \mathbf{r}_{i-1} = \|\mathbf{r}_{i-1}\|_2^2$, $1 \le i \le \mathrm{K}$, [SM9, (2.5.37)] follows

$$\phi_i = \frac{(\mathbf{v}_i^T \mathbf{r}_{i-1})^2}{\mathbf{v}_i^T \mathbf{A} \mathbf{v}_i}, \qquad 1 \le i \le \mathrm{K}.$$

Applying the normalization $\tilde{\mathbf{v}}_i = \mathbf{v}_i / \sqrt{\mathbf{v}_i^T \mathbf{A} \mathbf{v}_i}$ and the fact $\tilde{\mathbf{v}}_i^T \mathbf{r}_{i-1} = \tilde{\mathbf{v}}_i^T \mathbf{r}_0$, $1 \le i \le \mathrm{K}$, [SM3, (11)] to the previous equation gives

$$\phi_i = (\tilde{\mathbf{v}}_i^T \mathbf{r}_{i-1})^2 = (\tilde{\mathbf{v}}_i^T \mathbf{r}_0)^2, \qquad 1 \le i \le \mathrm{K}.$$ □

Equation (SM3.7) provides a geometric interpretation of $\mathbf{\Phi}$. It shows that ϕ_i is the squared $\mathbf{A}$-norm length of error $\mathbf{x}_* - \mathbf{x}_0$ in the direction $\tilde{\mathbf{v}}_i$, $1 \le i \le \mathrm{K}$.

In finite precision, the definition of $\mathbf{\Phi}$ in Theorem 3.7 and Algorithm 3.1 is preferable over (SM3.7). This is because (3.18) in Theorem 3.7 requires only local orthogonality of CG [SM18, Section 10], while (SM3.7) requires global orthogonality due to its reliance on the equalities $\mathbf{v}_i^T \mathbf{r}_{i-1} = \cdots = \mathbf{v}_i^T \mathbf{r}_0$.

SM4. Error Estimation and the Krylov Posterior. We investigate performance of estimating the error in CG by sampling from the Krylov posterior distribution. We do this with the sampling based error estimate

$$S \equiv \|X - \mathbf{x}_m\|_{\mathbf{A}}^2, \quad X \sim \mathcal{N}(\mathbf{x}_0, \widehat{\mathbf{\Gamma}}_0), \tag{SM4.1}$$

introduced in section 3.3. Additionally, in section SM4.1 we develop a $\alpha\%$ credible interval of (SM4.1) that can be computed without sampling. In section SM4.2, we compare the performance of section SM4.1 and its analytic credible interval to two existing CG error estimation techniques.

REMARK SM4.1. *Even though we are estimating CG error in this section, we remind the reader that the purpose of (SM4.1) in sections 3.3 and 4 in the main part of paper is not to estimate the error, it is to determine if the posterior is informative.*

SM4.1. Credible Interval of Sampling Based Error Estimate. The exact distribution of the sampling based error estimate (SM4.1) is a generalized chi-squared distribution and does not have a known closed form. We present an approximation that avoids the cost of sampling without losing accuracy. Compared to the many existing approximations [SM1, SM7, SM20] for distributions of Gaussian quadratic forms, our approximation is simple and designed to be computable within CG.

First we approximate (SM4.1) by a Gaussian distribution $\mathcal{N}(\mu, \sigma^2)$. Because (SM4.1) is a quadratic form, we can compute its mean and variance [SM10, Sections 3.2b.1–3.2b.3] (see also Lemma B.2). From Lemma 3.5, Theorem 3.7, and (3.19) follows that

$$\mu \equiv \operatorname{trace}(\mathbf{A}\widehat{\mathbf{\Gamma}}_m) = \sum_{i=m+1}^{m+d} \gamma_i \|\mathbf{r}_{i-1}\|_2^2 \approx \|\mathbf{x}_* - \mathbf{x}_m\|. \tag{SM4.2}$$

Following a similar argument with the variance formula in Lemma B.2 gives

$$\sigma^2 \equiv 2\ \operatorname{trace}((\mathbf{A}\widehat{\mathbf{\Gamma}}_m)^2) = 2\sum_{i=m+1}^{m+d} \gamma_i^2 \|\mathbf{r}_{i-1}\|_2^4.$$

