

---

# Probabilistic Linear Solvers for Machine Learning

---

Jonathan Wenger    Philipp Hennig

University of Tübingen

Max Planck Institute for Intelligent Systems

Tübingen, Germany

{jonathan.wenger, philipp.hennig}@uni-tuebingen.de

## Abstract

Linear systems are the bedrock of virtually all numerical computation. Machine learning poses specific challenges for the solution of such systems due to their scale, characteristic structure, stochasticity and the central role of uncertainty in the field. Unifying earlier work we propose a class of probabilistic linear solvers which jointly infer the matrix, its inverse and the solution from matrix-vector product observations. This class emerges from a fundamental set of desiderata which constrains the space of possible algorithms and recovers the method of conjugate gradients under certain conditions. We demonstrate how to incorporate prior spectral information in order to calibrate uncertainty and experimentally showcase the potential of such solvers for machine learning.

## 1 Introduction

Arguably one of the most fundamental problems in machine learning, statistics and scientific computation at large is the solution of linear systems of the form  $\mathbf{A}\mathbf{x}_* = \mathbf{b}$ , where  $\mathbf{A} \in \mathbb{R}_{\text{sym}}^{n \times n}$  is a symmetric positive definite matrix [1–3]. Such matrices usually arise in the context of second-order or quadratic optimization problems and as Gram matrices. Some of the numerous application areas in machine learning and related fields are least-squares regression [4], kernel methods [5], Kalman filtering [6], Gaussian (process) inference [7], spectral graph theory [8], (linear) differential equations [9] and (stochastic) second-order methods [10].

Linear systems in machine learning are typically large-scale, have characteristic structure arising from generative processes, and are subject to noise. These distinctive features call for linear solvers that can explicitly make use of such structural information. While classic solvers are highly optimized for general problems, they lack key functionality for machine learning. In particular, they do not consider generative prior information about the matrix.

An important example are kernel Gram matrices, which exhibit specific sparsity structure and spectral properties, depending on the kernel choice and the generative process of the data. Exploiting such prior information is a prime application for probabilistic linear solvers, which aim to quantify numerical uncertainty arising from limited computational resources. Another key challenge, which we will not yet address here, are noisy matrix evaluations arising from data subsampling. Ultimately, linear algebra for machine learning should integrate all sources of uncertainty in a computational pipeline – aleatoric, epistemic and numerical – into one coherent probabilistic framework.

**Contribution** This paper sets forth desiderata for probabilistic linear solvers which establish first principles for such methods. From these, we derive an algorithm incorporating prior information on the matrix  $\mathbf{A}$  or its inverse  $\mathbf{A}^{-1}$ , which jointly estimates both via repeated application of  $\mathbf{A}$ . This results in posterior beliefs over the two operators and the solution which quantify numerical uncertainty. Our approach unifies and extends earlier formulations and constitutes a new way of




Figure 1: *Illustration of a probabilistic linear solver.* Given a prior for  $\mathbf{A}$  or  $\mathbf{H}$  modelling the linear operator  $\mathbf{A}$  and its inverse  $\mathbf{A}^{-1}$ , posterior beliefs are inferred via observations  $\mathbf{y}_i = \mathbf{A}\mathbf{s}_i$ . This induces a distribution on the solution  $\mathbf{x}_*$ , quantifying numerical uncertainty arising from finite computation. The plot shows  $k = 3$  iterations of Algorithm 1 on a toy problem of dimension  $n = 5$ .

interpreting linear solvers. Further, we propose a prior covariance class which recovers the method of conjugate gradients as its posterior mean and uses prior spectral information for uncertainty calibration, one of the primary shortcomings of probabilistic linear solvers. We conclude by presenting simplified examples of promising applications of such solvers within machine learning.

## 2 Probabilistic Linear Solvers

Let  $\mathbf{A}\mathbf{x}_* = \mathbf{b}$  be a linear system with  $\mathbf{A} \in \mathbb{R}_{\text{sym}}^{n \times n}$  positive definite and  $\mathbf{b} \in \mathbb{R}^n$ . *Probabilistic linear solvers* (PLS) [11–13] iteratively build a model for the linear operator  $\mathbf{A}$ , its inverse  $\mathbf{H} = \mathbf{A}^{-1}$  or the solution  $\mathbf{x}_*$ , represented by random variables  $\mathbf{A}$ ,  $\mathbf{H}$  or  $\mathbf{x}$ . In the framework of probabilistic numerics [14, 15] such solvers can be seen as Bayesian agents performing *inference* via linear *observations*  $\mathbf{Y} = [\mathbf{y}_1, \dots, \mathbf{y}_k] \in \mathbb{R}^{n \times k}$  resulting from *actions*  $\mathbf{S} = [\mathbf{s}_1, \dots, \mathbf{s}_k] \in \mathbb{R}^{n \times k}$  given by an internal *policy*  $\pi(\mathbf{s} \mid \mathbf{A}, \mathbf{H}, \mathbf{x}, \mathbf{A}, \mathbf{b})$ . For a matrix-variate prior  $p(\mathbf{A})$  or  $p(\mathbf{H})$  encoding prior (generative) information, our solver computes posterior beliefs over the matrix, its inverse and the solution of the linear system. An illustration of a probabilistic linear solver is given in Figure 1.

**Desiderata** We begin by stipulating a fundamental set of desiderata for probabilistic linear solvers. To our knowledge such a list has not been collated before. Connecting previously disjoint threads, the following presents a roadmap for the development of these methods. Probabilistic linear solvers modelling  $\mathbf{A}$  and  $\mathbf{A}^{-1}$  must assume matrix-variate distributions which are expressive enough to capture structure and generative prior information either for  $\mathbf{A}$  or its inverse. The distribution choice must also allow computationally efficient sampling and density evaluation. It should encode symmetry and positive definiteness and must be closed under positive linear combinations. Further, the two models for the system matrix or its inverse should be translatable into and consistent with each other. Actions  $\mathbf{s}_i$  of a PLS should be model-based and induce a tractable distribution on linear observations  $\mathbf{y}_i = \mathbf{A}\mathbf{s}_i$ . Since probabilistic linear solvers are low-level procedures, their inference procedure must be computationally lightweight. Given (noise-corrupted) observations this requires tractable posteriors over  $\mathbf{A}$ ,  $\mathbf{H}$  and  $\mathbf{x}$ , which are calibrated in the sense that at convergence the true solution  $\mathbf{x}_*$  represents a draw from the posterior  $p(\mathbf{x} \mid \mathbf{Y}, \mathbf{S})$ . Finally, such solvers need to allow preconditioning of the problem and ideally should return beliefs over non-linear properties of the system matrix extending the functionality of classic methods. These desiderata are summarized concisely in Table 1.

### 2.1 Bayesian Inference Framework

Guided by these desiderata, we will now outline the inference framework for  $\mathbf{A}$ ,  $\mathbf{H}$  and  $\mathbf{x}$  forming the base of the algorithm. The choice of a matrix-variate prior distribution is severely limited by the desideratum that conditioning on linear observations  $\mathbf{y}_i = \mathbf{A}\mathbf{s}_i$  must be tractable. This reduces the choice to stable distributions [16] and thus excludes candidates such as the Wishart, which has measure zero outside the cone of symmetric positive semi-definite matrices. For symmetric matrices, this essentially forces use of the symmetric matrix-variate normal distribution, introduced in this context by Hennig [11]. Given  $\mathbf{A}_0, \mathbf{W}_0^{\mathbf{A}} \in \mathbb{R}_{\text{sym}}^{n \times n}$ , assume a prior distribution

$$p(\mathbf{A}) = \mathcal{N}(\mathbf{A}; \mathbf{A}_0, \mathbf{W}_0^{\mathbf{A}} \otimes \mathbf{W}_0^{\mathbf{A}}),$$

Table 1: *Desired properties of probabilistic linear solvers.* Symbols ( $\times$ ,  $\sim$ ,  $\checkmark$ ) indicate which properties are encoded in our proposed solver (see Algorithm 1) and to what degree.

No.	Property	Formulation	
(1)	distribution over matrices	$\mathbf{A} \sim \mathcal{D}, p_{\mathcal{D}}(\mathbf{A})$	$\checkmark$
(2)	symmetry	$\mathbf{A} = \mathbf{A}^{\top}$ a.s.	$\checkmark$
(3)	positive definiteness	$\forall \mathbf{v} \neq 0 : \mathbf{v}^{\top} \mathbf{A} \mathbf{v} > 0$ a.s.	$\sim$
(4)	positive linear combination in same distribution family	$\forall \alpha_j > 0 : \sum_j \alpha_j \mathbf{A}_j \sim \mathcal{D}$	$\checkmark$
(5)	corresponding priors on the matrix and its inverse	$p(\mathbf{A}) \longleftrightarrow p(\mathbf{H})$	$\checkmark$
(6)	model-based policy	$s_i \sim \pi(s   \mathbf{A}, \mathbf{b}, \mathbf{A}, \mathbf{H}, \mathbf{x})$	$\checkmark$
(7)	matrix-vector product in tractable distribution family	$\mathbf{A} \mathbf{s} \sim \mathcal{D}'$	$\checkmark$
(8)	noisy observations	$p(\mathbf{Y}   \mathbf{A}, \mathbf{S}) = \mathcal{N}(\mathbf{Y}; \mathbf{A} \mathbf{S}, \Lambda)$	$\times$
(9)	tractable posterior	$p(\mathbf{A}   \mathbf{Y}, \mathbf{S})$ or $p(\mathbf{H}   \mathbf{Y}, \mathbf{S})$	$\checkmark$
(10)	calibrated uncertainty	$\mathbf{x}_* \sim \mathcal{N}(\mathbb{E}[\mathbf{x}], \text{Cov}[\mathbf{x}])$	$\sim$
(11)	preconditioning	$(\mathbf{P}^{-\top} \mathbf{A} \mathbf{P}^{-1}) \mathbf{P} \mathbf{x}_* = \mathbf{P}^{-\top} \mathbf{b}$	$\checkmark$
(12)	distributions over non-linear derived quantities of $\mathbf{A}$	$\det(\mathbf{A}), \sigma(\mathbf{A}), \mathbf{A} = \mathbf{L}^{\top} \mathbf{L}, \dots$	$\times$

where  $\otimes$  denotes the symmetric Kronecker product [17].<sup>1</sup> The symmetric matrix-variate Gaussian induces a Gaussian distribution on linear observations. While it has non-zero measure only for symmetric matrices, its support is not the positive definite cone. However, positive definiteness can still be enforced post-hoc (see Proposition 1). We assume noise-free linear observations of the form  $\mathbf{y}_i = \mathbf{A} \mathbf{s}_i$ , leading to a Dirac likelihood

$$p(\mathbf{Y} | \mathbf{A}, \mathbf{S}) = \lim_{\varepsilon \downarrow 0} \mathcal{N}(\mathbf{Y}; \mathbf{A} \mathbf{S}, \varepsilon^2 \mathbf{I} \otimes \mathbf{I}) = \delta(\mathbf{Y} - \mathbf{A} \mathbf{S}).$$

The posterior distribution follows from the properties of Gaussians [4] and has been investigated in detail in previous work [18, 11, 13]. It is given by  $p(\mathbf{A} | \mathbf{S}, \mathbf{Y}) = \mathcal{N}(\mathbf{A}; \mathbf{A}_k, \Sigma_k)$  with

$$\begin{aligned} \mathbf{A}_k &= \mathbf{A}_0 + \Delta_0^{\mathbf{A}} \mathbf{U}^{\top} + \mathbf{U} (\Delta_0^{\mathbf{A}})^{\top} - \mathbf{U} \mathbf{S}^{\top} \Delta_0^{\mathbf{A}} \mathbf{U}^{\top} \\ \Sigma_k &= \mathbf{W}_0^{\mathbf{A}} (\mathbf{I}_n - \mathbf{S} \mathbf{U}^{\top}) \otimes \mathbf{W}_0^{\mathbf{A}} (\mathbf{I}_n - \mathbf{S} \mathbf{U}^{\top}) \end{aligned}$$

where  $\Delta_0^{\mathbf{A}} = \mathbf{Y} - \mathbf{A}_0 \mathbf{S}$  and  $\mathbf{U} = \mathbf{W}_0^{\mathbf{A}} \mathbf{S} (\mathbf{S}^{\top} \mathbf{W}_0^{\mathbf{A}} \mathbf{S})^{-1}$ . We aim to construct a probabilistic model  $\mathbf{H}$  for the inverse  $\mathbf{H} = \mathbf{A}^{-1}$  consistent with the model  $\mathbf{A}$  as well. However, not even in the scalar case does the inverse of a Gaussian have finite mean. We ask instead what Gaussian model for  $\mathbf{H}$  is as consistent as possible with our observational model for  $\mathbf{A}$ . For a prior of the form  $p(\mathbf{H}) = \mathcal{N}(\mathbf{H}; \mathbf{H}_0, \mathbf{W}_0^{\mathbf{H}} \otimes \mathbf{W}_0^{\mathbf{H}})$  and likelihood  $p(\mathbf{S} | \mathbf{H}, \mathbf{Y}) = \delta(\mathbf{S} - \mathbf{H} \mathbf{Y})$ , we analogously to the  $\mathbf{A}$ -model obtain a posterior distribution  $p(\mathbf{H} | \mathbf{S}, \mathbf{Y}) = \mathcal{N}(\mathbf{H}; \mathbf{H}_k, \Sigma_k^{\mathbf{H}})$  with

$$\begin{aligned} \mathbf{H}_k &= \mathbf{H}_0 + \Delta_0^{\mathbf{H}} (\mathbf{U}^{\mathbf{H}})^{\top} + \mathbf{U}^{\mathbf{H}} (\Delta_0^{\mathbf{H}})^{\top} - \mathbf{U}^{\mathbf{H}} \mathbf{Y}^{\top} \Delta_0^{\mathbf{H}} (\mathbf{U}^{\mathbf{H}})^{\top} \\ \Sigma_k^{\mathbf{H}} &= \mathbf{W}_0^{\mathbf{H}} (\mathbf{I}_n - \mathbf{Y} (\mathbf{U}^{\mathbf{H}})^{\top}) \otimes \mathbf{W}_0^{\mathbf{H}} (\mathbf{I}_n - \mathbf{Y} (\mathbf{U}^{\mathbf{H}})^{\top}) \end{aligned}$$

where  $\Delta_0^{\mathbf{H}} = \mathbf{S} - \mathbf{H}_0 \mathbf{Y}$  and  $\mathbf{U}^{\mathbf{H}} = \mathbf{W}_0^{\mathbf{H}} \mathbf{Y} (\mathbf{Y}^{\top} \mathbf{W}_0^{\mathbf{H}} \mathbf{Y})^{-1}$ . In Section 3 we will derive a covariance class, which establishes correspondence between the two Gaussian viewpoints for the linear operator and its inverse and is consistent with our desiderata.

## 2.2 Algorithm

The above inference procedure leads to Algorithm 1. The degree to which the desiderata are encoded in our formulation of a PLS can be found in Table 1. We will now go into more detail about the policy, the choice of step size, stopping criteria and the implementation.

