# Probabilistic Numerical Linear Solvers 

## Ilse C.F. Ipsen

## NC STATE UNIVERSITY

Raleigh, NC, USA

Research supported in part by NSF DMS: RTG \& FRG

## Collaborators



Jon Cockayne
Alan Turing Institute, UK


Chris Oates
University of Newcastle upon Tyne, UK


Tim Reid

North Carolina State University, USA

## Probabilistic Numerics

- Statistical treatment of approximation errors in deterministic numerical methods:

Assign probability distributions to quantities of interest Express methods as probabilistic inference

- Model uncertainty due to limited computational resources (Truncation errors in discretizations, termination of iterative methods)


## Probabilistic Numerics

- Statistical treatment of approximation errors in deterministic numerical methods:

Assign probability distributions to quantities of interest Express methods as probabilistic inference

- Model uncertainty due to limited computational resources
(Truncation errors in discretizations, termination of iterative methods)
But why??? Want error measures that
- are more informative than traditional, pessimistic bounds
- can be propagated through computational pipelines


## Probabilistic Numerics

- Statistical treatment of approximation errors in deterministic numerical methods:

Assign probability distributions to quantities of interest Express methods as probabilistic inference

- Model uncertainty due to limited computational resources
(Truncation errors in discretizations, termination of iterative methods)
But why??? Want error measures that
- are more informative than traditional, pessimistic bounds
- can be propagated through computational pipelines

Probabilistic numerical methods have been developed for:
Approximation, quadrature, numerical solution of ordinary and partial differential equations, optimization

Here: Computational kernels and linear algebra

# Our Focus: Accurate Error Estimation for Iterative Solution of Linear Systems <br> [Disclaimer: Research still at proof-of-concept stage] 

Given: Nonsingular real matrix $\boldsymbol{A}$, vector $\boldsymbol{b}$
Want: Solution $\boldsymbol{x}_{*}$ of linear system $\boldsymbol{A} \boldsymbol{x}_{*}=\boldsymbol{b}$
Iterative solver:
From user-specified initial guess $x_{0}$
computes iterates $\boldsymbol{x}_{1}, \boldsymbol{x}_{2}, \ldots, \boldsymbol{x}_{m} \rightarrow \boldsymbol{x}_{*}$

- Prior Distribution

Reflects initial knowledge about $\boldsymbol{x}_{*}$

- Posterior Distribution at iteration $m$

Reflects knowledge about $\boldsymbol{x}_{*}$ after $m$ iterations

## What We Can Do So Far (Details Later)



Exact error $\left\|\boldsymbol{x}_{m}-\boldsymbol{x}_{*}\right\| /\left\|\boldsymbol{x}_{*}\right\|$
Traditional error bounds $\geq 1$, not informative

## What We Can Do So Far (Details Later)



Exact error $\left\|\boldsymbol{x}_{m}-\boldsymbol{x}_{*}\right\| /\left\|\boldsymbol{x}_{*}\right\|$
Traditional error bounds $\geq 1$, not informative
Our S statistic from posterior distribution $\approx$ exact error

## Overview

(1) Probabilistic Numerics: The Big Picture
(2) Gaussian Probability Distributions
(3) BayesCG, a Probabilistic Numerical Linear Solver
(9) Errors, and Metrics on Probability Distributions
(5) Two Priors that Reproduce Conjugate Gradient

Inverse prior
Krylov priors
(0) Low-Rank Approximations of Krylov Priors

## 1. Probabilistic Numerics: The Big Picture


http://probabilistic-numerics.org/

## Early Probabilistic Numerics

- H. Poincaré (1896, 1912): Calcul des Probabilités
- A.V. Sul'din (1959): Wiener Measure and its Application to Approximation Methods, I. Izv. Vysš. Učben. Zayed. Mat.
- F. M. Larkin (1972): Gaussian Measure in Hilbert space and Applications in Numerical Analysis, Rocky Mountain J. Math.
- G.S. Kimeldorf, G. Wahba (1970): A Correspondence between Bayesian Estimation on Stochastic Processes and Smoothing by Splines, Ann. Math. Stat.
- P. Diaconis (1988): Bayesian Numerical Analysis, Stat. Decis. Theory Relat. Top. IV
- A. O'Hagan (1992): Some Bayesian Numerical Analysis, Bayesian. Stat.


## Perspectives on Modern Probabilistic Numerics

- Bayesian inference in inverse problems A.M. Stuart (2010): Inverse Problems: A Bayesian Perspective, Acta Numer.
- 'A call to arms for probabilistic numerical methods' P. Hennig, M.A Osborne, M. Girolami (2015): Probabilistic Numerics and Uncertainty in Computations, Proc. R. Soc. A
- Statistical foundations for non-linear, non-Gaussian problems J. Cockayne, C.J. Oates, T.J. Sullivan, M. Girolami (2019): Bayesian Probabilistic Numerical Methods, SIAM Rev.
- Historical development of probabilistic numerics C.J. Oates, T.J. Sullivan (2019): A Modern Retrospective on Probabilistic Numerics, Stat. Comput.
- Probabilistic Numerics at SIAM Conference on UQ http://probabilistic-numerics.org/meetings/SIAMUQ2020/


# Application: Global Optimization 



## Application: Numerical Integration for Rendering of Glossy Surfaces



Fig. 1. Indirect radiance component for the Room scene rendered with low discrepancy Monte Carlo Importance Sampling (LDIS, (b)) and BMC (c). The shown images have been multiplied by a factor of 4. The RMSE and the time are computed by considering only the glossy component. 16 ray samples per visible point were used for the materials with a glossy BRDF (teapot, tea cup and fruit-dish), while 64 samples were used for the diffuse BRDFs.
> R. Marques, C. Bouville, M. Ribardiere, L.P. Santos, K. Bouatouch (2013), IEEE Trans. Vis. Comput. Graph.