Next we determine an '$\alpha\%$ credible interval' of $\mathcal{N}(\mu, \sigma^2)$ for some $0 < \alpha < 100$. A *credible interval* is a band around the mean μ whose width is a multiple of the standard deviation σ. Since μ is an underestimate of the error, we only need the upper one-sided upper credible interval $[\mu, S(\alpha)]$ where

$$S(\alpha) \equiv \mu + h(\alpha)\,\sigma \qquad \text{and} \qquad h(\alpha) \equiv \sqrt{2}\,\operatorname{erf}^{-1}(\alpha/100). \tag{SM4.3}$$

The *error function* erf is associated with the integral over the probability density of the normal distribution, and erf^{-1} is its inverse[1], that is $\operatorname{erf}^{-1}(\operatorname{erf}(z)) = z$.

The one-sided credible interval $[\mu, S_\alpha]$ becomes wider for large α, and narrower for small α. In section SM4.2 we select the popular choice $\alpha = 95$, and illustrate that $[\mu, S(95)]$ represents an estimate whose quality is comparable to (SM4.1).

SM4.2. Numerical Experiments. We perform numerical experiments to illustrate the accuracy of credible interval bound $S(95)$ by comparing it to the mean and samples of the sampling based error estimate (SM4.1), an empirical version of the credible interval, and state-of-the-art CG error estimators from [SM12, SM13]. After describing the setup for the numerical experiments, we present results for matrices with small dimension and large dimension, followed by a summary.

[1]The function `erfinv` is implemented in Matlab, Python's `scipy.special` library, and Julia's SpecialFunctions package.

SM4.2.1. Setup for the Numerical Experiments. We describe the setup for the numerical experiments. These estimates are plotted in each iteration m, but we suppress the explicit dependence on m to keep the notation simple.[2]

One-sided Credible Interval. We plot the upper 95% one-sided credible interval. This interval is the band between the mean μ from Theorem SM4.2 and bound $S(95)$ from (SM4.3) with $\sqrt{2}\,\mathrm{erf}(.95) = 1.96$,

$$\mu = \sum_{i=m+1}^{m+d} \gamma_i \|\mathbf{r}_i\|_2^2 \quad \text{and} \quad S(95) = \mu + 1.96\sqrt{2\sum_{i=m+1}^{m+d} \gamma_i^2 \|\mathbf{r}_{i-1}\|_2^4}. \tag{SM4.4}$$

While μ represents the known underestimate (3.19), we are not aware of other estimates of the type $S(95)$. As mentioned in Remark 3.3.3, the mean μ is equal to the CG error estimate derived from Gaussian quadrature [SM18, Section 3].

We also plot empirically computed credible interval $[\hat{\mu}, \hat{S}(95)]$ with bounds from the 10 samples of (SM4.1), where

$$\hat{\mu} = \tfrac{1}{10}\sum_{i=1}^{10} s_i \qquad \text{and} \qquad \hat{S}_{95} = \hat{\mu} + 1.96\sqrt{\tfrac{1}{9}\sum_{i=1}^{10} (s_i - \hat{\mu})^2}. \tag{SM4.5}$$

Gauss-Radau Estimates. We employ two different estimates.

(a) Gauss-Radau Upper bound [SM12, Section 4].
This is an upper bound on CG error computed with the CGQ algorithm [SM12, Section 4]. It requires a user-specified lower bound on the smallest eigenvalue of $\mathbf{A}$.

(b) Gauss-Radau Approximation [SM13, Sections 6 and 8.2].
This is an approximation of the Gauss-Radau upper bound (a) and it can underestimate the error [SM13, Section 8.2]. It does not require a bound for the smallest eigenvalue of $\mathbf{A}$, and instead approximates the smallest Ritz value of the tridiagonal matrix in CG [SM13, Section 5].

Both error estimates require running d additional CG iterations to be computed. The additional amount of iterations is called the *delay* and is analogous to the rank of the approximate Krylov posterior covariance matrix. The Gauss-Radau approximation (b) does not require a delay, however we use a delay by computing the estimate with the Ritz value from iteration $m + d$. More discussion about CG error estimates can be found in Remark 3.3.3 in the main part of the paper.