**Policy and Step Size** In each iteration our solver collects information about the linear operator  $\mathbf{A}$  via actions  $s_i$  determined by the policy  $\pi(s | \mathbf{A}, \mathbf{H}, \mathbf{x}, \mathbf{A}, \mathbf{b})$ . The next action  $s_i = -\mathbb{E}[\mathbf{H}] \mathbf{r}_{i-1}$  is

<sup>1</sup>See Sections S2 and S3 of the supplementary material for more detail on Kronecker-type products and matrix-variate normal distributions.

---

Algorithm 1: Probabilistic Linear Solver with Uncertainty Calibration

---

```

1 procedure PROBLINSOLVE( $\mathbf{A}(\cdot)$ ,  $\mathbf{b}$ ,  $\mathbf{A}$ ,  $\mathbf{H}$ )                                # prior for  $\mathbf{A}$  or  $\mathbf{H}$ 
2    $\mathbf{x}_0 \leftarrow \mathbb{E}[\mathbf{H}]\mathbf{b}$                                                     # initial guess
3    $\mathbf{r}_0 \leftarrow \mathbf{A}\mathbf{x}_0 - \mathbf{b}$ 
4   while  $\min(\sqrt{\text{tr}(\text{Cov}[\mathbf{x}]}), \|\mathbf{r}_i\|_2) > \max(\delta_{\text{rtol}}\|\mathbf{b}\|_2, \delta_{\text{atol}})$  do      # stopping criteria
5      $\mathbf{s}_i \leftarrow -\mathbb{E}[\mathbf{H}]\mathbf{r}_{i-1}$                                           # compute action via policy
6      $\mathbf{y}_i \leftarrow \mathbf{A}\mathbf{s}_i$                                                 # make observation
7      $\alpha_i \leftarrow -\mathbf{s}_i^\top \mathbf{r}_{i-1} (\mathbf{s}_i^\top \mathbf{y}_i)^{-1}$                     # optimal step size
8      $\mathbf{x}_i \leftarrow \mathbf{x}_{i-1} + \alpha_i \mathbf{s}_i$                                   # update solution estimate
9      $\mathbf{r}_i \leftarrow \mathbf{r}_{i-1} + \alpha_i \mathbf{y}_i$                                 # update residual
10     $\mathbf{A} \leftarrow \text{INFER}(\mathbf{A}, \mathbf{s}_i, \mathbf{y}_i)$                                 # infer posterior distributions
11     $\mathbf{H} \leftarrow \text{INFER}(\mathbf{H}, \mathbf{s}_i, \mathbf{y}_i)$                                 # (see Section 2.1)
12     $\Phi, \Psi \leftarrow \text{CALIBRATE}(\mathbf{S}, \mathbf{Y})$                                 # calibrate uncertainty
13     $\mathbf{x} \leftarrow \mathcal{N}(\mathbf{x}_k, \text{Cov}[\mathbf{H}\mathbf{b}])$                                 # belief over solution
14  return  $(\mathbf{x}, \mathbf{A}, \mathbf{H})$ 

```

---

chosen based on the current belief about the inverse. If  $\mathbb{E}[\mathbf{H}] = \mathbf{A}^{-1}$ , i.e. if the solver’s estimate for the inverse equals the true inverse, then Algorithm 1 converges in a single step since

$$\mathbf{x}_{i-1} + \mathbf{s}_i = \mathbf{x}_{i-1} - \mathbb{E}[\mathbf{H}]\mathbf{r}_{i-1} = \mathbf{x}_{i-1} - \mathbf{A}^{-1}(\mathbf{A}\mathbf{x}_{i-1} - \mathbf{b}) = \mathbf{A}^{-1}\mathbf{b} = \mathbf{x}_*.$$

The step size minimizing the quadratic  $q(\mathbf{x}_i + \alpha\mathbf{s}_i) = \frac{1}{2}(\mathbf{x}_i + \alpha\mathbf{s}_i)^\top \mathbf{A}(\mathbf{x}_i + \alpha\mathbf{s}_i) - \mathbf{b}^\top(\mathbf{x}_i + \alpha\mathbf{s}_i)$  along the action  $\mathbf{s}_i$  is given by  $\alpha_i = \arg \min_{\alpha} q(\mathbf{x}_i + \alpha\mathbf{s}_i) = \mathbf{s}_i^\top (\mathbf{b} - \mathbf{A}\mathbf{x}_i) (\mathbf{s}_i^\top \mathbf{A}\mathbf{s}_i)^{-1}$ .

**Stopping Criteria** Classic linear solvers typically use stopping criteria based on the current residual of the form  $\|\mathbf{A}\mathbf{x}_i - \mathbf{b}\|_2 \leq \max(\delta_{\text{rtol}}\|\mathbf{b}\|_2, \delta_{\text{atol}})$  for relative and absolute tolerances  $\delta_{\text{rtol}}$  and  $\delta_{\text{atol}}$ . However, this residual may oscillate or even increase in all but the last step even if the error  $\|\mathbf{x}_* - \mathbf{x}_i\|_2$  is monotonically decreasing [19, 20]. From a probabilistic point of view, we should stop if our posterior uncertainty is sufficiently small. Assuming the posterior covariance is calibrated, it holds that  $(\mathbb{E}_{\mathbf{x}_*} [\|\mathbf{x}_* - \mathbb{E}[\mathbf{x}]\|_2])^2 \leq \mathbb{E}_{\mathbf{x}_*} [\|\mathbf{x}_* - \mathbb{E}[\mathbf{x}]\|_2^2] = \text{tr}(\text{Cov}[\mathbf{x}])$ . Hence given calibration, we can bound the expected (relative) error between our estimate and the true solution by terminating when  $\sqrt{\text{tr}(\text{Cov}[\mathbf{x}])} \leq \max(\delta_{\text{rtol}}\|\mathbf{b}\|_2, \delta_{\text{atol}})$ . A probabilistic criterion is also necessary for an extension to the noisy setting, where classic convergence criteria become stochastic. However, probabilistic linear solvers typically suffer from miscalibration [21], an issue we will address in Section 3.

**Implementation** We provide an open-source implementation of Algorithm 1 as part of **PROBNUM**, a Python package implementing probabilistic numerical methods, in an online code repository:

$$f(\Lambda) \quad \text{https://github.com/probabilistic-numerics/probnum}$$

The mean and covariance up- and downdates in Section 2.1 when performed iteratively are of low rank. In order to maintain numerical stability these updates can instead be performed for their respective Cholesky factors [22]. This also enables computationally efficient sampling or evaluation of probability density functions downstream.

### 2.3 Theoretical Properties

This section details some theoretical properties of our method such as its convergence behavior and computational complexity. In particular we demonstrate that for a specific prior choice Algorithm 1 recovers the method of conjugate gradients as its solution estimate. All proofs of results in this section and the next can be found in the supplementary material. We begin by establishing that our solver is a *conjugate directions method* and therefore converges in at most  $n$  steps in exact arithmetic.

**Theorem 1** (Conjugate Directions Method)

Given a prior  $p(\mathbf{H}) = \mathcal{N}(\mathbf{H}; \mathbf{H}_0, \mathbf{W}_0^{\mathbf{H}} \otimes \mathbf{W}_0^{\mathbf{H}})$  such that  $\mathbf{H}_0, \mathbf{W}_0^{\mathbf{H}} \in \mathbb{R}_{\text{sym}}^{n \times n}$  positive definite, then actions  $\mathbf{s}_i$  of Algorithm 1 are  $\mathbf{A}$ -conjugate, i.e. for  $0 \leq i, j \leq k$  with  $i \neq j$  it holds that  $\mathbf{s}_i^\top \mathbf{A}\mathbf{s}_j = 0$ .

We can obtain a better convergence rate by placing stronger conditions on the prior covariance class as outlined in Section 3. Given these assumptions, Algorithm 1 recovers the iterates of (preconditioned) CG and thus inherits its favorable convergence behavior (overviews in [23, 10]).

**Theorem 2** (Connection to the Conjugate Gradient Method)

Given a scalar prior mean  $\mathbf{A}_0 = \mathbf{H}_0^{-1} = \alpha \mathbf{I}$  with  $\alpha > 0$ , assume (1) and (2) hold, then the iterates  $\mathbf{x}_i$  of Algorithm 1 are identical to the ones produced by the conjugate gradient method.

A common phenomenon observed when implementing conjugate gradient methods is that due to cancellation in the computation of the residuals, the search directions  $\mathbf{s}_i$  lose  $\mathbf{A}$ -conjugacy [24, 25, 3]. In fact, they can become independent up to working precision for  $i$  large enough [25]. One way to combat this is to perform complete reorthogonalization of the search directions in each iteration as originally suggested by Lanczos [26]. Algorithm 1 does this *implicitly* via its choice of policy which depends on all previous search directions as opposed to just  $\mathbf{s}_{i-1}$  for (naive) CG.

**Computational Complexity** The solver has time complexity  $\mathcal{O}(kn^2)$  for  $k$  iterations without uncertainty calibration. Compared to CG, inferring the posteriors in Section 2.1 adds an overhead of four outer products and four matrix-vector products per iteration, given (1) and (2). Uncertainty calibration outlined in Section 3 adds between  $\mathcal{O}(1)$  and  $\mathcal{O}(k^3)$  per iteration depending on the sophistication of the scheme. Already for moderate  $n$  this is dominated by the iteration cost. In practice, means and covariances do not need to be formed in memory. Instead they can be evaluated lazily as linear operators  $\mathbf{v} \mapsto \mathbf{L}\mathbf{v}$ , if  $\mathbf{S}$  and  $\mathbf{Y}$  are stored. This results in space complexity  $\mathcal{O}(kn)$ .

**2.4 Related Work**

Numerical methods for the solution of linear systems have been studied in great detail since the last century. Standard texts [1, 2, 10, 3] give an in-depth overview. The conjugate gradient method recovered by our algorithm for a specific choice of prior was introduced by Hestenes and Stiefel [19]. Recently, randomization has been exploited to develop improved algorithms for large-scale problems arising from machine learning [27, 28]. The key difference to our approach is that we do not rely on sampling to approximate large-scale matrices, but instead perform probabilistic inference. Our approach is based on the framework of probabilistic numerics [14, 15] and is a natural continuation of previous work on probabilistic linear solvers. In historical order, Hennig and Kiefel [18] provided a probabilistic interpretation of Quasi-Newton methods, which was expanded upon in [11]. This work also relied on the symmetric matrix-variate Gaussian as used in our paper. Bartels and Hennig [29] estimate numerical error in approximate least-squares solutions by using a probabilistic model. More recently, Cockayne et al. [21] proposed a Bayesian conjugate gradient method performing inference on the solution of the system. This was connected to the matrix-based view by Bartels et al. [13].

**3 Prior Covariance Class**

Having outlined the proposed algorithm, this section derives a prior covariance class which satisfies nearly all desiderata, connects the two modes of prior information and allows for calibration of uncertainty by appropriately choosing remaining degrees of freedom in the covariance. The third desideratum posited that  $\mathbf{A}$  and  $\mathbf{H}$  should be almost surely positive definite. This evidently does not hold for the matrix-variate Gaussian. However, we can restrict the choice of admissible  $\mathbf{W}_0^{\mathbf{A}}$  to act like  $\mathbf{A}$  on  $\text{span}(\mathbf{S})$ . This in turn induces a positive definite posterior mean.

**Proposition 1** (Hereditary Positive Definiteness [30, 18])

Let  $\mathbf{A}_0 \in \mathbb{R}_{\text{sym}}^{n \times n}$  be positive definite. Assume the actions  $\mathbf{S}$  are  $\mathbf{A}$ -conjugate and  $\mathbf{W}_0^{\mathbf{A}} \mathbf{S} = \mathbf{Y}$ , then for  $i \in \{0, \dots, k - 1\}$  it holds that  $\mathbf{A}_{i+1}$  is symmetric positive definite.

Prior information about the linear system usually concerns the matrix  $\mathbf{A}$  itself and not its inverse, but the inverse is needed to infer the solution  $\mathbf{x}_*$  of the linear problem. So a way to translate between a Gaussian distribution on  $\mathbf{A}$  and  $\mathbf{H}$  is crucial. Previous works generally committed to either one view or the other, potentially discarding available information. Below, we show that the two correspond, if we allow ourselves to constrain the space of possible models. We impose the following condition.

**Definition 1**

Let  $\mathbf{A}_i$  and  $\mathbf{H}_i$  be the means of  $\mathbf{A}$  and  $\mathbf{H}$  at step  $i$ . We say a prior induces *posterior correspondence* if  $\mathbf{A}_i^{-1} = \mathbf{H}_i$  for all  $0 \leq i \leq k$ . If only  $\mathbf{A}_i^{-1} \mathbf{Y} = \mathbf{H}_i \mathbf{Y}$ , *weak posterior correspondence* holds.

The following theorem establishes a sufficient condition for weak posterior correspondence. For an asymmetric prior model one can establish the stronger notion of posterior correspondence. A proof is included in the supplements.

**Theorem 3** (Weak Posterior Correspondence)

Let  $\mathbf{W}_0^{\mathbf{H}} \in \mathbb{R}_{\text{sym}}^{n \times n}$  be positive definite. Assume  $\mathbf{H}_0 = \mathbf{A}_0^{-1}$ , and that  $\mathbf{W}_0^{\mathbf{A}}, \mathbf{A}_0, \mathbf{W}_0^{\mathbf{H}}$  satisfy

$$\mathbf{W}_0^{\mathbf{A}} \mathbf{S} = \mathbf{Y}, \quad (1)$$

$$\mathbf{S}^{\top} (\mathbf{W}_0^{\mathbf{A}} \mathbf{A}_0^{-1} - \mathbf{A} \mathbf{W}_0^{\mathbf{H}}) = \mathbf{0}, \quad (2)$$

then weak posterior correspondence holds for the symmetric Kronecker covariance.

Given the above, let  $\mathbf{A}_0$  be a symmetric positive definite prior mean and  $\mathbf{H}_0 = \mathbf{A}_0^{-1}$ . Define the orthogonal projection matrices  $\mathbf{P}_{\mathbf{S}^{\perp}} = \mathbf{I} - \mathbf{S}(\mathbf{S}^{\top} \mathbf{S})^{-1} \mathbf{S}^{\top} \in \mathbb{R}_{\text{sym}}^{n \times n}$  and  $\mathbf{P}_{\mathbf{Y}^{\perp}} = \mathbf{I} - \mathbf{Y}(\mathbf{Y}^{\top} \mathbf{Y})^{-1} \mathbf{Y}^{\top} \in \mathbb{R}_{\text{sym}}^{n \times n}$  mapping to the spaces  $\text{span}(\mathbf{S})^{\perp}$  and  $\text{span}(\mathbf{Y})^{\perp}$ . We propose the following prior covariance class given by the prior covariance factors of the  $\mathbf{A}$  and  $\mathbf{H}$  view

$$\begin{aligned} \mathbf{W}_0^{\mathbf{A}} &= \mathbf{A} \mathbf{S} (\mathbf{S}^{\top} \mathbf{A} \mathbf{S})^{-1} \mathbf{S}^{\top} \mathbf{A} + \mathbf{P}_{\mathbf{S}^{\perp}} \mathbf{\Phi} \mathbf{P}_{\mathbf{S}^{\perp}}, \\ \mathbf{W}_0^{\mathbf{H}} &= \mathbf{A}_0^{-1} \mathbf{Y} (\mathbf{Y}^{\top} \mathbf{A}_0^{-1} \mathbf{Y})^{-1} \mathbf{Y}^{\top} \mathbf{A}_0^{-1} + \mathbf{P}_{\mathbf{Y}^{\perp}} \mathbf{\Psi} \mathbf{P}_{\mathbf{Y}^{\perp}}, \end{aligned} \quad (3)$$

where  $\mathbf{\Phi} \in \mathbb{R}^{n \times n}$  and  $\mathbf{\Psi} \in \mathbb{R}^{n \times n}$  are degrees of freedom. This choice of covariance class satisfies Theorem 1, Proposition 1, Theorem 3 and for a scalar mean also Theorem 2. Therefore, it produces symmetric realizations, has symmetric positive semi-definite means, it links the matrix and the inverse view and at any given time only needs access to  $\mathbf{v} \mapsto \mathbf{A} \mathbf{v}$  not  $\mathbf{A}$  itself. It is also compatible with a preconditioner by simply transforming the given linear problem.