## Probabilistic Numerics \& Other Areas

Connections to

- Information-based complexity
[Traub, Wasilkowski, Woźniakowski 1983]
- Average-case analysis
[Novak 1988], [Ritter 2000]
- Data assimilation, Kalman filters
[Law, Stuart, Zygalakis 2015], [Reich, Cotter 2015]
- Statistical learning
[Hastie, Tibshirani, Friedman 2009], [Rasmussen, Williams 2006]

Difference to Uncertainty Quantification [Smith 2014]

- UQ: Problems usually ill-posed, uncertainties from all sources: model, parameter, experimental, measurement
- PN: Problems are well-posed, algorithmic uncertainty due to limited computational resources


## Probabilistic Numerics in Linear Algebra

- P. Hennig (2015): Probabilistic Interpretation of Linear Solvers, SIAM J. Optim.
- S. Bartels, P. Hennig (2015): Probabilistic Approximate Least-Squares, Proc. Machine Learning Research
- F. Schäfer, T.J. Sullivan, H. Owhadi (2017): Compression, Inversion, and Approximate PCA of Dense Kernel Matrices at Near-Linear Computational Complexity, arXiv
- J. Cockayne, C.J. Oates, I.C.F. Ipsen, M. Girolami (2019): A Bayesian Conjugate Gradient Method, Bayesian Anal.
- S. Bartels, J. Cockayne, I.C.F. Ipsen, P. Hennig (2019): Probabilistic Linear Solvers: A Unifying View, Stat. Comput.


## 2. Gaussian Probability Distributions


[The MathWorks R2020a]

## Our Probability Distributions are Gaussians

Gaussian $=$ multi-variate normal distribution
Gaussian random vector $\boldsymbol{x} \sim \mathcal{N}\left(\mu_{0}, \Sigma_{0}\right)$
Mean $\mu_{0} \in \mathbb{R}^{d}$
Covariance $\Sigma_{0} \in \mathbb{R}^{d \times d}$ symmetric positive semi-definite
If $\Sigma_{0}$ also positive definite, then probability density function is

$$
\begin{aligned}
& \qquad p(\boldsymbol{x})=\frac{1}{\sqrt{\operatorname{det}\left(\Sigma_{0}\right)(2 \pi)^{d}}} \exp \left(-\frac{1}{2}\left\|\boldsymbol{x}-\mu_{0}\right\|_{\Sigma_{0}^{-1}}^{2}\right) \\
& \text { where }\left\|\boldsymbol{x}-\mu_{0}\right\|_{\Sigma_{0}^{-1}}^{2}=\left(\boldsymbol{x}-\mu_{0}\right)^{T} \Sigma_{0}^{-1}\left(\boldsymbol{x}-\mu_{0}\right) \text { and } T \text { is transpose }
\end{aligned}
$$

Why Gaussians?
Stability: Linear transformations preserve Gaussianity
If $\boldsymbol{x} \sim \mathcal{N}\left(\mu_{0}, \Sigma_{0}\right)$ then $B x+z \sim \mathcal{N}\left(B \mu_{0}+z, B \Sigma_{0} B^{T}\right)$
[Muirhead 1982], [Stuart 2010]

## Stability of Gaussians for Computing Posteriors

If $\boldsymbol{x} \sim \mathcal{N}\left(\mu_{0}, \Sigma_{0}\right)$ then $z+B x \sim \mathcal{N}\left(z+B \mu_{0}, B \Sigma_{0} B^{T}\right)$

- Prior: $\boldsymbol{x} \sim \mathcal{N}\left(\mu_{0}, \Sigma_{0}\right)$
- Condition on linear information $\boldsymbol{y}=\boldsymbol{B} \boldsymbol{x}$
- Bayesian inference

Posterior: $\boldsymbol{x} \mid \boldsymbol{y} \sim \mathcal{N}(\boldsymbol{\mu}, \Sigma)$ with

$$
\begin{aligned}
\mu & =\mu_{0}+\Sigma_{0} B^{T}\left(B \Sigma_{0} B^{T}\right)^{-1}\left(\boldsymbol{y}-B \mu_{0}\right) \\
\Sigma & =\Sigma_{0}-\Sigma_{0} B^{T}\left(B \Sigma_{0} B^{T}\right)^{-1} B \Sigma_{0}
\end{aligned}
$$

Connections to: Schur complements, projectors

Conditioning Gaussian random vector $\boldsymbol{x} \sim \mathcal{N}\left(\boldsymbol{x}_{0}, \Sigma_{0}\right)$ on linear information $\boldsymbol{y}=\boldsymbol{B} \boldsymbol{x}$ gives another Gaussian random vector $\boldsymbol{x} \mid \boldsymbol{y} \sim \mathcal{N}(\boldsymbol{\mu}, \Sigma)$