Relative Accuracy of Estimates. We plot the relative difference between an estimate E and the squared $\mathbf{A}$-norm error $\|\mathbf{x}_* - \mathbf{x}_m\|_{\mathbf{A}}^2$,

$$\rho(E) = \frac{\left|E - \|\mathbf{x}_* - \mathbf{x}_m\|_{\mathbf{A}}^2\right|}{\min\{E,\ \|\mathbf{x}_* - \mathbf{x}_m\|_{\mathbf{A}}^2\}}, \tag{SM4.6}$$

where E can be μ, $S(95)$, or one of the Gauss-Radau estimators. The minimum in the denominator avoids favoring underestimate or overestimates, so that smaller values $\rho(E)$ indicate more accurate estimators E.

Inputs. The linear systems $\mathbf{A}\mathbf{x}_* = \mathbf{b}$ have a size $n = 48$ or $n = 11948$ symmetric positive definite matrix $\mathbf{A} \in \mathbb{R}^{n\times n}$, solution vector of all ones $\mathbf{x}_* = \mathbf{1} \in \mathbb{R}^n$, and right-hand side vector $\mathbf{b} = \mathbf{A}\mathbf{1}$. The initial guess $\mathbf{x}_0 = \mathbf{0} \in \mathbb{R}^n$ is the zero vector.

[2] The Python code used in the numerical experiments can be found at https://github.com/treid5/ProbNumCG_Supp

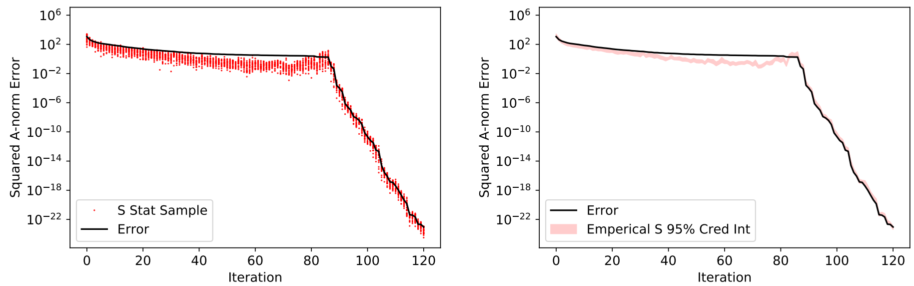

Figure SM4.1: Squared $\mathbf{A}$-norm error $\|\mathbf{x}_* - \mathbf{x}_m\|_{\mathbf{A}}^2$ versus iteration m for the matrix $\mathbf{A}$ with eigenvalue distribution (SM4.7). On the left: samples s_i from (SM4.1). On the right: empirical upper credible interval $[\hat{\mu}, \hat{S}(95)]$ from (SM4.5).

SM4.2.2. Matrix with Small Dimension. We first examine the error estimates on a size $n = 48$ random matrix $\mathbf{A} = \mathbf{Q}\mathbf{D}\mathbf{Q}^T$ [SM6, Section 2], where $\mathbf{Q}$ is a random orthogonal matrix with Haar distribution [SM16, Section 3] and $\mathbf{D}$ is a diagonal matrix with eigenvalues [SM17]

$$d_{ii} = 0.1 + \frac{i-1}{n-1}\left(10^4 - 0.1\right)(0.9)^{n-i}, \qquad 1 \leq i \leq 48. \tag{SM4.7}$$

The eigenvalue distribution is chosen to increase round off errors in CG, and is similar to the one in [SM18, Section 11] for testing (3.19). The two-norm condition number is $\kappa_2(\mathbf{A}) = 10^5$.

Figures SM4.1 and SM4.2 display the squared $\mathbf{A}$-norm error $\|\mathbf{x}_* - \mathbf{x}_m\|_{\mathbf{A}}^2$ and the estimates over 120 iterations. The delay used to compute the error estimates and posterior covariance rank is $d = 4$.

Figure SM4.1 plots the samples s_i from (SM4.1) on the left, and the empirical upper credible interval $[\hat{\mu}, \hat{S}(95)]$ from (SM4.5) on the right. Both underestimate the error in the initial period of slow convergence, cover the error during fast convergence, and underestimate the error once maximal attainable accuracy has been reached. The upper credible intervals appear deceptively thinner because of the logarithmic scale on the vertical axis.