This class can be interpreted as follows. The derived covariance factor  $\mathbf{W}_0^{\mathbf{A}}$  acts like  $\mathbf{A}$  on the space  $\text{span}(\mathbf{S})$  explored by the algorithm. On the remaining space its uncertainty is determined by the degrees of freedom in  $\mathbf{\Phi}$ . Likewise, our best guess for  $\mathbf{A}^{-1}$  is  $\mathbf{A}_0^{-1}$  on the space spanned by  $\mathbf{Y}$ . On the orthogonal space  $\text{span}(\mathbf{Y})^{\perp}$  the uncertainty is determined by  $\mathbf{\Psi}$ . Note that the prior depends on actions and observations collected during a run of Algorithm 1, hence one might call this an empirical Bayesian approach. This begs the question how the algorithm is realizable for the proposed prior (3) given its dependence on future data. Notice that the posterior mean in Section 2.1 only depends on  $\mathbf{W}_0^{\mathbf{A}} \mathbf{S} = \mathbf{Y}$  not on  $\mathbf{W}_0^{\mathbf{A}}$  alone. Using eq. (3), at iteration  $i$  we have  $\mathbf{W}_0^{\mathbf{A}} \mathbf{S}_{1:i} = \mathbf{Y}_{1:i}$ , i.e. the observations made up to this point. Similar reasoning applies for the inverse. Now, the posterior covariances do depend on  $\mathbf{W}_0^{\mathbf{A}}$ , respectively  $\mathbf{W}_0^{\mathbf{H}}$  alone, but prior to convergence we only require  $\text{tr}(\text{Cov}[\mathbf{x}])$  for the stopping criterion. We show in Section S4.3 under the assumptions of Theorem 2 how to compute this at any iteration  $i$  independent of future actions and observations.

**Uncertainty Calibration** Generally the actions of Algorithm 1 identify eigenpairs  $(\lambda_i, \mathbf{v}_i)$  in descending order of  $\lambda_i \mathbf{v}_i^{\top} \mathbf{r}_0$  which is a well-known behavior of CG (see eqn. 5.29 in [10]). In part, since this dynamic of the underlying Krylov subspace method is not encoded in the prior, the solver in its current form is typically miscalibrated (see also [21]). While this non-linear information is challenging to include in the Gaussian framework, we can choose  $\mathbf{\Phi}$  and  $\mathbf{\Psi}$  in (3) to empirically calibrate uncertainty. This can be interpreted as a form of hyperparameter optimization similar to optimization of kernel parameters in GP regression.

We would like to encode prior knowledge about the way  $\mathbf{A}$  and  $\mathbf{H}$  act in the respective orthogonal spaces  $\text{span}(\mathbf{S})^{\perp}$  and  $\text{span}(\mathbf{Y})^{\perp}$ . For the Rayleigh quotient  $R(\mathbf{A}, \mathbf{v}) = (\mathbf{v}^{\top} \mathbf{A} \mathbf{v}) / (\mathbf{v}^{\top} \mathbf{v})^{-1}$  it holds that  $\lambda_{\min}(\mathbf{A}) \leq R(\mathbf{A}, \mathbf{v}) \leq \lambda_{\max}(\mathbf{A})$ . Hence for vectors  $\mathbf{v}$  lying in the respective null spaces of  $\mathbf{S}$  and  $\mathbf{Y}$  our uncertainty should be determined by the not yet explored eigenvalues  $\lambda_{k+1}, \dots, \lambda_n$  of  $\mathbf{A}$  and  $\mathbf{H}$ . Without prior information about the eigenspaces, we choose  $\mathbf{\Phi} = \phi \mathbf{I}$  and  $\mathbf{\Psi} = \psi \mathbf{I}$ . If a priori we know the respective spectra, a straightforward choice is

$$\phi = \psi^{-1} = \frac{1}{n - k} \sum_{i=k+1}^n \lambda_i(\mathbf{A}).$$

In the absence of prior spectral information we can make use of already collected quantities during a run of Algorithm 1. We build a one-dimensional regression model  $p(\ln R_i \mid \mathbf{Y}, \mathbf{S})$  for the ln-Rayleigh quotient  $\ln R(\mathbf{A}, \mathbf{s}_i)$  given actions  $\mathbf{s}_i$ . Such a model can then encode the well studied behaviour of CG, whose Rayleigh coefficients rapidly decay at first, followed by a slower continuous




Figure 2: *Rayleigh regression.* Uncertainty calibration via GP regression on  $\{\ln R(\mathbf{A}, \mathbf{s}_i)\}_{i=1}^k$  after  $k = 91$  iterations of Algorithm 1 on an  $n = 1000$  dimensional Matérn32 kernel matrix inversion problem. The degrees of freedom  $\phi = \psi^{-1} > 0$  are set based on the average predicted Rayleigh quotient for the remaining  $n - k = 909$  dimensions.

Table 2: *Uncertainty calibration for kernel matrices.* Monte Carlo estimate  $\bar{w} \approx \mathbb{E}_{\mathbf{x}_*} [w(\mathbf{x}_*)]$  measuring calibration given  $10^5/n$  sampled linear problems of the form  $(\mathbf{K} + \varepsilon^2 \mathbf{I})\mathbf{x}_* = \mathbf{b}$  for each kernel and calibration method. For  $\bar{w} \approx 0$  the solver is well calibrated, for  $\bar{w} \gg 0$  underconfident and for  $\bar{w} \ll 0$  overconfident.

Kernel	$n$	none	Rayleigh	$\varepsilon^2$	$\bar{\lambda}_{k+1:n}$
Matérn32	$10^2$	-5.99	-0.24	0.32	0.09
Matérn32	$10^3$	-1.93	7.53	4.26	4.19
Matérn32	$10^4$	3.87	17.16	8.48	8.47
Matérn52	$10^2$	-7.84	-1.01	-0.76	-0.80
Matérn52	$10^3$	-4.63	1.43	-0.80	-0.81
Matérn52	$10^4$	-4.34	10.81	0.80	0.80
RBF	$10^2$	-7.53	-0.70	-0.84	-0.87
RBF	$10^3$	-4.94	6.60	0.77	0.77
RBF	$10^4$	0.14	21.32	2.92	2.92

decay [10]. Figure 2 illustrates this approach using a GP regression model. At convergence, we use the prediction of the Rayleigh quotient for the remaining  $n - k$  dimensions by choosing

$$\phi = \psi^{-1} = \exp\left(\frac{1}{n-k} \sum_{i=k+1}^n \mathbb{E}[\ln R_i | \mathbf{A}, \mathbf{S}]\right),$$

i.e. uncertainty about actions in  $\text{span}(\mathbf{S})^\perp$  is calibrated to be the average Rayleigh quotient as an approximation to the spectrum. Depending on the application a simple or more complex model may be useful. For large problems, where generally  $k \ll n$ , more sophisticated schemes become computationally feasible. However, these do not necessarily need to be computationally demanding due to the simple nature of this one-dimensional regression problem with few data. For example, approximate [31] or even exact GP regression [32] is possible in  $\mathcal{O}(k)$  using a Kalman filter.

## 4 Experiments

This section demonstrates the functionality of Algorithm 1. We choose some – deliberately simple – example problems from machine learning and scientific computation, where the solver can be used to quantify uncertainty induced by finite computation, solve multiple consecutive linear systems, and propagate information between problems.

**Gaussian Process Regression** GP regression [7] infers a latent function  $f : \mathbb{R}^N \rightarrow \mathbb{R}$  from data  $\mathbf{D} = (\mathbf{X}, \mathbf{y})$ , where  $\mathbf{X} \in \mathbb{R}^{n \times N}$  and  $\mathbf{y} \in \mathbb{R}^n$ . Given a prior  $p(f) = \mathcal{GP}(f; 0, k)$  with kernel  $k$  for the unknown function  $f$ , the posterior mean and marginal variance at  $m$  new inputs  $\tilde{\mathbf{x}} \in \mathbb{R}^{N \times m}$  are  $\mathbb{E}[\tilde{\mathbf{f}}] = \tilde{\mathbf{k}}^\top (\mathbf{K} + \varepsilon^2 \mathbf{I})^{-1} \mathbf{y}$  and  $\mathbb{V}[\tilde{\mathbf{f}}] = k(\tilde{\mathbf{x}}, \tilde{\mathbf{x}}) - \tilde{\mathbf{k}}^\top (\mathbf{K} + \varepsilon^2 \mathbf{I})^{-1} \tilde{\mathbf{k}}$ , where  $\mathbf{K} = k(\mathbf{X}, \mathbf{X}) \in \mathbb{R}^{n \times n}$  is the Gram matrix of the kernel and  $\tilde{\mathbf{k}} = k(\mathbf{X}, \tilde{\mathbf{x}}) \in \mathbb{R}^{n \times m}$ . The bulk of computation during prediction arises from solving the linear system  $(\mathbf{K} + \varepsilon^2 \mathbf{I})\mathbf{z} = \mathbf{b}$  for some right-hand side  $\mathbf{b} \in \mathbb{R}^n$  repeatedly. When using a probabilistic linear solver for this task, we can quantify the uncertainty arising from finite computation as well as the belief of the solver about the shape of the GP at a set of not yet computed inputs. Figure 3 illustrates this. In fact, we can estimate the marginal variance of the GP without solving the linear system again by multiplying  $\tilde{\mathbf{k}}$  with the estimated inverse of  $\mathbf{K} + \varepsilon^2 \mathbf{I}$ . In large-scale applications, we can trade off computational expense for increased uncertainty arising from the numerical approximation and quantified by the probabilistic linear solver. By assessing the numerical uncertainty arising from not exploring the full space, we can judge the quality of the estimated GP mean and marginal variance.

**Kernel Gram Matrix Inversion** Consider a linear problem  $\mathbf{K}\mathbf{x}_* = \mathbf{b}$ , where  $\mathbf{K}$  is generated by a Mercer kernel. For a  $\nu$ -times continuously differentiable kernel the eigenvalues  $\lambda_n(\mathbf{K})$  decay approximately as  $|\lambda_n| \in \mathcal{O}(n^{-\nu-\frac{1}{2}})$  [33]. We can make use of this generative prior information

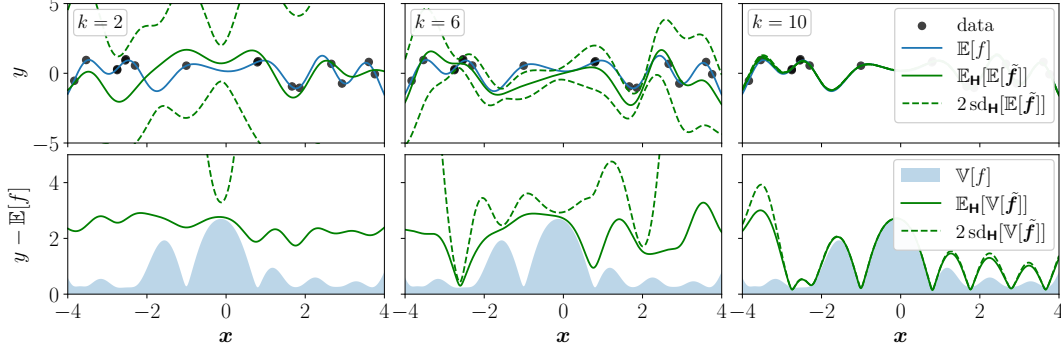


Figure 3: *Numerical uncertainty in GP inference.* Computing posterior mean and covariance of a GP regression using a PLS. *Top:* GP mean for a toy data set ( $n = 16$ ) computed with increasing number of iterations  $k$  of Algorithm 1. The numerical estimate of the GP mean approaches the true mean. Note that the numerical variance is different from the marginal variance of the GP. *Bottom:* GP variance and estimate of GP variance with numerical uncertainty. The GP variance estimate is computed using the estimated inverse from computing  $\mathbb{E}[\hat{\mathbf{f}}]$  without any additional solver iterations.

by specifying a parametrized prior mean  $\mu(n) = \ln(\theta_0 n^{-\theta_1}) = \theta_0 - \theta_1 \ln(n)$  for the ln-Rayleigh quotient model. Typically, such Gram matrices are ill-conditioned and therefore  $\mathbf{K}' = \mathbf{K} + \varepsilon^2 \mathbf{I}$  is used instead, implying  $\lambda(\mathbf{K}')_i \geq \varepsilon^2$ . In order to assess calibration we apply various differentiable kernels to the airline delay dataset from January 2020 [34]. We compute the ln-ratio statistic  $w(\mathbf{x}_*) = \frac{1}{2} \ln(\text{tr}(\text{Cov}[\mathbf{x}])) - \ln(\|\mathbf{x}_* - \mathbb{E}[\mathbf{x}]\|_2)$  for no calibration, calibration via Rayleigh quotient GP regression using  $\mu(n)$  as a prior mean, calibration by setting  $\phi = \varepsilon^2$  and calibration using the average spectrum  $\phi = \bar{\lambda}_{k+1:n}$ . The average  $\bar{w}$  for  $10^5/n$  randomly sampled test problems is shown in Table 2.<sup>2</sup> Without any calibration the solver is generally overconfident. All tested calibration procedures reverse this, resulting in more cautious uncertainty estimates. We observe that Rayleigh quotient regression overcorrects for larger problems. This is due to the fact that its model correctly predicts  $\mathbf{K}$  to be numerically singular from the dominant Rayleigh quotients, however it misses the information that the spectrum of  $\mathbf{K}'$  is bounded from below by  $\varepsilon^2$ . If we know the (average) of the remaining spectrum, significantly better calibration can be achieved, but often this information is not available. Nonetheless, since in this setting the majority of eigenvalues satisfy  $\lambda(\mathbf{K}')_i \approx \varepsilon^2$  by choosing  $\phi = \psi^{-1} = \varepsilon^2$ , we can get to the same degree of calibration. Therefore, we can improve the solver's uncertainty calibration at constant cost  $\mathcal{O}(1)$  per iteration. For more general problems involving Gram matrices without damping we may want to rely on Rayleigh regression instead.