## Stability of Gaussians for Computational Sampling

- Want: Sample $X \sim \mathcal{N}(\mu, \Sigma)$
- Stability: If $\boldsymbol{x} \sim \mathcal{N}(0, \boldsymbol{I})$ then $\mu+\boldsymbol{L} \boldsymbol{x} \sim \mathcal{N}\left(\mu, L L^{T}\right)$
- Possible factorizations $\Sigma=\boldsymbol{L L}^{T}$

Square root $\boldsymbol{L}=\Sigma^{1 / 2}$
Cholesky factorization, if $\Sigma$ nonsingular
Thin Cholesky: $L$ has $\operatorname{rank}(\Sigma)$ columns

- Matlab

$$
X=\boldsymbol{\mu}+\boldsymbol{L} * \underbrace{\operatorname{randn}(\operatorname{size}(\boldsymbol{L}, 2), 1)}_{\text {vector } \sim \mathcal{N}(0, \boldsymbol{I})}
$$

3. BayesCG, a Probabilistic Numerical Linear Solver

## Probabilistic Linear System Solution

Given: Symmetric positive-definite $\boldsymbol{A} \in \mathbb{R}^{d \times d}$, vector $\boldsymbol{b} \in \mathbb{R}^{d}$ Want: Solution of $\boldsymbol{A} \boldsymbol{x}_{*}=\boldsymbol{b}$
Initial guess $x_{0}$ : $\boldsymbol{A}\left(\boldsymbol{x}_{*}-\boldsymbol{x}_{0}\right)=\boldsymbol{b}-\boldsymbol{A} \boldsymbol{x}_{0}$
Solver: Computes iterates $\boldsymbol{x}_{1}, \boldsymbol{x}_{2}, \ldots, \boldsymbol{x}_{m} \rightarrow \boldsymbol{x}_{*}$
Solution-based inference
Probability distribution over solution space $x \in \mathbb{R}^{d}$
Keeping it simple: $\boldsymbol{x}_{m}=\mu_{m}, m \geq 0$
Iterates coincide with means of probability distribution

- User specifies Gaussian prior $\mathcal{N}\left(\boldsymbol{x}_{0}, \Sigma_{0}\right)$

Prior reflects initial knowledge about $\boldsymbol{x}_{*}$

- Solver computes Gaussian posteriors $\mathcal{N}\left(\boldsymbol{x}_{m}, \Sigma_{m}\right)$

Posteriors reflect knowledge about $\boldsymbol{x}_{*}$ after $m$ iterations

## Computing the Posteriors in Iteration $m \geq 1$

[Cockayne, Oates, Sullivan, Girolami 2019]

Information about $\boldsymbol{x}_{*}$ provided by $m$ search directions

$$
\boldsymbol{S}_{m}=\left(\begin{array}{lll}
\boldsymbol{s}_{1} & \cdots & \boldsymbol{s}_{m}
\end{array}\right) \in \mathbb{R}^{d \times m} \quad \operatorname{rank}\left(\boldsymbol{S}_{m}\right)=m
$$

Condition on $\boldsymbol{y}_{m}=\boldsymbol{S}_{m}^{T} \boldsymbol{A} \boldsymbol{x}_{*}=\boldsymbol{S}_{m}^{T} \boldsymbol{b}$
Exists unique Bayesian method that outputs posterior distribution

$$
\boldsymbol{x} \mid \boldsymbol{y}_{m} \sim \mathcal{N}\left(\boldsymbol{x}_{m}, \Sigma_{m}\right)
$$

with

$$
\begin{aligned}
\boldsymbol{x}_{m} & =\boldsymbol{x}_{0}+\Sigma_{0} \boldsymbol{A} \boldsymbol{S}_{m}\left(\boldsymbol{S}_{m}^{T} \boldsymbol{A} \Sigma_{0} \boldsymbol{A} \boldsymbol{S}_{m}\right)^{-1} \boldsymbol{S}_{m}^{T}\left(\boldsymbol{b}-\boldsymbol{A} \boldsymbol{x}_{0}\right) \\
\Sigma_{m} & =\Sigma_{0}-\Sigma_{0} \boldsymbol{A} \boldsymbol{S}_{m}\left(\boldsymbol{S}_{m}^{T} \boldsymbol{A} \Sigma_{0} \boldsymbol{A} \boldsymbol{S}_{m}\right)^{-1} \boldsymbol{S}_{m}^{T} \boldsymbol{A} \Sigma_{0}
\end{aligned}
$$

Choose $\mathbf{A} \Sigma_{0} \mathbf{A}$-orthogonal search directions: $S_{m}^{T} \mathbf{A} \Sigma_{0} \mathbf{A} \boldsymbol{S}_{m}=\boldsymbol{I}_{m}$