The left part of Figure SM4.2 plots the credible interval $[\mu, S(95)]$ from (SM4.4); as well as the Gauss-Radau bound (a) and approximation (b). The Gauss-Radau bound is computed with a lower bound of $9.99 \cdot 10^{-2}$ for the smallest eigenvalue 0.1 of $\mathbf{A}$. The upper credible interval $[\mu, S(95)]$ behaves like its empirical version $[\hat{\mu}, \hat{S}(95)]$ in Figure SM4.1, and therefore represents an accurate approximation. The Gauss-Radau bound (a) overestimates the error, and the Gauss-Radau approximation (b) underestimates the error when convergence is slow and overestimates it when convergence is fast. Note that the bound $S(95)$ underestimates the error during slow convergence and overestimates it during fast convergence.

The right part of Figure SM4.2 plots the relative accuracy (SM4.6) for the mean μ from (SM4.4), the bound $S(95)$ from (SM4.4), the Gauss-Radau bound (a) and the Gauss-Radau approximation (b). During the initial period of slow convergence, the bound $S(95)$ starts out as the most accurate until iteration 75 when the Gauss-Radau bound (a) becomes the most accurate. During fast convergence, after iteration 90, the mean μ is most accurate.

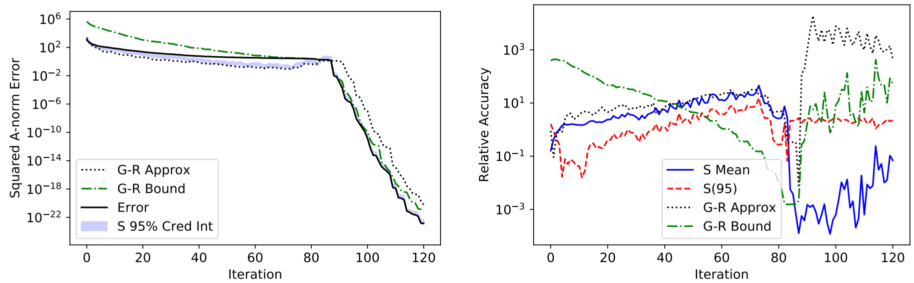

Figure SM4.2: Squared **A**-norm error $\|\mathbf{x}_* - \mathbf{x}_m\|_{\mathbf{A}}^2$ and relative accuracy versus iteration m for the matrix **A** with eigenvalue distribution (SM4.7). On the left: upper credible interval $[\mu, S(95)]$ from (SM4.4), Gauss-Radau bound (a), and Gauss-Radau approximation (b). On the right: relative accuracy ρ from (SM4.6) for the mean μ and bound $S(95)$ from (SM4.4) as well as the Gauss-Radau bound (a) and approximation (b).

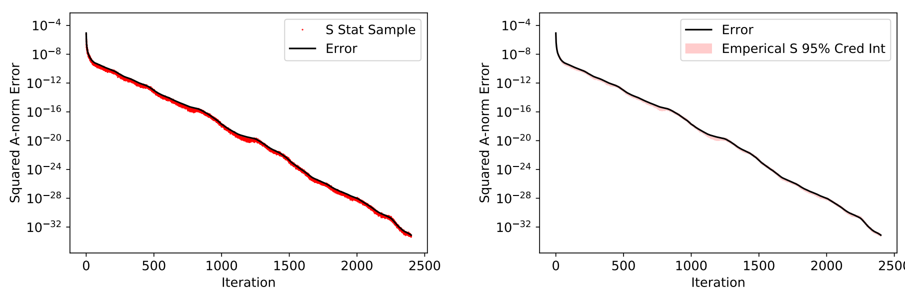

Figure SM4.3: Squared **A**-norm error $\|\mathbf{x}_* - \mathbf{x}_m\|_{\mathbf{A}}^2$ versus iteration m for the matrix **A** based on `BCSSTK18`. On the left: samples s_i from (SM4.1). On the right: empirical upper credible interval $[\hat{\mu}, \hat{S}(95)]$ from (SM4.5).