**Galerkin's Method for PDEs** In the spirit of applying machine learning approaches to problems in the physical sciences and vice versa [35], we use Algorithm 1 for the approximate solution of a PDE via Galerkin's method [9]. Consider the Dirichlet problem for the Poisson equation given by

$$\begin{cases} -\Delta u(x, y) = f(x, y) & (x, y) \in \text{int } \Omega \\ u(x, y) = u_{\partial\Omega}(x, y) & (x, y) \in \partial\Omega \end{cases}$$

where  $\Omega$  is a connected open region with sufficiently regular boundary and  $u_{\partial\Omega} : \partial\Omega \rightarrow \mathbb{R}$  defines the boundary conditions. One obtains an approximate solution by projecting the weak formulation of the PDE to a finite dimensional subspace. This results in the *Galerkin equation*  $\mathbf{A}\mathbf{u} = \mathbf{f}$ , i.e. a linear system where  $\mathbf{A}$  is the Gram matrix of the associated bilinear form. Figure 4 shows the induced uncertainty on the solution of the Dirichlet problem for  $f(x, y) = 15$  and  $u_{\partial\Omega}(x, y) = (x^2 - 2y)^2(1 + \sin(2\pi x))$ . The mesh and corresponding Gram matrix were computed using FENICS [36]. We can exploit two properties of Algorithm 1 in this setting. First, if we need to solve multiple related problems  $(\mathbf{A}_j, \mathbf{f}_j)_j$ , by solving a single problem we obtain an estimate of the solution to all other problems. We can successively use the posterior over the inverse as a prior for the next problem. This approach is closely related to subspace recycling in numerical linear algebra [37, 38]. Second, suppose we first compute a solution in a low-dimensional subspace corresponding to a coarse

<sup>2</sup>We decrease the number of samples with the dimension because forming *dense* kernel matrices in memory and computing their eigenvalues becomes computationally prohibitive – *not* because of the cost of our solver.






Figure 4: *Solving the Dirichlet problem with a probabilistic linear solver.* Figures 4a and 4b show the ground truth and mean of the solution computed with Algorithm 1 after  $k = 23$  iterations along with samples from the posterior. The posterior on the coarse mesh can be used to assess uncertainty about the solution on a finer mesh. The signed error computed on the coarse mesh in Figure 4c shows that the approximation is better near the top boundary of  $\Omega$ . Given perfect uncertainty calibration, Figure 4d represents a sample from  $\mathcal{N}(\mathbf{0}, \mathbf{I})$ . The apparent structure in the plot and smaller than expected deviations in the upper part of  $\Omega$  indicate the conservative confidence estimate of the solver.

discretization for computational efficiency. We can then leverage the estimated solution to extrapolate to an (adaptively) refined discretization based on the posterior uncertainty. In machine learning lingo these two approaches can be viewed as forms of *transfer learning*.

## 5 Conclusion

In this work, we condensed a line of previous research on probabilistic linear algebra into a self-contained algorithm for the solution of linear problems in machine learning. We proposed first principles to constrain the space of possible generative models and derived a suitable covariance class. In particular, our proposed framework incorporates prior knowledge on the system matrix or its inverse and performs inference for both in a *consistent* fashion. Within our framework we identified parameter choices that recover the iterates of conjugate gradients in the mean, but add calibrated uncertainty around them in a computationally lightweight manner. To our knowledge our solver, available as part of the **PROBNUM** package, is the first practical implementation of this kind. In the final parts of this paper we showcased applications like kernel matrix inversion, where prior spectral information can be used for uncertainty calibration and outlined example use-cases for propagation of numerical uncertainty through computations. Naturally, there are also limitations remaining. While our theoretical framework can incorporate noisy matrix-vector product evaluations into its inference procedure via a Gaussian likelihood, practically *tractable* inference in the inverse model is more challenging. Our solver also opens up new research directions. In particular, our outlined regression model on the Rayleigh quotient may lead to a probabilistic model of the eigenspectrum. Finally, the matrix-based view of probabilistic linear solvers could inform probabilistic approaches to matrix decompositions, analogous to the way Lanczos methods are used in the classical setting.

## Broader Impact

Our research on probabilistic linear solvers is primarily aimed at members of the machine learning field working on uncertainty estimation which use linear solvers as part of their toolkit. We are convinced that numerical uncertainty induced by finite computational resources is a key missing component to be quantified in machine learning settings. By making numerical uncertainty explicit like our solver does, holistic probabilistic models incorporating all sources of uncertainty become possible. In fact, we hope that this line of work stimulates further research into numerical linear algebra for machine learning, a topic that has been largely considered solved by the community.

This is first and foremost a methods paper aiming to improve the quantification of numerical uncertainty in linear problems. While methodological papers may seem far removed from application and questions of ethical and societal impact, this is not the case. Precisely due to the general nature of the problem setting, the linear solver presented in this work is applicable to a broad range of applications, from regression on flight data, to optimization in robotics, to the solution of PDEs in meteorology.

The flip-side of this potential impact is that arguably, down the line, methodological research suffers from dual use more than any specialized field. While we cannot control the use of a probabilistic linear solver due to its general applicability, we have tried, to the best of our ability, to ensure it performs as intended.

We are hopeful that no specific population group is put at a disadvantage through this research. We are providing an open-source implementation of our method and of all experiments contained in this work. Therefore anybody with access to the internet is able to retrieve and reproduce our findings. In this manner we hope to address the important issues of accessibility and reproducibility.

## Acknowledgments and Disclosure of Funding

The authors gratefully acknowledge financial support by the European Research Council through ERC StG Action 757275 / PANAMA; the DFG Cluster of Excellence “Machine Learning - New Perspectives for Science”, EXC 2064/1, project number 390727645; the German Federal Ministry of Education and Research (BMBF) through the Tübingen AI Center (FKZ: 01IS18039A); and funds from the Ministry of Science, Research and Arts of the State of Baden-Württemberg.

JW is grateful to the International Max Planck Research School for Intelligent Systems (IMPRS-IS) for support.

We thank the reviewers for helpful comments and suggestions. JW would also like to thank Alexandra Gessner and Felix Dangel for a careful reading of an earlier version of this manuscript.

## References

- [1] Youcef Saad. *Numerical methods for large eigenvalue problems*. Manchester University Press, 1992.
- [2] Lloyd N. Trefethen and David Bau. *Numerical Linear Algebra*. Society for Industrial and Applied Mathematics, 1997.
- [3] Gene H. Golub and Charles F. van Loan. *Matrix Computations*. JHU Press, fourth edition, 2013.
- [4] Christopher M. Bishop. *Pattern Recognition and Machine Learning (Information Science and Statistics)*. Springer-Verlag, 2006.
- [5] Thomas Hofmann, Bernhard Schölkopf, and Alexander J. Smola. Kernel methods in machine learning. *The Annals of Statistics*, pages 1171–1220, 2008.
- [6] Rudolph E. Kalman. A new approach to linear filtering and prediction problems. *Journal of Basic Engineering*, 82(1):35–45, 1960.
- [7] Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian Processes for Machine Learning*. The MIT Press, 2006.
- [8] Fan R. K. Chung. *Spectral graph theory*. American Mathematical Society, 1997.
- [9] Clive A. J. Fletcher. *Computational Galerkin methods*. Springer, 1984.
- [10] Jorge Nocedal and Stephen Wright. *Numerical optimization*. Springer Science & Business Media, 2006.
- [11] Philipp Hennig. Probabilistic interpretation of linear solvers. *SIAM Journal on Optimization*, 25(1):234–260, 2015.
- [12] Jon Cockayne, Chris Oates, Tim J. Sullivan, and Mark Girolami. Bayesian probabilistic numerical methods. *SIAM Review*, 61(4):756–789, 2019.
- [13] Simon Bartels, Jon Cockayne, Ilse C. Ipsen, and Philipp Hennig. Probabilistic linear solvers: A unifying view. *Statistics and Computing*, 29(6):1249–1263, 2019.

- [14] Philipp Hennig, Mike A. Osborne, and Mark Girolami. Probabilistic numerics and uncertainty in computations. *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 471(2179), 2015.
- [15] Chris Oates and Tim J. Sullivan. A modern retrospective on probabilistic numerics. *Statistics and Computing*, 10 2019.
- [16] Paul Lévy. *Calcul des probabilités*. J. Gabay, 1925.
- [17] Charles F. Van Loan. The ubiquitous Kronecker product. *Journal of Computational and Applied Mathematics*, 123(1-2):85–100, 2000.
- [18] Philipp Hennig and Martin Kiefel. Quasi-Newton method: A new direction. *Journal of Machine Learning Research*, 14(Mar):843–865, 2013.
- [19] Magnus Rudolph Hestenes and Eduard Stiefel. Methods of conjugate gradients for solving linear systems. *Journal of Research of the National Bureau of Standards*, 49, 1952.
- [20] Martin H. Gutknecht and Zdenek Strakos. Accuracy of two three-term and three two-term recurrences for Krylov space solvers. *SIAM Journal on Matrix Analysis and Applications*, 22(1):213–229, 2000.
- [21] Jon Cockayne, Chris Oates, Ilse C. Ipsen, and Mark Girolami. A Bayesian conjugate gradient method. *Bayesian Analysis*, 14(3):937–1012, 2019.
- [22] Matthias Seeger. Low rank updates for the Cholesky decomposition. Technical report, University of California at Berkeley, 2008.
- [23] David G. Luenberger. *Introduction to Linear and Nonlinear Programming*. Addison-Wesley Publishing Company, 1973.
- [24] Christopher C. Paige. Computational variants of the Lanczos method for the eigenproblem. *IMA Journal of Applied Mathematics*, 10(3):373–381, 1972.
- [25] Horst D. Simon. Analysis of the symmetric Lanczos algorithm with reorthogonalization methods. *Linear algebra and its applications*, 61:101–131, 1984.
- [26] Cornelius Lanczos. *An iteration method for the solution of the eigenvalue problem of linear differential and integral operators*. United States Government Press Office Los Angeles, CA, 1950.
- [27] Petros Drineas and Michael W. Mahoney. RandNLA: randomized numerical linear algebra. *Communications of the ACM*, 59(6):80–90, 2016.
- [28] Alex Gittens and Michael W. Mahoney. Revisiting the Nyström method for improved large-scale machine learning. *Journal of Machine Learning Research*, 17(1):3977–4041, January 2016.
- [29] Simon Bartels and Philipp Hennig. Probabilistic approximate least-squares. In *Proceedings of the 20th International Conference on Artificial Intelligence and Statistics (AISTATS)*, volume 51 of *Proceedings of Machine Learning Research*, pages 676–684, Cadiz, Spain, 09–11 May 2016. PMLR.
- [30] John E. Dennis, Jr and Jorge J. Moré. Quasi-Newton methods, motivation and theory. *SIAM review*, 19(1):46–89, 1977.
- [31] Toni Karvonen and Simo Sarkkå. Approximate state-space Gaussian processes via spectral transformation. In *2016 IEEE 26th International Workshop on Machine Learning for Signal Processing (MLSP)*, pages 1–6, 2016.
- [32] Arno Solin, James Hensman, and Richard E. Turner. Infinite-horizon Gaussian processes. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 3486–3495, 2018.
- [33] Hermann Weyl. Das asymptotische Verteilungsgesetz der Eigenwerte linearer partieller Differentialgleichungen (mit einer Anwendung auf die Theorie der Hohlraumstrahlung). *Mathematische Annalen*, 71(4):441–479, 1912.

- [34] US Department of Transportation. Airline on-time performance data. <https://www.transtats.bts.gov/>, 2020. Accessed: 2020-05-26.
- [35] Giuseppe Carleo, Ignacio Cirac, Kyle Cranmer, Laurent Daudet, Maria Schuld, Naftali Tishby, Leslie Vogt-Maranto, and Lenka Zdeborová. Machine learning and the physical sciences. *Reviews of Modern Physics*, 91(4):045002, 2019.
- [36] Martin Alnæs, Jan Blechta, Johan Hake, August Johansson, Benjamin Kehlet, Anders Logg, Chris Richardson, Johannes Ring, Marie E. Rognes, and Garth N. Wells. The FEniCS project version 1.5. *Archive of Numerical Software*, 3(100), 2015.
- [37] Michael L. Parks, Eric De Sturler, Greg Mackey, Duane D. Johnson, and Spandan Maiti. Recycling Krylov subspaces for sequences of linear systems. *SIAM Journal on Scientific Computing*, 28(5):1651–1674, 2006.
- [38] Filip de Roos and Philipp Hennig. Krylov subspace recycling for fast iterative least-squares in machine learning. *arXiv pre-print*, 2017. URL <http://arxiv.org/abs/1706.00241>.
- [39] Frederick Michael Larkin. Estimation of a non-negative function. *BIT Numerical Mathematics*, 9(1):30–52, 1969.
- [40] Persi Diaconis. Bayesian numerical analysis. *Statistical Decision Theory and Related Topics IV*, 1:163–175, 1988.
- [41] Anthony O’Hagan. Some Bayesian numerical analysis. *Bayesian Statistics*, 4:345–363, 1992.
- [42] Harold V Henderson and Shayle R Searle. The vec-permutation matrix, the vec operator and Kronecker products: A review. *Linear and multilinear algebra*, 9(4):271–288, 1981.
- [43] Farid Alizadeh, Jean-Pierre A. Haeberly, and Michael L. Overton. Primal-dual interior-point methods for semidefinite programming: convergence rates, stability and numerical results. *SIAM Journal on Optimization*, 8(3):746–768, 1998.
- [44] Kathrin Schäcke. On the Kronecker product. Technical report, University of Waterloo, 2013.
- [45] Peder A. Olsen, Steven J. Rennie, and Vaibhava Goel. Efficient automatic differentiation of matrix functions. In *Recent Advances in Algorithmic Differentiation*, pages 71–81. Springer, 2012.
- [46] Arjun K. Gupta and Daya K. Nagar. *Matrix-variate distributions*. Chapman and Hall/CRC, 2000.
- [47] Larry Nazareth. A relationship between the BFGS and conjugate gradient algorithms and its implications for new algorithms. *SIAM Journal on Numerical Analysis*, 16(5):794–800, 1979.
- [48] Keith T. Conrad. The minimal polynomial and some applications. Technical report, University of Connecticut, 2008.
- [49] Mark Girolami, Eky Febrianto, Ge Yin, and Fehmi Cirak. The statistical finite element method (statFEM) for coherent synthesis of observation data and model predictions. *arXiv pre-print*, art. arXiv:1905.06391, May 2020.
- [50] Vladimir Igorevich Bogachev. *Gaussian measures*, volume 62 of *Mathematical Surveys and Monographs*. American Mathematical Society, 1998.
- [51] Pieter Wesseling. *An Introduction to Multigrid Methods*. R.T. Edwards, 2004.