## Bayesian Conjugate Gradient Method (BayesCG)

[Cockayne, Oates, Ipsen, Girolami 2019]
$\boldsymbol{r}_{0}=\boldsymbol{b}-\boldsymbol{A} \boldsymbol{x}_{0}$
$\boldsymbol{s}_{1}=\boldsymbol{r}_{0} / \sqrt{\boldsymbol{r}_{0}{ }^{T} \boldsymbol{A} \Sigma_{0} \boldsymbol{A} \boldsymbol{r}_{0}} \quad$ \{initial search direction $\}$
while not converged do

$$
\begin{aligned}
& \qquad \begin{array}{ll}
\Sigma_{m}=\Sigma_{m-1}-\left(\Sigma_{0} \boldsymbol{A} \boldsymbol{s}_{m}\right)\left(\Sigma_{0} \boldsymbol{A} \boldsymbol{s}_{m}\right)^{T} & \{\text { next posterior }\} \\
\boldsymbol{x}_{m}=\boldsymbol{x}_{m-1}+\Sigma_{0} \boldsymbol{A} \boldsymbol{s}_{m}\left(\boldsymbol{r}_{m-1}^{T} \boldsymbol{s}_{m}\right) & \{\text { next iterate }\} \\
\boldsymbol{r}_{m}=\boldsymbol{b}-\boldsymbol{A} \boldsymbol{x}_{m} & \{\text { next residual }\} \\
\tilde{\boldsymbol{s}}_{m+1}=\boldsymbol{r}_{m}-\boldsymbol{s}_{m}\left(\boldsymbol{r}_{m}^{T} \boldsymbol{A} \Sigma_{0} \boldsymbol{A} \boldsymbol{s}_{m}\right) & \{\text { next search direction }\} \\
\boldsymbol{s}_{m+1}=\tilde{\boldsymbol{s}}_{m+1} / \sqrt{\tilde{\boldsymbol{s}}_{m+1}^{T} \boldsymbol{A} \Sigma_{0} \boldsymbol{A} \tilde{\boldsymbol{s}}_{m+1}} & \text { \{normalize search dir\}} \\
\text { end while } &
\end{array} \text { }
\end{aligned}
$$

## Properties of BayesCG

- Krylov space for iterates: $\boldsymbol{x}_{m} \in \mathcal{K}_{m}^{*}$

$$
\mathcal{K}_{m}^{*}=\boldsymbol{x}_{0}+\mathcal{K}_{m}\left(\Sigma_{0} \boldsymbol{A}^{2}, \Sigma_{0} \boldsymbol{A}\left(\boldsymbol{b}-\boldsymbol{A} \boldsymbol{x}_{0}\right)\right)
$$

- Error minimization in $\Sigma_{0}^{-1}$ norm

$$
\left\|\boldsymbol{x}_{m}-\boldsymbol{x}_{*}\right\|_{\Sigma_{0}^{-1}}=\min _{\boldsymbol{x} \in \mathcal{K}_{m}^{*}}\left\|\boldsymbol{x}-\boldsymbol{x}_{*}\right\|_{\Sigma_{0}^{-1}}
$$

- Convergence in $\Sigma_{0}^{-1}$ norm depends on conditioning of $\Sigma_{0} \boldsymbol{A}^{2}$

$$
\left\|\boldsymbol{x}_{m}-\boldsymbol{x}_{*}\right\|_{\Sigma_{0}^{-1}} \leq 2\left(\frac{\sqrt{\kappa_{2}\left(\Sigma_{0} \boldsymbol{A}^{2}\right)}-1}{\sqrt{\kappa_{2}\left(\Sigma_{0} \boldsymbol{A}^{2}\right)}+1}\right)^{m}\left\|\boldsymbol{x}_{0}-\boldsymbol{x}_{*}\right\|_{\Sigma_{0}^{-1}}
$$

- Prior: If $\Sigma_{0}=\boldsymbol{A}^{-1}$ then BayesCG $=C G$ $C G=$ Bayesian inference with prior $\mathcal{N}\left(\boldsymbol{x}_{0}, \boldsymbol{A}^{-1}\right)$


## Summary: BayesCG

Solve symmetric positive-definite system $\boldsymbol{A} \boldsymbol{x}_{*}=\boldsymbol{b}$

- BayesCG = probabilistic extension of Conjugate Gradient Models uncertainty in solution $\boldsymbol{x}_{*}$ due to early termination
- Input: Gaussian prior $\mathcal{N}\left(\boldsymbol{x}_{0}, \Sigma_{0}\right)$ Models initial uncertainty about solution $\boldsymbol{x}_{*}$ Mean $x_{0}$ identical to initial guess for iterates
- Output: Gaussian posterior $\mathcal{N}\left(\boldsymbol{x}_{m}, \Sigma_{m}\right)$ at iteration $m$ Models uncertainty about solution $\boldsymbol{x}_{*}$ after $m$ iterations Mean identical to iterate $\boldsymbol{x}_{m} \approx \boldsymbol{x}_{*}$

Next: How to use the BayesCG posteriors?

## 4. Errors, and <br> Metrics on Probability Distributions

## Metric on Set of Probability Distributions

Measure distance between two probability distributions

- Any two Gaussian distributions

$$
\mathcal{G}_{m}=\mathcal{N}\left(\boldsymbol{x}_{m}, \Sigma_{m}\right) \text { and } \mathcal{G}_{*}=\mathcal{N}\left(\boldsymbol{x}_{*}, \Sigma_{*}\right)
$$

- Wasserstein 2-norm metric: Distance between $\mathcal{G}_{m}$ and $\mathcal{G}_{*}$

$$
\left[W_{2}\left(\mathcal{G}_{m}, \mathcal{G}_{*}\right)\right]^{2}=\left\|\boldsymbol{x}_{m}-\boldsymbol{x}_{*}\right\|_{2}^{2}+\operatorname{trace}\left(\Sigma_{m}+\Sigma_{*}-2\left(\Sigma_{m} \Sigma_{*}\right)^{1 / 2}\right)
$$