SM4.2.3. Matrix with Large Dimension. We now examine the error estimates on the same $n = 11948$ matrix as in section 4.3.

Figures SM4.3 and SM4.4 display the squared **A**-norm error $\|\mathbf{x}_* - \mathbf{x}_m\|_{\mathbf{A}}^2$ and the estimates over 2,700 iterations. The delay and posterior covariance has rank is $d = 50$.

Figure SM4.3 plots the samples s_i from (SM4.1) on the left, and the empirical credible interval $[\hat{\mu}, \hat{S}(95)]$ from (SM4.5) on the right. Both behave as in Figure SM4.1 and closely underestimate the error.

The left part of Figure SM4.4 plots the credible interval $[\mu, S(95)]$ from (SM4.4); as well as the Gauss-Radau bound (a) and approximation (b). The Gauss-Radau bound is computed with a lower bound of $9 \cdot 10^{-14}$ for the smallest eigenvalue of **A**. Again, the behavior is similar as in Figure SM4.2.

The right part of Figure SM4.4 plots the relative accuracy (SM4.6) for the mean μ from (SM4.4), the bound $S(95)$ from (SM4.4), and the Gauss-Radau approximation (b). As before, the bound $S(95)$ is generally the most accurate, followed by the

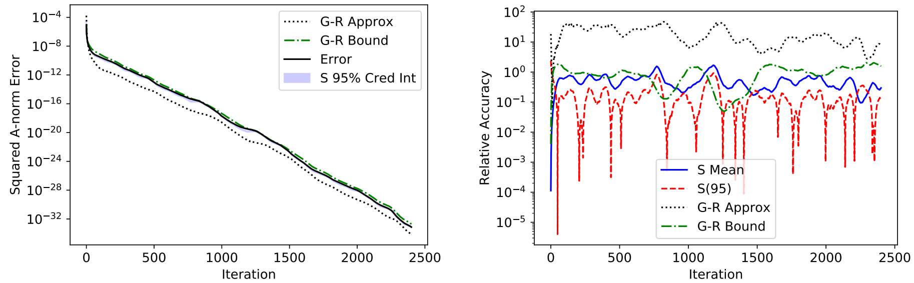

Figure SM4.4: Squared $\mathbf{A}$-norm error $\|\mathbf{x}_* - \mathbf{x}_m\|_{\mathbf{A}}^2$ and relative accuracy versus iteration m for the matrix $\mathbf{A}$ based on BCSSTK18. On the left: upper credible interval $[\mu, S(95)]$ from (SM4.4), and Gauss-Radau bound (a) and approximation (b). On the right: relative accuracy ρ of the error estimates.

mean μ.

SM4.2.4. Summary of the Experiments. Numerical experiments confirm that the sampling based error estimate (SM4.1) performs as expected. In particular, the upper credible interval $[\mu, S(95)]$ in (SM4.4) is an accurate approximation of the empirical upper credible interval $[\hat{\mu}, \hat{S}(95)]$ in (SM4.5).

The speed of convergence impacts the effectiveness of (SM4.1) as an error estimate. The credible interval $[\mu, S(95)]$ (SM4.4) depends on the mean μ, and the distance between μ and the error depends on convergence speed. As a consequence, the mean and credible interval are far from the error when convergence is slow.

Convergence speed can also affect the Gauss-Radau approximation (b). The convergence rate of the smallest Ritz value to the smallest eigenvalue is usually related to convergence of the $\mathbf{A}$-norm error [SM13, Section 8.1 and Figures 3 and 4]. Slow convergence of the $\mathbf{A}$-norm means the Ritz value has not converged to the smallest eigenvalue, and this causes the Gauss-Radau approximation (b) to be less accurate.

In general, the bound $S(95)$ tends to underestimate the error during slow convergence and to cover the error during fast convergence. The distance between $S(95)$ and the error is competitive with the Gauss-Radau estimates.

REFERENCES

[SM1] T. Y. Al-Naffouri, M. Moinuddin, N. Ajeeb, B. Hassibi, and A. L. Moustakas, *On the distribution of indefinite quadratic forms in Gaussian random variables*, IEEE Trans. Commun., 64 (2016), pp. 153–165.