This supplement complements the paper *Probabilistic Linear Solvers for Machine Learning* and is structured as follows. Section S1 explains the approach of probabilistic numerics to model (deterministic) numerical problems probabilistically in more depth. Section S2 introduces different variants of Kronecker products used to define matrix-variate normal distributions in Section S3. Section S4 details the matrix-based inference procedure of probabilistic linear solvers based on matrix-vector product observations. It also contains some more explanation regarding prior construction and stopping criteria. Section S5 and Section S6 outline theoretical results from the paper and properties of the proposed covariance class, in particular detailed proofs. Finally, Section S7 provides some background for the application of probabilistic linear solvers to the solution of discretized partial differential equations. To provide a clear exposition to the reader in some sections we restate results from the literature. References referring to sections, equations or theorem-type environments within this document are tagged with ‘S’, while references to, or results from the main paper are stated as is.

**Preliminaries and Notation** We consider the linear system  $\mathbf{A}\mathbf{x}_* = \mathbf{b}$ , where  $\mathbf{A} \in \mathbb{R}_{\text{sym}}^{n \times n}$  is symmetric positive definite. The random variables  $\mathbf{A}$ ,  $\mathbf{H}$  and  $\mathbf{x}$  model the linear operator  $\mathbf{A}$ , its inverse  $\mathbf{H} = \mathbf{A}^{-1}$  and the solution  $\mathbf{x}_*$ . Algorithm 1 chooses actions  $\mathbf{S} = [\mathbf{s}_1, \dots, \mathbf{s}_k] \in \mathbb{R}^{n \times k}$  given by its policy  $\pi(\mathbf{s} \mid \mathbf{A}, \mathbf{H}, \mathbf{x}, \mathbf{A}, \mathbf{b})$  and computes observations  $\mathbf{Y} = [\mathbf{y}_1, \dots, \mathbf{y}_k] \in \mathbb{R}^{n \times k}$  given by a linear projection  $\mathbf{y}_i = \mathbf{A}\mathbf{s}_i$  in each iteration  $0 < i \leq k$ .

## S1 Probabilistic Modelling of Deterministic Problems

At first glance it might seem counterintuitive to frame a numerical problem in the language of probability theory. After all, when considering the exact problem  $\mathbf{A}\mathbf{x}_* = \mathbf{b}$  all quantities involved  $\mathbf{A}$ ,  $\mathbf{x}_*$ , and  $\mathbf{b}$  are deterministic. However, the distribution of the random variables  $\mathbf{A}$ ,  $\mathbf{H}$  and  $\mathbf{x}$  represents *epistemic uncertainty* arising from finite computational resources. With a finite budget only a limited amount of information can be obtained about  $\mathbf{A}$  (e.g. via matrix-vector products). In particular, for a sufficiently large problem a priori the inverse  $\mathbf{H} = \mathbf{A}^{-1}$  and the solution  $\mathbf{x}_*$ , while deterministic and computable in finite time, are not known. This uncertainty about the inverse is captured by the prior distribution of  $\mathbf{H}$ . In the Bayesian framework the belief about the inverse  $\mathbf{H}$  is then iteratively updated given new observations  $\mathbf{y}_i = \mathbf{A}\mathbf{s}_i$ .

The motivation for also estimating  $\mathbf{A}$  becomes clear if one considers the following. Usually in large-scale applications, the matrix  $\mathbf{A}$  is never actually formed in memory due to computational constraints. Instead only the matrix-vector product  $\mathbf{v} \mapsto \mathbf{A}\mathbf{v}$  is available. Therefore without further computation, the value of any given matrix entry  $\mathbf{A}_{ij}$  is in fact uncertain. Further, generally other properties of the matrix  $\mathbf{A}$  such as its eigenspectrum are also not readily available. The probabilistic framework provides a principled way of incorporating prior knowledge about  $\mathbf{A}$  and makes assumptions about the problem explicit. Relating the prior model  $\mathbf{A}$  and  $\mathbf{H}$  is important here to allow Algorithm 1 to take such prior information into account in its policy. Finally, the strongest argument for a model  $\mathbf{A}$  may yet be the incorporation of noise. Suppose we only have access to  $\mathbf{y}_i = (\mathbf{A} + \mathbf{E}_i)\mathbf{s}_i$  with additive noise  $\mathbf{E}_i$ . This is a common occurrence in application, where the linear system to be solved arises from an approximation itself or if  $\mathbf{A}$  is constructed from data. Concrete examples are batched empirical risk minimization problems or stochastic quadratic optimization. In this setting the probabilistic linear solver must estimate the true  $\mathbf{A}$  via its observations.

The application of probabilistic inference to numerical problems goes back well into the last century [39–41] and has recently seen a resurgence in research interest in the form of *probabilistic numerics*. Overviews discussing motivations and historical perspectives can be found in Hennig et al. [14] and Oates and Sullivan [15]. Hennig [11] gives additional insight into the statistical interpretation of linear systems.

## S2 The Kronecker Product and its Variants

We will now introduce different types of Kronecker products needed for constructing covariances for matrix-variate distributions. In order to transfer results from probabilistic modelling of vector-variate random variables to the matrix-variate case, we need two types of vectorization operations, i.e. bijections between spaces of matrices and vector spaces.

Let  $\text{vec} : \mathbb{R}^{m \times n} \rightarrow \mathbb{R}^{mn}$ , denote the *column-wise stacking operator* [42], defined as

$$\text{vec}(\mathbf{X}) = (X_{11}, X_{21}, \dots, X_{m1}, X_{12}, \dots, X_{mn})^\top \in \mathbb{R}^{mn}.$$

Further, define  $\text{svec} : \mathbb{R}_{\text{sym}}^{n \times n} \rightarrow \mathbb{R}^{\frac{1}{2}n(n+1)}$ , the *column-wise symmetric stacking operator* [43] given by

$$\text{svec}(\mathbf{X}) = (X_{11}, \sqrt{2}X_{21}, \dots, \sqrt{2}X_{n1}, X_{22}, \sqrt{2}X_{32}, \dots, \sqrt{2}X_{n2}, \dots, X_{nn})^\top \in \mathbb{R}^{\frac{1}{2}n(n+1)}.$$

To translate between the two representations following Schacke [44] we also define the matrix  $\mathbf{Q} \in \mathbb{R}^{\frac{1}{2}n(n+1) \times n^2}$  such that for all symmetric matrices  $\mathbf{X} \in \mathbb{R}_{\text{sym}}^{n \times n}$ , we have  $\mathbf{Q} \text{vec}(\mathbf{X}) = \text{svec}(\mathbf{X})$  and  $\text{vec}(\mathbf{X}) = \mathbf{Q}^\top \text{svec}(\mathbf{X})$ . Note, that  $\mathbf{Q}$  has orthonormal rows, i.e.  $\mathbf{Q}\mathbf{Q}^\top = \mathbf{I}$ . For convenience we also name the inverse operations  $\text{mat} := \text{vec}^{-1}$  and  $\text{smat} := \text{svec}^{-1}$ .

## S2.1 Kronecker Product

We make extensive use of Kronecker-type structures for covariance matrices of matrix-variate distributions in this paper. The *Kronecker product*  $\mathbf{A} \otimes \mathbf{B}$  [17] of two matrices  $\mathbf{A} \in \mathbb{R}^{m_1 \times n_1}$  and  $\mathbf{B} \in \mathbb{R}^{m_2 \times n_2}$  is given by

$$\mathbf{A} \otimes \mathbf{B} = \begin{pmatrix} \mathbf{A}_{11}\mathbf{B} & \dots & \mathbf{A}_{1n_1}\mathbf{B} \\ \vdots & \ddots & \vdots \\ \mathbf{A}_{m_11}\mathbf{B} & \dots & \mathbf{A}_{m_1n_1}\mathbf{B} \end{pmatrix} \in \mathbb{R}^{(m_1m_2) \times (n_1n_2)}$$

The Kronecker product satisfies the characteristic property

$$(\mathbf{A} \otimes \mathbf{B}) \text{vec}(\mathbf{X}) = \text{vec}(\mathbf{B}\mathbf{X}\mathbf{A}^\top), \quad (\text{S4})$$

for  $\mathbf{X} \in \mathbb{R}^{n_2 \times n_1}$ . Characteristic properties of Kronecker-type products are useful to turn matrix equations into vector equations. We state a set of properties of the Kronecker product next without proof. More detail on Kronecker products can be found in Van Loan [17].

**Proposition S2** (Properties of the Kronecker Product [17])

*The Kronecker product satisfies the following identities:*

$$\exists \mathbf{A}, \mathbf{B} : \mathbf{A} \otimes \mathbf{B} \neq \mathbf{B} \otimes \mathbf{A} \quad (\text{S5})$$

$$(\mathbf{A} \otimes \mathbf{B})^\top = \mathbf{A}^\top \otimes \mathbf{B}^\top \quad (\text{S6})$$

$$(\mathbf{A} \otimes \mathbf{B})^{-1} = \mathbf{A}^{-1} \otimes \mathbf{B}^{-1} \quad (\text{S7})$$

$$(\mathbf{A} + \mathbf{B}) \otimes \mathbf{C} = \mathbf{A} \otimes \mathbf{C} + \mathbf{B} \otimes \mathbf{C} \quad (\text{S8})$$

$$(\mathbf{A} \otimes \mathbf{B})(\mathbf{C} \otimes \mathbf{D}) = (\mathbf{A}\mathbf{C}) \otimes (\mathbf{B}\mathbf{D}) \quad (\text{S9})$$

$$\text{tr}(\mathbf{A} \otimes \mathbf{B}) = \text{tr}(\mathbf{A}) \text{tr}(\mathbf{B}) \quad (\text{S10})$$

$$\mathbf{A} \in \mathbb{R}_{\text{sym}}^{m \times m}, \mathbf{B} \in \mathbb{R}_{\text{sym}}^{n \times n} \implies \mathbf{A} \otimes \mathbf{B} \in \mathbb{R}_{\text{sym}}^{mn \times mn} \quad (\text{S11})$$

$$\mathbf{A} \otimes \mathbf{B} = (\mathbf{L}_\mathbf{A} \mathbf{L}_\mathbf{A}^\top) \otimes (\mathbf{L}_\mathbf{B} \mathbf{L}_\mathbf{B}^\top) = (\mathbf{L}_\mathbf{A} \otimes \mathbf{L}_\mathbf{B})(\mathbf{L}_\mathbf{A}^\top \otimes \mathbf{L}_\mathbf{B}^\top) \quad (\text{S12})$$

$$\mathbf{A} \otimes \mathbf{B} = (\mathbf{U}_\mathbf{A} \mathbf{\Lambda}_\mathbf{A} \mathbf{U}_\mathbf{A}^\top) \otimes (\mathbf{U}_\mathbf{B} \mathbf{\Lambda}_\mathbf{B} \mathbf{U}_\mathbf{B}^\top) = (\mathbf{U}_\mathbf{A} \otimes \mathbf{U}_\mathbf{B})(\mathbf{\Lambda}_\mathbf{A} \otimes \mathbf{\Lambda}_\mathbf{B})(\mathbf{U}_\mathbf{A}^\top \otimes \mathbf{U}_\mathbf{B}^\top) \quad (\text{S13})$$

## S2.2 Box Product

The *box product*  $\mathbf{A} \boxtimes \mathbf{B} \in \mathbb{R}^{(m_1m_2) \times (n_1n_2)}$  can be defined via its characteristic property

$$(\mathbf{A} \boxtimes \mathbf{B}) \text{vec}(\mathbf{Y}) = \text{vec}(\mathbf{B}\mathbf{Y}^\top \mathbf{A}^\top) \quad (\text{S14})$$

for  $\mathbf{Y} \in \mathbb{R}^{n_1 \times n_2}$ . See also Olsen et al. [45] for details.

**Proposition S3** (Properties of the Box Product [45])

*The box product satisfies the following identities:*

$$\exists \mathbf{A}, \mathbf{B} : \mathbf{A} \boxtimes \mathbf{B} \neq \mathbf{B} \boxtimes \mathbf{A} \quad (\text{S15})$$

$$(\mathbf{A} \boxtimes \mathbf{B})^\top = \mathbf{B}^\top \boxtimes \mathbf{A}^\top \quad (\text{S16})$$

$$(\mathbf{A} \boxtimes \mathbf{B})^{-1} = \mathbf{B}^{-1} \boxtimes \mathbf{A}^{-1} \quad (\text{S17})$$

$$(\mathbf{A} + \mathbf{B}) \boxtimes \mathbf{C} = \mathbf{A} \boxtimes \mathbf{C} + \mathbf{B} \boxtimes \mathbf{C} \quad (\text{S18})$$

$$(\mathbf{A} \boxtimes \mathbf{B})(\mathbf{C} \boxtimes \mathbf{D}) = (\mathbf{A}\mathbf{C}) \boxtimes (\mathbf{B}\mathbf{D}) \quad (\text{S19})$$

$$(\mathbf{A} \boxtimes \mathbf{B})(\mathbf{C} \otimes \mathbf{D}) = (\mathbf{A}\mathbf{D}) \boxtimes (\mathbf{B}\mathbf{C}) \quad (\text{S20})$$

$$(\mathbf{A} \otimes \mathbf{B})(\mathbf{C} \boxtimes \mathbf{D}) = (\mathbf{A}\mathbf{C}) \boxtimes (\mathbf{B}\mathbf{D}) \quad (\text{S21})$$

$$\text{tr}(\mathbf{A} \boxtimes \mathbf{B}) = \text{tr}(\mathbf{A}\mathbf{B}) \quad (\text{S22})$$

### S2.3 Symmetric Kronecker Product

The *symmetric Kronecker product*  $A \otimes B$  of two square matrices  $A, B \in \mathbb{R}^{n \times n}$  is defined via its characteristic property for  $X \in \mathbb{R}_{\text{sym}}^{n \times n}$  as

$$(A \otimes B) \text{svec}(X) = \frac{1}{2} \text{svec}(BXA^\top + AXB^\top) \quad (\text{S23})$$

or equivalently

$$A \otimes B = \frac{1}{2} Q(A \otimes B + B \otimes A) Q^\top.$$

**Proposition S4** (Properties of the Symmetric Kronecker Product [43, 44])

The symmetric Kronecker product satisfies the following identities:

$$A \otimes B = B \otimes A \quad (\text{S24})$$

$$(A \otimes B)^\top = A^\top \otimes B^\top \quad (\text{S25})$$

$$(A \otimes A)^{-1} = A^{-1} \otimes A^{-1} \quad (\text{S26})$$

$$(A + B) \otimes C = A \otimes C + B \otimes C \quad (\text{S27})$$

$$(A \otimes B)(C \otimes D) = \frac{1}{2}(AC \otimes BD + AD \otimes BC) \quad (\text{S28})$$

$$A \in \mathbb{R}_{\text{sym}}^{n \times n}, B \in \mathbb{R}_{\text{sym}}^{n \times n} \implies A \otimes B \in \mathbb{R}_{\text{sym}}^{\frac{1}{2}n(n+1) \times \frac{1}{2}n(n+1)} \quad (\text{S29})$$

$$A \otimes A = (L_A L_A^\top) \otimes (L_A L_A^\top) = (L_A \otimes L_A)(L_A^\top \otimes L_A^\top) \quad (\text{S30})$$

$$A \otimes A = (U_A \Lambda_A U_A^\top) \otimes (U_A \Lambda_A U_A^\top) = (U_A \otimes U_A)(\Lambda_A \otimes \Lambda_A)(U_A^\top \otimes U_A^\top) \quad (\text{S31})$$

Note, that the symmetric Kronecker product represented as a  $\frac{1}{2}n(n+1) \times \frac{1}{2}n(n+1)$  matrix is in general not symmetric.