[Dowson, Landau 1982]

- Extension to general weighted norms ( $\boldsymbol{B}$ is spd)

$$
\begin{aligned}
& {\left[W_{B}\left(\mathcal{G}_{m}, \mathcal{G}_{*}\right)\right]^{2}=\left\|x_{m}-\boldsymbol{x}_{*}\right\|_{\boldsymbol{B}}^{2}} \\
& \quad+\operatorname{trace}\left(\boldsymbol{B}\left(\Sigma_{m}+\Sigma_{*}\right)-2\left(\boldsymbol{B}^{1 / 2} \Sigma_{m} \boldsymbol{B} \Sigma_{*} \boldsymbol{B}^{1 / 2}\right)^{1 / 2}\right)
\end{aligned}
$$

## Application of Wasserstein Metric to BayesCG

Input: Prior $\mathcal{G}_{0}=\mathcal{N}\left(x_{0}, \Sigma_{0}\right)$
Output: Posteriors $\mathcal{G}_{m}=\mathcal{N}\left(\boldsymbol{x}_{m}, \Sigma_{m}\right)$
Errors: $\left\|\boldsymbol{x}_{m}-\boldsymbol{x}_{*}\right\|_{\Sigma_{0}^{-1}}$
Represent solution $\boldsymbol{x}_{*}$ as point distribution $\mathcal{N}\left(\boldsymbol{x}_{*}, 0\right)$

$$
\begin{aligned}
{\left[W_{\Sigma_{0}^{-1}}\left(\mathcal{G}_{m}, x_{*}\right)\right]^{2} } & =\left\|x_{m}-x_{*}\right\|_{\Sigma_{0}^{-1}}^{2}+\operatorname{trace}\left(\Sigma_{0}^{-1} \Sigma_{m}\right) \\
& =\left\|x_{m}-x_{*}\right\|_{\Sigma_{0}^{-1}}^{2}+(d-m)
\end{aligned}
$$

$\Sigma_{0}^{-1}$ distance of posterior $\mathcal{G}_{m}$ to solution $\boldsymbol{x}_{*}$ equals $\Sigma_{0}^{-1}$ norm error + dimension of unexplored space

Computational error estimation with "S-statistic"

$$
\left\|\boldsymbol{x}_{m}-\boldsymbol{x}_{*}\right\|_{\Sigma_{0}^{-1}}^{2} \approx\left\|\boldsymbol{x}_{m}-X\right\|_{\Sigma_{0}^{-1}}^{2} \quad \text { where } \quad X \sim \mathcal{G}_{m}
$$

## Possible Prior Distributions

- "Noninformative": $\Sigma_{0}=\boldsymbol{I}_{d}$
- Inverse: $\Sigma_{0}=\boldsymbol{A}^{-1}$, BayesCG $=$ CG
- Natural: $\Sigma_{0}=\boldsymbol{A}^{-2}$, convergence in 1 iteration
- Preconditioner: $\Sigma_{0}=\left(\boldsymbol{P}^{T} \boldsymbol{P}\right)^{-1} \approx \boldsymbol{A}^{-2}$
- Krylov ${ }^{1}: \Sigma_{0}=\boldsymbol{V} \Phi \boldsymbol{V}^{\top}$, where columns of $\boldsymbol{V}$ are basis for Krylov space
- Hierarchical: $\Sigma_{0}=\nu \hat{\Sigma}_{0}$ with Jeffrey's improper $p(\nu) \sim \nu^{-1}$

Priors that reproduce CG: Inverse and Krylov ${ }^{1}$ Impractical "academic" priors that already contain all information to compute solution immediately
But: We study them to calibrate our expectations In order to: Develop practical, low-rank approximations

[^0]
## 5a. Priors that Reproduce CG Inverse Prior

## Wasserstein Metric for BayesCG under Inverse Prior

- Prior $\mathcal{G}_{0}=\mathcal{N}\left(\boldsymbol{x}_{0}, \Sigma_{0}\right)$ with $\Sigma_{0}=\boldsymbol{A}^{-1}$
- BayesCG under inverse prior $=C G+$ posteriors
- Posteriors $\mathcal{G}_{m}=\mathcal{N}\left(\boldsymbol{x}_{m}, \Sigma_{m}\right)$ with $\Sigma_{m}=\Sigma_{m-1}-\boldsymbol{s}_{m} \boldsymbol{s}_{m}^{T}$ Down dating with search directions
Distance between posterior and solution

$$
\left[W_{\boldsymbol{A}}\left(\mathcal{G}_{m}, \boldsymbol{x}_{*}\right)\right]^{2}=\left\|\boldsymbol{x}_{m}-\boldsymbol{x}_{*}\right\|_{\boldsymbol{A}}^{2}+(d-m)
$$

Error estimation with $S$ statistic

$$
\left\|\boldsymbol{x}_{m}-\boldsymbol{x}_{*}\right\|_{\boldsymbol{A}}^{2} \approx\left\|\boldsymbol{x}_{m}-X\right\|_{\boldsymbol{A}}^{2} \quad \text { where } \quad X \sim \mathcal{G}_{m}
$$

Experiments: $\boldsymbol{A} \boldsymbol{x}_{*}=\boldsymbol{b}$
Dimension $d=100, \boldsymbol{x}_{*}=\mathbb{1}, \kappa_{2}(\boldsymbol{A}) \approx 4 \cdot 10^{4}$
$S$ statistic samples per iteration: 10