[SM2] D. Calvetti, *Contributed discussion for "A Bayesian conjugate gradient method"*, Bayesian Anal., 14 (2019), pp. 937–1012, https://doi.org/10.1214/19-BA1145.

[SM3] J. Cockayne, C. J. Oates, I. C. F. Ipsen, and M. Girolami, *A Bayesian conjugate gradient method (with discussion)*, Bayesian Anal., 14 (2019), pp. 937–1012, https://doi.org/10.1214/19-BA1145. Includes 6 discussions and a rejoinder from the authors.

[SM4] J. Cockayne, C. J. Oates, I. C. F. Ipsen, and M. Girolami, *Rejoinder for "A Bayesian conjugate gradient method"*, Bayesian Anal., 14 (2019), pp. 937–1012, https://doi.org/10.1214/19-BA1145.

[SM5] J. Cockayne, C. J. Oates, I. C. F. Ipsen, and M. Girolami, *Supplementary material for 'A Bayesian conjugate-gradient method'*, Bayesian Anal., (2019), https://doi.org/10.1214/19-BA1145SUPP.

[SM6] A. Greenbaum and Z. Strakoš, *Predicting the behavior of finite precision Lanczos and* *conjugate gradient computations*, SIAM J. Matrix Anal. Appl., 13 (1992), pp. 121–137, https://doi.org/10.1137/0613011.

[SM7] D. Jensen and H. Solomon, *A Gaussian approximation to the distribution of a definite quadratic form*, J. Amer. Statist. Assoc., 67 (1972), pp. 898–902.

[SM8] L. Li and E. X. Fang, *Invited discussion for "A Bayesian conjugate gradient method"*, Bayesian Anal., 14 (2019), pp. 937–1012, https://doi.org/10.1214/19-BA1145.

[SM9] J. Liesen and Z. Strakos, *Krylov Subspace Methods: Principles and Analysis*, Oxford University Press, 2013.

[SM10] A. M. Mathai and S. B. Provost, *Quadratic forms in random variables: Theory and applications*, Dekker, 1992.

[SM11] G. Meurant and Z. Strakoš, *The Lanczos and conjugate gradient algorithms in finite precision arithmetic*, Acta Numer., 15 (2006), pp. 471–542, https://doi.org/10.1017/S096249290626001X.

[SM12] G. Meurant and P. Tichý, *On computing quadrature-based bounds for the A-norm of the error in conjugate gradients*, Numer. Algorithms, 62 (2013), pp. 163–191, https://doi.org/10.1007/s11075-012-9591-9.

[SM13] G. Meurant and P. Tichý, *Approximating the extreme Ritz values and upper bounds for the A-norm of the error in CG*, Numer. Algorithms, 82 (2019), pp. 937–968, https://doi.org/10.1007/s11075-018-0634-8.

[SM14] R. J. Muirhead, *Aspects of multivariate statistical theory*, John Wiley & Sons, Inc., New York, 1982. Wiley Series in Probability and Mathematical Statistics.

[SM15] D. V. Ouellette, *Schur complements and statistics*, Linear Algebra Appl., 36 (1981), pp. 187–295, https://doi.org/10.1016/0024-3795(81)90232-9.

[SM16] G. W. Stewart, *The efficient generation of random orthogonal matrices with an application to condition estimators*, SIAM J. Numer. Anal., 17 (1980), pp. 403–409, https://doi.org/10.1137/0717034.

[SM17] Z. Strakoš, *On the real convergence rate of the conjugate gradient method*, Linear Algebra Appl., 154/156 (1991), pp. 535–549, https://doi.org/10.1016/0024-3795(91)90393-B.

[SM18] Z. Strakoš and P. Tichý, *On error estimation in the conjugate gradient method and why it works in finite precision computations*, Electron. Trans. Numer. Anal., 13 (2002), pp. 56–80.

[SM19] A. M. Stuart, *Inverse problems: a Bayesian perspective*, Acta Numer., 19 (2010), pp. 451–559, https://doi.org/10.1017/S0962492910000061.

[SM20] G. G. Tziritas, *On the distribution of positive-definite Gaussian quadratic forms*, IEEE Trans. Inform. Theory, 33 (1987), pp. 895–906.