Further properties can be found in Alizadeh et al. [43] and Schacke [44]. We prove the following technical results for mixed expressions of Kronecker-type products, which we will make use of later.

**Corollary S1** (Mixed Kronecker Product Identities)

Let  $A \in \mathbb{R}_{\text{sym}}^{n \times n}$ ,  $B, C \in \mathbb{R}^{n \times k}$  and  $X \in \mathbb{R}^{k \times k}$  such that  $(CXB^\top)^\top = CXB^\top$ , then it holds that

$$Q^\top(A \otimes A)Q(B \otimes C) \text{vec}(X) = \frac{1}{2}(AB \otimes AC + AC \boxtimes AB) \text{vec}(X) \quad (\text{S32})$$

$$(B^\top \otimes C^\top)Q^\top(A \otimes A)Q = \frac{1}{2}(B^\top A \otimes C^\top A + B^\top A \boxtimes C^\top A). \quad (\text{S33})$$

$$(B^\top \otimes C^\top)Q^\top(A \otimes A)Q(B \otimes C) \text{vec}(X) = \frac{1}{2}(B^\top AB \otimes C^\top AC + B^\top AC \boxtimes C^\top AB) \text{vec}(X). \quad (\text{S34})$$

Now, assume  $A$  to be invertible,  $\text{rank}(C) = k$  and  $Y \in \mathbb{R}^{k \times n}$  such that  $(YC)^\top = YC$ , then for

$$G = (I_n \otimes C^\top)Q^\top(A \otimes A)Q(I_n \otimes C)$$

$$G_{\text{right}}^{-1} = (2A^{-1} - C(C^\top AC)^{-1}C^\top) \otimes (C^\top AC)^{-1}$$

we have  $GG_{\text{right}}^{-1} \text{vec}(Y) = \text{vec}(Y)$ , i.e.  $G_{\text{right}}^{-1}$  is the right inverse of  $G$ . Finally, for  $D, E \in \mathbb{R}^{n \times n}$  and  $Z \in \mathbb{R}_{\text{sym}}^{n \times n}$  such that  $(EAZAD^\top)^\top = EAZAD^\top$ , we have

$$(A^\top \otimes A^\top)Q(D \otimes E)Q^\top(A \otimes A) \text{svec}(Z) = (A^\top DA) \otimes (A^\top EA) \text{svec}(Z). \quad (\text{S35})$$

*Proof.* Let  $X \in \mathbb{R}^{k \times k}$  such that  $(CXB^\top)^\top = CXB^\top$ , then

$$\begin{aligned} Q^\top(A \otimes A)Q(B \otimes C) \text{vec}(X) &= Q^\top(A \otimes A)Q \text{vec}(CXB^\top) \\ &= Q^\top(A \otimes A) \text{svec}(CXB^\top) \\ &= Q^\top \text{svec}(ACXB^\top A) \\ &= \frac{1}{2} \text{vec}(ACXB^\top A + ABX^\top C^\top A) \\ &= \frac{1}{2}(AB \otimes AC + AC \boxtimes AB), \end{aligned}$$

further it holds for  $\mathbf{W} \in \mathbb{R}_{\text{sym}}^{n \times n}$

$$\begin{aligned}
(\mathbf{B}^\top \otimes \mathbf{C}^\top) \mathbf{Q}^\top (\mathbf{A} \otimes \mathbf{A}) \mathbf{Q} \text{vec}(\mathbf{W}) &= (\mathbf{B}^\top \otimes \mathbf{C}^\top) \mathbf{Q}^\top \text{svec}(\mathbf{A} \mathbf{W} \mathbf{A}) \\
&= \text{vec}(\mathbf{C}^\top \mathbf{A} \mathbf{W} \mathbf{A} \mathbf{B}) \\
&= \frac{1}{2} (\mathbf{C}^\top \mathbf{A} \mathbf{W} \mathbf{A} \mathbf{B} + \mathbf{C}^\top \mathbf{A}^\top \mathbf{W}^\top \mathbf{A}^\top \mathbf{B}) \\
&= \frac{1}{2} (\mathbf{B}^\top \mathbf{A} \otimes \mathbf{C}^\top \mathbf{A} + \mathbf{B}^\top \mathbf{A} \boxtimes \mathbf{C}^\top \mathbf{A}),
\end{aligned}$$

and using the properties of the Kronecker and the Box product we obtain

$$\begin{aligned}
(\mathbf{B}^\top \otimes \mathbf{C}^\top) \mathbf{Q}^\top (\mathbf{A} \otimes \mathbf{A}) \mathbf{Q} (\mathbf{B} \otimes \mathbf{C}) \text{vec}(\mathbf{X}) &= (\mathbf{B}^\top \otimes \mathbf{C}^\top) \frac{1}{2} (\mathbf{B}^\top \mathbf{A} \otimes \mathbf{C}^\top \mathbf{A} + \mathbf{B}^\top \mathbf{A} \boxtimes \mathbf{C}^\top \mathbf{A}) \text{vec}(\mathbf{X}) \\
&= \frac{1}{2} (\mathbf{B}^\top \mathbf{A} \otimes \mathbf{C}^\top \mathbf{A} + \mathbf{B}^\top \mathbf{A} \boxtimes \mathbf{C}^\top \mathbf{A}) \text{vec}(\mathbf{X}).
\end{aligned}$$

Now let  $\mathbf{A}$  be invertible, let  $\mathbf{C}$  have full rank and choose  $\mathbf{Y} \in \mathbb{R}^{k \times n}$  arbitrarily such that  $(\mathbf{Y} \mathbf{C})^\top = \mathbf{Y} \mathbf{C}$ . Then using Proposition S2 and Proposition S3 we obtain

$$\begin{aligned}
(\mathbf{I}_n \otimes \mathbf{C}^\top) \mathbf{Q}^\top (\mathbf{A} \otimes \mathbf{A}) \mathbf{Q} (\mathbf{I}_n \otimes \mathbf{C}) (2\mathbf{A}^{-1} - \mathbf{C}(\mathbf{C}^\top \mathbf{A} \mathbf{C})^{-1} \mathbf{C}^\top) \otimes (\mathbf{C}^\top \mathbf{A} \mathbf{C})^{-1} \text{vec}(\mathbf{Y}) \\
&= \frac{1}{2} (\mathbf{A} \otimes \mathbf{C}^\top \mathbf{A} \mathbf{C} + \mathbf{A} \mathbf{C} \boxtimes \mathbf{C}^\top \mathbf{A}) (2\mathbf{A}^{-1} - \mathbf{C}(\mathbf{C}^\top \mathbf{A} \mathbf{C})^{-1} \mathbf{C}^\top) \otimes (\mathbf{C}^\top \mathbf{A} \mathbf{C})^{-1} \text{vec}(\mathbf{Y}) \\
&= (\mathbf{I}_n \otimes \mathbf{I}_k - \frac{1}{2} \mathbf{A} \mathbf{C} (\mathbf{C}^\top \mathbf{A} \mathbf{C})^{-1} \mathbf{C}^\top \otimes \mathbf{I}_k + \mathbf{A} \mathbf{C} (\mathbf{C}^\top \mathbf{A} \mathbf{C})^{-1} \boxtimes \mathbf{C}^\top - \frac{1}{2} \mathbf{A} \mathbf{C} (\mathbf{C}^\top \mathbf{A} \mathbf{C})^{-1} \boxtimes \mathbf{C}^\top) \text{vec}(\mathbf{Y}) \\
&= (\mathbf{I}_n \otimes \mathbf{I}_k - \frac{1}{2} \mathbf{A} \mathbf{C} (\mathbf{C}^\top \mathbf{A} \mathbf{C})^{-1} \mathbf{C}^\top \otimes \mathbf{I}_k + \frac{1}{2} \mathbf{A} \mathbf{C} (\mathbf{C}^\top \mathbf{A} \mathbf{C})^{-1} \boxtimes \mathbf{C}^\top) \text{vec}(\mathbf{Y}) \\
&= \text{vec}(\mathbf{Y}) - \frac{1}{2} (\mathbf{Y} \mathbf{C} (\mathbf{C}^\top \mathbf{A} \mathbf{C})^{-1} \mathbf{C}^\top \mathbf{A} - \mathbf{C}^\top \mathbf{Y}^\top (\mathbf{C}^\top \mathbf{A} \mathbf{C})^{-1} \mathbf{C}^\top \mathbf{A}) \\
&= \text{vec}(\mathbf{Y})
\end{aligned}$$

Lastly, by assumption it holds that

$$\begin{aligned}
(\mathbf{A}^\top \otimes \mathbf{A}^\top) \mathbf{Q} (\mathbf{D} \otimes \mathbf{E}) \mathbf{Q}^\top (\mathbf{A} \otimes \mathbf{A}) \text{svec}(\mathbf{Z}) &= (\mathbf{A} \otimes \mathbf{A}) \mathbf{Q} \text{vec}(\mathbf{E} \mathbf{A} \mathbf{Z} \mathbf{A} \mathbf{D}^\top) \\
&= \text{svec}(\mathbf{A} \mathbf{E} \mathbf{A} \mathbf{Z} \mathbf{A} \mathbf{D}^\top \mathbf{A}) \\
&= \frac{1}{2} (\mathbf{A} \mathbf{E} \mathbf{A} \mathbf{Z} \mathbf{A} \mathbf{D}^\top \mathbf{A} + \mathbf{A} \mathbf{D} \mathbf{A} \mathbf{Z} \mathbf{A} \mathbf{E}^\top \mathbf{A}) \\
&= (\mathbf{A} \mathbf{D} \mathbf{A} \otimes \mathbf{A} \mathbf{E} \mathbf{A}) \text{svec}(\mathbf{Z}).
\end{aligned}$$

This concludes the proof.  $\square$

### S3 The Matrix-variate Normal Distribution

In order for our probabilistic linear solvers to infer the true latent  $\mathbf{A}$  or its inverse  $\mathbf{H} = \mathbf{A}^{-1}$ , we need a distribution expressing the belief of the solver over those latent quantities at any given point. A Gaussian distribution over matrices will play this role, motivated by the linear nature of the observations. This section closely follows Gupta and Nagar [46].

**Definition S2** (Matrix-variate Normal Distribution [46])

Let  $\mathbf{X}_0 \in \mathbb{R}^{m \times n}$  and let  $\mathbf{V} \in \mathbb{R}_{\text{sym}}^m$  and  $\mathbf{W} \in \mathbb{R}_{\text{sym}}^{n \times n}$  be positive-definite. We say a random matrix  $\mathbf{X}$  has a *matrix-variate normal distribution* with mean  $\mathbf{X}_0$  and covariance  $\mathbf{V} \otimes \mathbf{W}$ , iff

$$\text{vec}(\mathbf{X}^\top) \sim \mathcal{N}_{mn}(\text{vec}(\mathbf{X}_0^\top), \mathbf{V} \otimes \mathbf{W}).$$

We write as a shorthand  $\mathbf{X} \sim \mathcal{N}(\mathbf{X}_0, \mathbf{V} \otimes \mathbf{W})$ .

Note, that the matrices  $\mathbf{V}$  and  $\mathbf{W}$  represent the covariance between rows and columns of  $\mathbf{X}$ , respectively. Since we model symmetric matrices in this work, we also introduce a Gaussian distribution over  $\mathbb{R}_{\text{sym}}^{n \times n}$ .

**Definition S3** (Symmetric Matrix-variate Normal Distribution [46])

Let  $\mathbf{X}_0, \mathbf{W} \in \mathbb{R}_{\text{sym}}^{n \times n}$  such that  $\mathbf{W}$  is positive-definite, then the random matrix  $\mathbf{X}$  has a *symmetric matrix-variate normal distribution*, iff

$$\text{svec}(\mathbf{X}) \sim \mathcal{N}_{\frac{1}{2}n(n+1)}(\text{svec}(\mathbf{X}_0), \mathbf{W} \otimes \mathbf{W}).$$

We write  $\mathbf{X} \sim \mathcal{N}(\mathbf{X}_0, \mathbf{W} \otimes \mathbf{W})$ .



It follows immediately from the definition that realizations of a symmetric matrix-variate normal distribution are symmetric matrices. This distribution also emerges naturally by conditioning a matrix-variate normal distribution on the linear constraint  $\mathbf{X} = \mathbf{X}^\top$ .

## S4 Probabilistic Linear Solvers

Probabilistic linear solvers (PLS) [11, 21, 13] infer posterior beliefs over the matrix  $\mathbf{A}$ , its inverse  $\mathbf{H}$  or the solution  $\mathbf{x}_* = \mathbf{H}\mathbf{b}$  of a linear system via linear observations  $\mathbf{Y} = \mathbf{A}\mathbf{S}$ . We consider matrix-based inference [13] in this work. Assuming a prior  $p(\mathbf{A})$  or  $p(\mathbf{H})$ , actions  $\mathbf{S}$  and linear observations  $\mathbf{Y}$  such methods return posterior distributions  $p(\mathbf{A} \mid \mathbf{S}, \mathbf{Y})$  or  $p(\mathbf{H} \mid \mathbf{S}, \mathbf{Y})$ .

### S4.1 Matrix-based Inference

The generic matrix-based inference procedure of probabilistic linear solvers is a consequence of the matrix-variate version of the following standard result for Gaussian inference under linear observations.