## BayesCG under Inverse Prior Absolute Error, Two-norm Bound, S Statistic



Squared absolute error $\left\|\boldsymbol{x}_{m}-\boldsymbol{x}_{*}\right\|_{\boldsymbol{A}}^{2}$
S statistic samples
Two-norm bound squared

## BayesCG under Inverse Prior Preliminary Conclusions

## Pros

- Iterates and convergence identical to that of CG
- Samples from S statistic more accurate than norm-wise forward error bounds

Cons

- Preserving semi-definiteness of posterior during down dates $\Sigma_{m}=\Sigma_{m-1}-\boldsymbol{s}_{m} \boldsymbol{s}_{m}^{T}$
- Prior is not practical


## Existing Work on A-Norm Errors in CG

Estimation approaches:
Ritz values, Gauss-Radau rules, incremental norm estimation, maximal attainable accuracy estimates

Golub, Strakoš 1994<br>Greenbaum 1997<br>Meurant 1998<br>Golub, Meurant 1997<br>Strakoš, Tichý 2002<br>Strakoš, Tichý 2005<br>Meurant, Tichý 2013<br>Liesen, Strakoš 2016<br>Meurant, Tichý 2019

Probabilistic numerical solvers:
Designed to capture computational uncertainty

## 5b. Priors that Reproduce CG Krylov Priors

## Krylov Priors $\Gamma_{0}=\boldsymbol{V} \Phi \boldsymbol{V}^{T}$

Columns of $\boldsymbol{V}$ are $\boldsymbol{A}$-orthonormal basis
$\left\{\boldsymbol{V}^{T} \boldsymbol{A} \boldsymbol{V}=\boldsymbol{I}_{n}\right.$, normalized CG search directions $\}$
for Krylov space of maximal dimension $n \leq d$

$$
\mathcal{K}_{n}\left(\boldsymbol{A}, \boldsymbol{r}_{0}\right)=\operatorname{span}\left\{\boldsymbol{r}_{0}, \boldsymbol{A} \boldsymbol{r}_{0}, \ldots, \boldsymbol{A}^{n-1} \boldsymbol{r}_{0}\right\}
$$

$\Phi$ is any diagonal matrix $\Phi$ with positive diagonal
Advantages

- BayesCG under Krylov prior $=C G+$ posteriors
- Trailing submatrices: No explicit downdating of posteriors

$$
\Gamma_{m}=\boldsymbol{V}_{m+1: n} \Phi_{m+1: n, m+1: n} \boldsymbol{V}_{m+1: n}^{T}
$$

$\Gamma_{0}$ singular for $n<d$
$\Longrightarrow$ cannot use $\Gamma_{0}$ to define norm for Wasserstein distance

## Wasserstein Metric for BayesCG under Specific Krylov Prior

Specific Krylov prior: $\Gamma_{0}=\boldsymbol{V} \Phi \boldsymbol{V}^{\top}$ has diagonal elements

$$
\Phi_{m m}=\left(\boldsymbol{v}_{m}^{T} \boldsymbol{r}_{0}\right)^{2} \quad 1 \leq m \leq n
$$

\{Computed automatically in our CG implementation\}

- A-norm error with specific Krylov prior

$$
\left\|\boldsymbol{x}_{m}-\boldsymbol{x}_{*}\right\|_{\boldsymbol{A}}^{2}=\operatorname{trace}\left(\boldsymbol{A} \Gamma_{m}\right)
$$

- Wasserstein $\boldsymbol{A}$-norm metric for BayesCG under $\Gamma_{0}$

$$
\left[W_{\boldsymbol{A}}\left(\mathcal{G}_{m}, \boldsymbol{x}_{*}\right)\right]^{2}=\left\|\boldsymbol{x}_{m}-\boldsymbol{x}_{*}\right\|_{\boldsymbol{A}}^{2}+\operatorname{trace}\left(\boldsymbol{A} \Gamma_{m}\right)
$$

Wasserstein $\boldsymbol{A}$-norm metric $\triangleq \boldsymbol{A}$-norm error in iterate

$$
\left[W_{\boldsymbol{A}}\left(\mathcal{G}_{m}, \boldsymbol{x}_{*}\right)\right]^{2}=2\left\|\boldsymbol{x}_{m}-\boldsymbol{x}_{*}\right\|_{\boldsymbol{A}}^{2}
$$

## S Statistic for Specific Krylov Prior

Krylov posterior $\mathcal{G}_{m}=\mathcal{N}\left(\boldsymbol{x}_{m}, \Gamma_{m}\right)$

- Distance of posterior to solution $\triangleq$ absolute error

$$
\left[W_{\boldsymbol{A}}\left(\mathcal{G}_{m}, \boldsymbol{x}_{*}\right)\right]^{2}=2\left\|\boldsymbol{x}_{m}-\boldsymbol{x}_{*}\right\|_{\boldsymbol{A}}^{2}
$$

- Absolute error $\triangleq$ mean of $S$ statistic

$$
2\left\|\boldsymbol{x}_{m}-\boldsymbol{x}_{*}\right\|_{\boldsymbol{A}}^{2}=\mathbb{E}_{X \sim \mathcal{G}_{m}}\left[\left\|\boldsymbol{x}_{m}-X\right\|_{\boldsymbol{A}}^{2}\right]
$$

- Empirical mean of $S$ statistic from 10 samples

$$
\mathbb{E}_{X \sim \mathcal{G}_{m}}\left[\left\|\boldsymbol{x}_{m}-X\right\|_{\boldsymbol{A}}^{2}\right] \approx \frac{1}{10} \sum_{j=1}^{10} \frac{1}{2}\left\|\boldsymbol{x}_{m}-X_{j}\right\|_{\boldsymbol{A}}^{2} \quad X_{j} \sim \mathcal{G}_{m}
$$

Experiments: $\boldsymbol{A} \boldsymbol{x}_{*}=\boldsymbol{b}$
Dimension $d=100, \boldsymbol{x}_{*}=\mathbb{1}$, same $\boldsymbol{A}$ as before
Condition number $\kappa_{2}(\boldsymbol{A}) \approx 4 \cdot 10^{4}$
$S$ statistic samples per iteration: 10

## BayesCG under Specific Krylov Prior Empirical Mean of S Statistic $\approx$ Absolute Error



Squared absolute error $\left\|\boldsymbol{x}_{m}-\boldsymbol{x}_{*}\right\|_{\boldsymbol{A}}^{2}$
Empirical mean of $S$ statistic from 10 samples

## BayesCG under Specific Krylov Prior Relative Error, Bounds, S Statistic



Exact error $\left\|\boldsymbol{x}_{m}-\boldsymbol{x}_{*}\right\|_{\boldsymbol{A}} /\left\|\boldsymbol{x}_{*}\right\|_{\boldsymbol{A}}$
S statistic samples (square-root, normalized) $\approx$ exact error CG bound, and general 2-norm error bound

## BayesCG under Specific Krylov Prior Preliminary Conclusions

## Pros:

- Errors in iterates $\triangleq$ distances of posteriors to solution
- Errors in iterates $\triangleq$ mean of $S$ statistic
- S statistic produces exact magnitude of errors

Cons: BayesCG under Krylov priors requires
(1) Running maximal number $n$ of CG iterations
(2) Storing all $n$ search directions $V$ (for factored form $\Gamma_{0}=\boldsymbol{V} \Phi \boldsymbol{V}^{\top}$ )
(3) Matvecs with $\Gamma_{m} \in \mathbb{R}^{d \times(n-m+1)}$ in iteration $m$

Again, this is not practical
Next: Low-rank approximations of Krylov priors are promising
6. Low-Rank Approximations of Krylov Priors

## Low-Rank Approximation via Look-Ahead

- Run $k$ iterations of CG
- Look-ahead: Run $\ell=5$ more CG iterations
- Prior has rank $k+\ell=k+5$
- Estimate errors $e_{m}=\left\|\boldsymbol{x}_{m}-\boldsymbol{x}_{*}\right\|_{\boldsymbol{A}}^{2}$ at iteration $1 \leq m \leq k$ with empirical mean of 5 samples from $S$ statistic

$$
\sigma_{m}=\frac{1}{5} \sum_{j=1}^{5} \frac{1}{2}\left\|\boldsymbol{x}_{m}-X_{m j}\right\|_{\boldsymbol{A}}^{2} \quad X_{m j} \sim \mathcal{N}\left(\boldsymbol{x}_{m}, \hat{\Gamma}_{m}\right)
$$

Error of empirical mean in last iteration is $\left|\sigma_{k}-e_{k}\right| / e_{k}$

3 numerical examples: $\boldsymbol{A} \boldsymbol{x}_{*}=\boldsymbol{b}$ of dimension $d$
$\boldsymbol{A}=\boldsymbol{Q} \wedge \boldsymbol{Q}^{T}$ with $\boldsymbol{Q}$ random orthogonal [Stewart 1980]
Eigenvalue distributions from [Liesen, Strakoš 2013]

## Uniformly Distributed Eigenvalues, Dimension $d=10^{3}$ $k=80$ Iterations, Prior of Rank 85

$$
\kappa_{2}(\boldsymbol{A})=10^{5}, \quad \text { Eigenvalues } \lambda_{j}=1+\frac{j-1}{d-1}\left(10^{5}-1\right), \quad 1 \leq j \leq d
$$



Squared absolute error $\left\|\boldsymbol{x}_{m}-\boldsymbol{x}_{*}\right\|_{\boldsymbol{A}}^{2}$
Empirical mean of $S$ statistic, mean error in last iteration $\leq .72$

## Large Eigenvalue Cluster, Dimension $d=10^{3}$ $k=30$ Iterations, Prior of Rank 35

$$
\kappa_{2}(\boldsymbol{A})=10^{5}, \quad \text { Eigenvalues } \lambda_{j}=1+\frac{j-1}{d-1}\left(10^{5}-1\right) 0.65^{d-j}
$$



Squared absolute error $\left\|\boldsymbol{x}_{m}-\boldsymbol{x}_{*}\right\|_{\boldsymbol{A}}^{2}$
Empirical mean of $S$ statistic, mean error in last iteration $\leq .18$

## Smoothly Increasing Eigenvalues, Dimension $d=100$ $k=12$ Iterations, Prior of Rank 17

$\kappa_{2}(\boldsymbol{A}) \approx 7 \cdot 10^{16}, \quad$ Eigenvalues $\lambda_{j}=\log (j) / \log (d), \quad 1 \leq j \leq d$