**Theorem S4** (Linear Gaussian Inference [4])

Let  $\mathbf{v} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ , where  $\boldsymbol{\mu} \in \mathbb{R}^n$  and  $\boldsymbol{\Sigma} \in \mathbb{R}_{\text{sym}}^{n \times n}$  positive-definite, and assume we are given observations of the form

$$\mathbf{B}\mathbf{v} + \mathbf{b} = \mathbf{y} \in \mathbb{R}^m,$$

where  $\mathbf{B} \in \mathbb{R}^{m \times n}$  and  $\mathbf{b} \in \mathbb{R}^m$ . Assuming a Gaussian likelihood

$$p(\mathbf{y} \mid \mathbf{B}, \mathbf{v}, \mathbf{b}) = \mathcal{N}(\mathbf{y}; \mathbf{B}\mathbf{v} + \mathbf{b}, \boldsymbol{\Lambda}),$$

for  $\boldsymbol{\Lambda} \in \mathbb{R}_{\text{sym}}^m$  positive definite, results in the posterior distribution

$$p(\mathbf{v} \mid \mathbf{y}, \mathbf{B}, \mathbf{b}) = \mathcal{N}(\mathbf{v}; \boldsymbol{\mu} + \boldsymbol{\Sigma}\mathbf{B}^\top(\mathbf{B}\boldsymbol{\Sigma}\mathbf{B}^\top + \boldsymbol{\Lambda})^{-1}(\mathbf{y} - \mathbf{B}\boldsymbol{\mu} - \mathbf{b}), \\ \boldsymbol{\Sigma} - \boldsymbol{\Sigma}\mathbf{B}^\top(\mathbf{B}\boldsymbol{\Sigma}\mathbf{B}^\top + \boldsymbol{\Lambda})^{-1}\mathbf{B}\boldsymbol{\Sigma}).$$

Further, the marginal distribution of  $\mathbf{y}$  is given by

$$p(\mathbf{y}) = \mathcal{N}(\mathbf{y}; \mathbf{B}\boldsymbol{\mu} + \mathbf{b}, \mathbf{B}\boldsymbol{\Sigma}\mathbf{B}^\top + \boldsymbol{\Lambda}).$$

#### S4.1.1 Asymmetric Model

**Corollary S2** (Asymmetric matrix-based Gaussian Inference [18, 11, 13])

Assume a prior  $p(\mathbf{A}) = \mathcal{N}(\mathbf{A}; \mathbf{A}_0, \mathbf{V}_0 \otimes \mathbf{W}_0)$  and exact observations of the form  $\mathbf{Y} = \mathbf{A}\mathbf{S}$ , corresponding to a Dirac likelihood  $p(\mathbf{Y} \mid \mathbf{A}, \mathbf{S}) = \delta(\mathbf{Y} - \mathbf{A}\mathbf{S})$ , then the posterior  $p(\mathbf{A} \mid \mathbf{S}, \mathbf{Y}) = \mathcal{N}(\mathbf{A}; \mathbf{A}_k, \boldsymbol{\Sigma}_k)$  is given by

$$\mathbf{A}_k = \mathbf{A}_0 + \boldsymbol{\Delta}_0\mathbf{U}^\top \\ \boldsymbol{\Sigma}_k = \mathbf{V}_0 \otimes \mathbf{W}_0(\mathbf{I}_n - \mathbf{S}\mathbf{U}^\top)$$

where  $\boldsymbol{\Delta}_0 = \mathbf{Y} - \mathbf{A}_0\mathbf{S}$  and  $\mathbf{U} = \mathbf{W}_0\mathbf{S}(\mathbf{S}^\top\mathbf{W}_0\mathbf{S})^{-1}$ .

*Proof.* In vectorized form the likelihood is given by

$$p(\text{vec}(\mathbf{Y}^\top) \mid \text{vec}(\mathbf{A}^\top), \text{vec}(\mathbf{S}^\top)) = \delta(\text{vec}(\mathbf{Y}^\top) - \text{vec}(\mathbf{S}^\top\mathbf{A}^\top)) = \delta(\text{vec}(\mathbf{Y}^\top) - (\mathbf{I} \otimes \mathbf{S}^\top) \text{vec}(\mathbf{A}^\top))$$

Using the Definition S2 of the matrix-variate normal distribution, applying Theorem S4 and using property (S9) of the Kronecker product in Proposition S2 leads to

$$\text{vec}(\mathbf{A}_k^\top) = \text{vec}(\mathbf{A}_0^\top) + (\mathbf{V}_0 \otimes \mathbf{W}_0)(\mathbf{I} \otimes \mathbf{S})((\mathbf{I} \otimes \mathbf{S}^\top)(\mathbf{V}_0 \otimes \mathbf{W}_0)(\mathbf{I} \otimes \mathbf{S}))^{-1}(\text{vec}(\mathbf{Y}^\top) - (\mathbf{I} \otimes \mathbf{S}^\top) \text{vec}(\mathbf{A}_0^\top)) \\ = \text{vec}(\mathbf{A}_0^\top) + (\mathbf{V}_0 \otimes \mathbf{W}_0\mathbf{S})(\mathbf{V}_0 \otimes \mathbf{S}^\top\mathbf{W}_0\mathbf{S})^{-1} \text{vec}(\boldsymbol{\Delta}_0^\top) \\ = \text{vec}(\mathbf{A}_0^\top) + (\mathbf{I}_n \otimes \mathbf{W}_0\mathbf{S}(\mathbf{S}^\top\mathbf{W}_0\mathbf{S})^{-1}) \text{vec}(\boldsymbol{\Delta}_0^\top) \\ = \text{vec}(\mathbf{A}_0^\top + \mathbf{U}\boldsymbol{\Delta}_0^\top)$$

and further analogously, additionally using bilinearity of the Kronecker product, we obtain

$$\boldsymbol{\Sigma}_k = \mathbf{V}_0 \otimes \mathbf{W}_0 - (\mathbf{V}_0 \otimes \mathbf{W}_0)(\mathbf{I} \otimes \mathbf{S})((\mathbf{I} \otimes \mathbf{S}^\top)(\mathbf{V}_0 \otimes \mathbf{W}_0)(\mathbf{I} \otimes \mathbf{S}))^{-1}(\mathbf{I} \otimes \mathbf{S}^\top)(\mathbf{V}_0 \otimes \mathbf{W}_0) \\ = \mathbf{V}_0 \otimes \mathbf{W}_0 - (\mathbf{V}_0 \otimes \mathbf{W}_0\mathbf{S})(\mathbf{V}_0 \otimes \mathbf{S}^\top\mathbf{W}_0\mathbf{S})^{-1}(\mathbf{V}_0 \otimes \mathbf{S}^\top\mathbf{W}_0) \\ = \mathbf{V}_0 \otimes \mathbf{W}_0 - \mathbf{V}_0 \otimes (\mathbf{W}_0\mathbf{S}(\mathbf{S}^\top\mathbf{W}_0\mathbf{S})^{-1}\mathbf{S}^\top\mathbf{W}_0) \\ = \mathbf{V}_0 \otimes \mathbf{W}_0(\mathbf{I} - \mathbf{S}\mathbf{U}^\top).$$

This concludes the proof.  $\square$

### S4.1.2 Symmetric Model

**Corollary S3** (Symmetric Matrix-based Gaussian Inference [18, 11, 13])

Assume a symmetric prior  $p(\mathbf{A}) = \mathcal{N}(\mathbf{A}; \mathbf{A}_0, \mathbf{W}_0 \otimes \mathbf{W}_0)$  and exact observations of the form  $\mathbf{Y} = \mathbf{A}\mathbf{S}$ , corresponding to a Dirac likelihood  $p(\mathbf{Y} | \mathbf{A}, \mathbf{S}) = \delta(\mathbf{Y} - \mathbf{A}\mathbf{S})$ , then the posterior  $p(\mathbf{A} | \mathbf{S}, \mathbf{Y}) = \mathcal{N}(\mathbf{A}; \mathbf{A}_k, \mathbf{\Sigma}_k)$  is given by

$$\begin{aligned}\mathbf{A}_k &= \mathbf{A}_0 + \mathbf{\Delta}_0 \mathbf{U}^\top + \mathbf{U} \mathbf{\Delta}_0^\top - \mathbf{U} \mathbf{S}^\top \mathbf{\Delta}_0 \mathbf{U}^\top = \mathbf{A}_0 + \mathbf{U} \mathbf{V}^\top + \mathbf{V} \mathbf{U}^\top \\ \mathbf{\Sigma}_k &= \mathbf{W}_0 (\mathbf{I}_n - \mathbf{S} \mathbf{U}^\top) \otimes \mathbf{W}_0 (\mathbf{I}_n - \mathbf{S} \mathbf{U}^\top)\end{aligned}$$

where  $\mathbf{\Delta}_0 = \mathbf{Y} - \mathbf{A}_0 \mathbf{S}$ ,  $\mathbf{U} = \mathbf{W}_0 \mathbf{S} (\mathbf{S}^\top \mathbf{W}_0 \mathbf{S})^{-1}$  and  $\mathbf{V} = (\mathbf{I}_n - \frac{1}{2} \mathbf{U} \mathbf{S}^\top) \mathbf{\Delta}_0$ .

*Proof.* A proof can be found in the appendix of Hennig [11]. We rederive it here in our notation. By assumption the likelihood takes the vectorized form

$$p(\text{vec}(\mathbf{Y}^\top) | \text{svec}(\mathbf{A}), \text{vec}(\mathbf{S}^\top)) = \delta(\text{vec}(\mathbf{Y}^\top) - \text{vec}(\mathbf{S}^\top \mathbf{A}^\top)) = \delta(\text{vec}(\mathbf{Y}^\top) - (\mathbf{I} \otimes \mathbf{S}^\top) \mathbf{Q}^\top \text{svec}(\mathbf{A}))$$

Applying Theorem S4 gives

$$\begin{aligned}\text{svec}(\mathbf{A}_k) &= \text{svec}(\mathbf{A}_0) + (\mathbf{W}_0 \otimes \mathbf{W}_0) \mathbf{Q} (\mathbf{I}_n \otimes \mathbf{S}) \mathbf{G}^{-1} (\text{vec}(\mathbf{Y}^\top) - (\mathbf{I} \otimes \mathbf{S}^\top) \mathbf{Q}^\top \text{svec}(\mathbf{A}_0)) \\ &= \text{svec}(\mathbf{A}_0) + (\mathbf{W}_0 \otimes \mathbf{W}_0) \mathbf{Q} (\mathbf{I}_n \otimes \mathbf{S}) \mathbf{G}^{-1} \text{vec}(\mathbf{\Delta}_0^\top) \\ \mathbf{\Sigma}_k &= \mathbf{W}_0 \otimes \mathbf{W}_0 - (\mathbf{W}_0 \otimes \mathbf{W}_0) \mathbf{Q} (\mathbf{I}_n \otimes \mathbf{S}) \mathbf{G}^{-1} (\mathbf{I}_n \otimes \mathbf{S}^\top) \mathbf{Q}^\top (\mathbf{W}_0 \otimes \mathbf{W}_0),\end{aligned}$$

where  $\mathbf{\Delta}_0 = \mathbf{Y} - \mathbf{A}_0 \mathbf{S}$  and the Gram matrix is given by

$$\mathbf{G} = (\mathbf{I}_n \otimes \mathbf{S}^\top) \mathbf{Q}^\top (\mathbf{W}_0 \otimes \mathbf{W}_0) \mathbf{Q} (\mathbf{I}_n \otimes \mathbf{S}) \in \mathbb{R}^{nk \times nk}.$$

Now since  $(\mathbf{\Delta}_0^\top \mathbf{S})^\top = \mathbf{\Delta}_0^\top \mathbf{S}$ , we have by Corollary S1 that the right inverse of  $\mathbf{G}$  is given by

$$\mathbf{G}_{\text{right}}^{-1} = (2\mathbf{W}_0^{-1} - \mathbf{S} (\mathbf{S}^\top \mathbf{W}_0 \mathbf{S})^{-1} \mathbf{S}^\top) \otimes (\mathbf{S}^\top \mathbf{W}_0 \mathbf{S})^{-1}$$

and therefore using (S9) and (S32) we obtain

$$\begin{aligned}\text{svec}(\mathbf{A}_k) &= \text{svec}(\mathbf{A}_0) + (\mathbf{W}_0 \otimes \mathbf{W}_0) \mathbf{Q} (\mathbf{I}_n \otimes \mathbf{S}) \mathbf{G}_{\text{right}}^{-1} \text{vec}(\mathbf{\Delta}_0^\top) \\ &= \text{svec}(\mathbf{A}_0) + \mathbf{Q} \mathbf{Q}^\top (\mathbf{W}_0 \otimes \mathbf{W}_0) \mathbf{Q} (2\mathbf{W}_0^{-1} - \mathbf{S} (\mathbf{S}^\top \mathbf{W}_0 \mathbf{S})^{-1} \mathbf{S}^\top) \otimes \mathbf{S} (\mathbf{S}^\top \mathbf{W}_0 \mathbf{S})^{-1} \text{vec}(\mathbf{\Delta}_0^\top) \\ &= \text{svec}(\mathbf{A}_0) + \mathbf{Q} \frac{1}{2} ((2\mathbf{I} - \mathbf{U} \mathbf{S}^\top) \otimes \mathbf{U} + \mathbf{U} \boxtimes (2\mathbf{I} - \mathbf{U} \mathbf{S}^\top)) \text{vec}(\mathbf{\Delta}_0^\top) \\ &= \text{svec}(\mathbf{A}_0) + \text{svec}(\mathbf{U} \mathbf{\Delta}_0^\top (\mathbf{I} - \frac{1}{2} \mathbf{U} \mathbf{S}^\top)^\top + (\mathbf{I} - \frac{1}{2} \mathbf{U} \mathbf{S}^\top) \mathbf{\Delta}_0 \mathbf{U}^\top) \\ &= \text{svec}(\mathbf{A}_0 + \mathbf{\Delta}_0 \mathbf{U}^\top + \mathbf{U} \mathbf{\Delta}_0^\top - \mathbf{U} \mathbf{S}^\top \mathbf{\Delta}_0 \mathbf{U}^\top).\end{aligned}$$

Further by definition it holds that

$$\mathbf{U} \mathbf{V}^\top + \mathbf{V} \mathbf{U}^\top = \mathbf{U} \mathbf{\Delta}_0^\top (\mathbf{I}_n - \frac{1}{2} \mathbf{S} \mathbf{U}^\top) + (\mathbf{I}_n - \frac{1}{2} \mathbf{U} \mathbf{S}^\top) \mathbf{\Delta}_0 \mathbf{U}^\top = \mathbf{\Delta}_0 \mathbf{U}^\top + \mathbf{U} \mathbf{\Delta}_0^\top - \mathbf{U} \mathbf{S}^\top \mathbf{\Delta}_0 \mathbf{U}^\top.$$

For the covariance we obtain using the right inverse of the Gram matrix and (S35) that

$$\begin{aligned}\mathbf{\Sigma}_k &= \mathbf{W}_0 \otimes \mathbf{W}_0 - (\mathbf{W}_0 \otimes \mathbf{W}_0) \mathbf{Q} (\mathbf{I}_n \otimes \mathbf{S}) \mathbf{G}^{-1} (\mathbf{I}_n \otimes \mathbf{S}^\top) \mathbf{Q}^\top (\mathbf{W}_0 \otimes \mathbf{W}_0) \\ &= \mathbf{W}_0 \otimes \mathbf{W}_0 - (2\mathbf{W}_0 - \mathbf{W}_0 \mathbf{S} (\mathbf{S}^\top \mathbf{W}_0 \mathbf{S})^{-1} \mathbf{S}^\top \mathbf{W}_0) \otimes (\mathbf{W}_0 \mathbf{S} (\mathbf{S}^\top \mathbf{W}_0 \mathbf{S})^{-1} \mathbf{S}^\top \mathbf{W}_0) \\ &= (\mathbf{W}_0 - \mathbf{W}_0 \mathbf{S} (\mathbf{S}^\top \mathbf{W}_0 \mathbf{S})^{-1} \mathbf{S}^\top \mathbf{W}_0) \otimes (\mathbf{W}_0 - \mathbf{W}_0 \mathbf{S} (\mathbf{S}^\top \mathbf{W}_0 \mathbf{S})^{-1} \mathbf{S}^\top \mathbf{W}_0) \\ &= \mathbf{W}_0 (\mathbf{I}_n - \mathbf{S} \mathbf{U}^\top) \otimes \mathbf{W}_0 (\mathbf{I}_n - \mathbf{S} \mathbf{U}^\top).\end{aligned}$$

□

## S4.2 Matrix-variate Prior Construction

From a practical point of view it is important to be able to construct a prior for  $\mathbf{A}$  and  $\mathbf{H}$  from an initial guess  $\mathbf{x}_0$  for the solution. This reduces down to finding  $\mathbf{A}_0$  and  $\mathbf{H}_0$  symmetric positive definite, such that  $\mathbf{A}_0 = \mathbf{H}_0^{-1}$  and  $\mathbf{x}_0 = \mathbf{H}_0 \mathbf{b}$  for the covariance class derived in Section 3. We provide a computationally efficient construction of such a prior here.