Squared absolute error $\left\|\boldsymbol{x}_{m}-\boldsymbol{x}_{*}\right\|_{\boldsymbol{A}}^{2}$
Empirical mean of $S$ statistic, mean error in last iteration $\leq .2$

## Relative Errors for Previous Problem



Relative error $\left\|\boldsymbol{x}_{m}-\boldsymbol{x}_{*}\right\|_{\boldsymbol{A}} /\left\|\boldsymbol{x}_{*}\right\|_{\boldsymbol{A}}$, empirical mean of S statistic, 2-norm bound
Last iteration
Two-norm bound: No accuracy $\left(\kappa_{2}(\boldsymbol{A}) \approx 7 \cdot 10^{16}\right)$
Relative error: $\left\|\boldsymbol{x}_{12}-\boldsymbol{x}_{*}\right\|_{\boldsymbol{A}} /\left\|\boldsymbol{x}_{*}\right\|_{\boldsymbol{A}} \approx 1.5405 \cdot 10^{-5}$
S statistic empirical mean: $\sigma_{12} \approx 1.1956 \cdot 10^{-5}$

## Accuracy of Low-Rank Prior Approximations

Experiments: Error estimates have correct magnitude especially where traditional bounds are not informative

Exact prior $\mathcal{G}_{0}=\mathcal{N}\left(x_{0}, \Gamma_{0}\right)$

$$
\Gamma_{0}=\boldsymbol{V} \Phi \boldsymbol{V}^{T} \quad \operatorname{rank}\left(\Gamma_{0}\right)=n
$$

Rank $k+\ell$ approximation $\hat{\mathcal{G}}_{0}=\mathcal{N}\left(\boldsymbol{x}_{0}, \hat{\Gamma}_{0}\right)$

$$
\hat{\Gamma}_{0}=\boldsymbol{V}_{1: k+\ell} \Phi_{1: k+\ell, 1: k+\ell} \boldsymbol{V}_{1: k+\ell}^{T} \quad \operatorname{rank}\left(\hat{\Gamma}_{0}\right)=k+\ell
$$

Distance between low-rank and exact distributions

$$
\left[W_{\boldsymbol{A}}\left(\mathcal{G}_{m}, \hat{\mathcal{G}}_{m}\right)\right]^{2}=\left\|\boldsymbol{V}_{k+\ell+1: n}^{T} \boldsymbol{r}_{0}\right\|_{2}^{2} \quad 0 \leq m \leq k
$$

Distance $\approx$ contribution of $\boldsymbol{r}_{0}$ in search directions ignored by $\hat{\Gamma}_{0}$

## Practical Procedure: Error Estimate for Last Iterate

(1) Run CG until convergence, at some iteration $k$
(2) Run $\ell$ more CG iterations, store $\ell$ search directions $\boldsymbol{V}_{k+1: k+\ell}$
(3) Normalize factor for low-rank posterior $\hat{\Gamma}_{k}=\boldsymbol{L}_{k} \boldsymbol{L}_{k}^{T}$

$$
\boldsymbol{L}_{k}=\boldsymbol{V}_{k+1: k+\ell} \Phi_{k+1: k+\ell}^{1 / 2} \in \mathbb{R}^{d \times \ell}
$$

(9) S statistic samples $\sigma_{j}=\left\|\boldsymbol{x}_{k}-X_{j}\right\|_{\boldsymbol{A}}^{2}, \quad 1 \leq j \leq n_{s}$

$$
X_{j}=\boldsymbol{x}_{k}+\boldsymbol{L}_{k} \operatorname{randn}(\ell, 1)
$$

(5) Error estimate $=$ empirical mean of $S$ statistic samples

$$
\frac{1}{n_{s}} \sum_{j=1}^{n_{s}} \frac{1}{2} \sigma_{j} \approx\left\|x_{k}-x_{*}\right\|_{\boldsymbol{A}}^{2}
$$

Work in addition to $C G$ : $\ell$ iterations, $c n_{s}$ matvecs Additional storage: $\ell$ search directions

## Take-Home Message

Solution of $\boldsymbol{A} \boldsymbol{x}_{*}=\boldsymbol{b}$ with real symmetric positive definite matrix $\boldsymbol{A}$

- BayesCG: Probabilistic extension of Conjugate Gradient (CG) BayesCG under low-rank Krylov prior $=$ CG
- Accurate error estimates with a few more iterations \& matvecs S statistic: Produces error estimates of correct magnitude especially where traditional bounds are not informative
- Distance of posterior to solution $\triangleq$ absolute error in iterate

Not discussed here

- Numerical implementation
(reducing re-orthogonalizations, numerical accuracy of posteriors)
- Rigorous \& meaningful statistical setting (nonlinear dependence of CG on solution, degenerate distributions)


## Future Work

- Systematic low-rank approximation of priors
- Error bars/variance for S statistic and other statistics
- Hardwiring termination criteria into posteriors
- Probabilistic numerical Krylov solvers for indefinite/non-symmetric systems, and least squares
- Priors and posteriors designed for capturing numerical uncertainty (roundoff)
- Disintegration: Non-Gaussian distributions


[^0]:    ${ }^{1}$ Different from Krylov prior in [Cockayne, Oates, Ipsen, Girolami 2019]