**Proposition S5**

Let  $\mathbf{x}_0 \in \mathbb{R}^n$  and  $\mathbf{b} \in \mathbb{R}^n \setminus \{0\}$ . Assume  $\mathbf{x}_0^\top \mathbf{b} > 0$ , then for  $\alpha < \frac{\mathbf{b}^\top \mathbf{x}_0}{\mathbf{b}^\top \mathbf{b}}$ ,

$$\mathbf{H}_0 = \alpha \mathbf{I} + \frac{1}{(\mathbf{x}_0 - \alpha \mathbf{b})^\top \mathbf{b}} (\mathbf{x}_0 - \alpha \mathbf{b})(\mathbf{x}_0 - \alpha \mathbf{b})^\top$$

is symmetric positive definite and  $\mathbf{H}_0 \mathbf{b} = \mathbf{x}_0$ . Further it holds that

$$\mathbf{A}_0 = \mathbf{H}_0^{-1} = \alpha^{-1} \mathbf{I} - \frac{\alpha^{-1}}{(\mathbf{x}_0 - \alpha \mathbf{b})^\top \mathbf{x}_0} (\mathbf{x}_0 - \alpha \mathbf{b})(\mathbf{x}_0 - \alpha \mathbf{b})^\top.$$

If  $\mathbf{x}_0^\top \mathbf{b} < 0$  or  $\mathbf{x}_0^\top \mathbf{b} = 0$ , then for  $\mathbf{x}_1 = -\mathbf{x}_0$  or  $\mathbf{x}_1 = \frac{\mathbf{b}^\top \mathbf{b}}{\mathbf{b}^\top \mathbf{A} \mathbf{b}} \mathbf{b}$  respectively, it holds that  $\|\mathbf{x}_1 - \mathbf{x}_*\|_{\mathbf{A}}^2 < \|\mathbf{x}_0 - \mathbf{x}_*\|_{\mathbf{A}}^2$ , i.e.  $\mathbf{x}_1$  is a strictly better initialization than  $\mathbf{x}_0$ .

*Proof.* Let  $\mathbf{H}_0$  as above. Then  $\mathbf{H}_0 \mathbf{b} = \alpha \mathbf{b} + \mathbf{x}_0 - \alpha \mathbf{b} = \mathbf{x}_0$ . The second term of the sum in the form of  $\mathbf{H}_0$  is of rank 1. Its non-zero eigenvalue is given by

$$\lambda = \frac{1}{(\mathbf{x}_0 - \alpha \mathbf{b})^\top \mathbf{b}} (\mathbf{x}_0 - \alpha \mathbf{b})^\top (\mathbf{x}_0 - \alpha \mathbf{b}) = \frac{1}{\mathbf{x}_0^\top \mathbf{b} - \alpha \mathbf{b}^\top \mathbf{b}} \|\mathbf{x}_0 - \alpha \mathbf{b}\|_2^2 \geq 0$$

since by assumption  $\mathbf{x}_0^\top \mathbf{b} > 0$  and  $\alpha < \frac{\mathbf{b}^\top \mathbf{x}_0}{\mathbf{b}^\top \mathbf{b}}$ . Now by Weyl's theorem it holds that  $\lambda_{\min}(\mathbf{A}) + \lambda_{\min}(\mathbf{E}) \leq \lambda_{\min}(\mathbf{A} + \mathbf{E})$  and therefore  $\mathbf{H}_0$  is positive definite. By the matrix inversion lemma we have for  $\gamma = \frac{\alpha^{-1}}{(\mathbf{x}_0 - \alpha \mathbf{b})^\top \mathbf{b}}$  that

$$\begin{aligned} \mathbf{A}_0 &= \mathbf{H}_0^{-1} = \alpha^{-1} \left( \mathbf{I} - \frac{\gamma}{1 + \gamma \|\mathbf{x}_0 - \alpha \mathbf{b}\|_2^2} (\mathbf{x}_0 - \alpha \mathbf{b})(\mathbf{x}_0 - \alpha \mathbf{b})^\top \right) \\ &= \alpha^{-1} \mathbf{I} - \frac{\alpha^{-2}}{(\mathbf{x}_0 - \alpha \mathbf{b})^\top \mathbf{b} + \alpha^{-1} \|\mathbf{x}_0 - \alpha \mathbf{b}\|_2^2} (\mathbf{x}_0 - \alpha \mathbf{b})(\mathbf{x}_0 - \alpha \mathbf{b})^\top \\ &= \alpha^{-1} \mathbf{I} - \frac{\alpha^{-1}}{(\mathbf{x}_0 - \alpha \mathbf{b})^\top \mathbf{x}_0} (\mathbf{x}_0 - \alpha \mathbf{b})(\mathbf{x}_0 - \alpha \mathbf{b})^\top. \end{aligned}$$

Finally, we obtain

$$\|\mathbf{x}_0 - \mathbf{x}_*\|_{\mathbf{A}}^2 = (\mathbf{x}_0 - \mathbf{A}^{-1} \mathbf{b})^\top \mathbf{A} (\mathbf{x}_0 - \mathbf{A}^{-1} \mathbf{b}) = \mathbf{x}_0^\top \mathbf{A} \mathbf{x}_0 + \mathbf{b}^\top \mathbf{A}^{-1} \mathbf{b} - 2 \mathbf{b}^\top \mathbf{x}_0.$$

Therefore if either  $\mathbf{x}_0^\top \mathbf{b} < 0$  or  $\mathbf{x}_0^\top \mathbf{b} = 0$ , then  $\mathbf{x}_1 = -\mathbf{x}_0$  or  $\mathbf{x}_1 = \frac{\mathbf{b}^\top \mathbf{b}}{\mathbf{b}^\top \mathbf{A} \mathbf{b}} \mathbf{b}$ , respectively are closer to  $\mathbf{x}_*$  in  $\mathbf{A}$  norm by positive definiteness of  $\mathbf{A}$ . This concludes the proof.  $\square$

### S4.3 Stopping Criteria

In addition to the classic stopping criteria  $\|\mathbf{A} \mathbf{x}_k - \mathbf{b}\|_2 \leq \max(\delta_{\text{tol}} \|\mathbf{b}\|_2, \delta_{\text{atol}})$  it is natural from a probabilistic viewpoint to use the induced posterior covariance of  $\mathbf{x}$ . Let  $\mathbf{M} \in \mathbb{R}^{n \times n}$  be a positive-definite matrix, then by linearity and the cyclic property of the trace it holds that

$$\begin{aligned} \mathbb{E}_{\mathbf{x}_*} [\|\mathbf{x}_* - \mathbb{E}[\mathbf{x}]\|_{\mathbf{M}}^2] &= \mathbb{E}_{\mathbf{x}_*} [(\mathbf{x}_* - \mathbb{E}[\mathbf{x}])^\top \mathbf{M} (\mathbf{x}_* - \mathbb{E}[\mathbf{x}])] \\ &= \text{tr}(\mathbb{E}_{\mathbf{x}_*} [(\mathbf{x}_* - \mathbb{E}[\mathbf{x}])^\top \mathbf{M} (\mathbf{x}_* - \mathbb{E}[\mathbf{x}])]) \\ &= \mathbb{E}_{\mathbf{x}_*} [\text{tr}((\mathbf{x}_* - \mathbb{E}[\mathbf{x}])^\top \mathbf{M} (\mathbf{x}_* - \mathbb{E}[\mathbf{x}]))] \\ &= \mathbb{E}_{\mathbf{x}_*} [\mathbf{M} \text{tr}((\mathbf{x}_* - \mathbb{E}[\mathbf{x}]) (\mathbf{x}_* - \mathbb{E}[\mathbf{x}])^\top)] \\ &= \text{tr}(\mathbf{M} \mathbb{E}_{\mathbf{x}_*} [(\mathbf{x}_* - \mathbb{E}[\mathbf{x}]) (\mathbf{x}_* - \mathbb{E}[\mathbf{x}])^\top]) \\ &= \text{tr}(\mathbf{M} (\text{Cov}[\mathbf{x}_* - \mathbb{E}[\mathbf{x}]] + (\mathbb{E}_{\mathbf{x}_*} [\mathbf{x}_*] - \mathbb{E}[\mathbf{x}])^\top (\mathbb{E}_{\mathbf{x}_*} [\mathbf{x}_*] - \mathbb{E}[\mathbf{x}]))) \\ &= \text{tr}(\mathbf{M} \text{Cov}[\mathbf{x}_*]) + \|\mathbb{E}_{\mathbf{x}_*} [\mathbf{x}_*] - \mathbb{E}[\mathbf{x}]\|_{\mathbf{M}}^2. \end{aligned}$$

Assuming calibration holds, i.e.  $\mathbf{x}_* \sim \mathcal{N}(\mathbb{E}[\mathbf{x}], \text{Cov}[\mathbf{x}])$ , we can bound the (relative) error by terminating when  $\text{tr}(\mathbf{M} \text{Cov}[\mathbf{x}]) \leq \max(\delta_{\text{tol}} \|\mathbf{b}\|, \delta_{\text{atol}})$  either in  $l_2$ -norm for  $\mathbf{M} = \mathbf{I}$  or in  $\mathbf{A}$ -norm for  $\mathbf{M} = \mathbf{A}$ .

We can efficiently evaluate the required  $\text{tr}(\mathbf{M} \text{Cov}[\mathbf{x}])$  without ever forming  $\text{Cov}[\mathbf{x}]$  in memory from already computed quantities. At iteration  $k$  we have  $\text{Cov}[\mathbf{x}] = \text{Cov}[\mathbf{H} \mathbf{b}] = \frac{1}{2} (\mathbf{W}_k^\mathbf{H} (\mathbf{b}^\top \mathbf{W}_k^\mathbf{H} \mathbf{b}) + (\mathbf{W}_k^\mathbf{H} \mathbf{b})(\mathbf{W}_k^\mathbf{H} \mathbf{b})^\top)$  and therefore

$$\text{tr}(\mathbf{M} \text{Cov}[\mathbf{x}]) = \frac{1}{2} ((\mathbf{b}^\top \mathbf{W}_k^\mathbf{H} \mathbf{b}) \text{tr}(\mathbf{M} \mathbf{W}_k^\mathbf{H}) + (\mathbf{W}_k^\mathbf{H} \mathbf{b})^\top \mathbf{M} (\mathbf{W}_k^\mathbf{H} \mathbf{b})).$$

Given the update for the covariance of the inverse view, we obtain the following recursion for its trace

$$\text{tr}(M\mathbf{W}_k^{\mathbf{H}}) = \text{tr}(M\mathbf{W}_{k-1}^{\mathbf{H}}) - \frac{1}{\mathbf{y}_k^{\top}\mathbf{W}_{k-1}^{\mathbf{H}}\mathbf{y}_k} \text{tr}((\mathbf{W}_{k-1}^{\mathbf{H}}\mathbf{y}_k)^{\top}M(\mathbf{W}_{k-1}^{\mathbf{H}}\mathbf{y}_k)).$$

Computing the trace in this iterative fashion adds at most three matrix-vector products and three inner products for arbitrary  $M$  all other quantities are computed for the covariance update anyhow.

For our proposed covariance class (3) we obtain for  $M = \mathbf{I}$  and  $\Psi = \psi\mathbf{I}$  that

$$\begin{aligned} \text{tr}(\mathbf{W}_0^{\mathbf{H}}) &= \text{tr}(\mathbf{A}_0^{-1}\mathbf{Y}(\mathbf{Y}^{\top}\mathbf{A}_0^{-1}\mathbf{Y})^{-1}\mathbf{Y}^{\top}\mathbf{A}_0^{-1} + (\mathbf{I} - \mathbf{Y}(\mathbf{Y}^{\top}\mathbf{Y})^{-1}\mathbf{Y}^{\top})\Psi(\mathbf{I} - \mathbf{Y}(\mathbf{Y}^{\top}\mathbf{Y})^{-1}\mathbf{Y}^{\top})) \\ &= \text{tr}((\mathbf{Y}^{\top}\mathbf{A}_0^{-1}\mathbf{Y})^{-1}\mathbf{Y}^{\top}\mathbf{A}_0^{-1}\mathbf{A}_0^{-1}\mathbf{Y}) + \psi \text{tr}((\mathbf{I} - \mathbf{Y}(\mathbf{Y}^{\top}\mathbf{Y})^{-1}\mathbf{Y}^{\top})(\mathbf{I} - \mathbf{Y}(\mathbf{Y}^{\top}\mathbf{Y})^{-1}\mathbf{Y}^{\top})) \\ &= \text{tr}((\mathbf{Y}^{\top}\mathbf{A}_0^{-1}\mathbf{Y})^{-1}\mathbf{Y}^{\top}\mathbf{A}_0^{-1}\mathbf{A}_0^{-1}\mathbf{Y}) + \psi \text{tr}(\mathbf{I} - \mathbf{Y}(\mathbf{Y}^{\top}\mathbf{Y})^{-1}\mathbf{Y}^{\top}) \\ &= \text{tr}((\mathbf{Y}^{\top}\mathbf{A}_0^{-1}\mathbf{Y})^{-1}\mathbf{Y}^{\top}\mathbf{A}_0^{-1}\mathbf{A}_0^{-1}\mathbf{Y}) + \psi(n - k), \end{aligned}$$

which for a scalar prior mean  $\mathbf{A}_0 = \alpha\mathbf{I}$  reduces to  $\text{tr}(\mathbf{W}_0^{\mathbf{H}}) = \alpha^{-1}k + \psi(n - k)$ .

#### S4.4 Implementation

In order to maintain numerical stability when performing low rank updates to symmetric positive definite matrices, as is the case in Algorithm 1 for the mean and covariance estimates, it is advantageous use a representation based on the Cholesky decomposition. One can perform the rank-2 update for the mean estimate and the rank-1 downdate for the covariance in Corollary S3 in each iteration of the algorithm for their respective Cholesky factors instead (see also Seeger [22]). The rank-2 update can be seen as a combination of a rank-1 up- and downdate by recognizing that

$$\mathbf{u}\mathbf{v}^{\top} + \mathbf{v}\mathbf{u}^{\top} = \frac{1}{2}((\mathbf{u} + \mathbf{v})(\mathbf{u} + \mathbf{v})^{\top} - (\mathbf{u} - \mathbf{v})(\mathbf{u} - \mathbf{v})^{\top}).$$

Similar updates arise in Quasi-Newton methods for the approximate (inverse) Hessian [10]. Having Cholesky factors of the mean and covariance available has the additional advantage that downstream sampling or the evaluation of the probability density function is computationally cheap.

## S5 Theoretical Properties: Proofs for Section 2.3

In this section we provide detailed proofs for the theoretical results on convergence and the connection of Algorithm 1 to the method of conjugate gradients. We restate each theorem here as a reference to the reader. We begin by proving an intermediate result giving an interpretation to the posterior mean of  $\mathbf{A}$  and  $\mathbf{H}$  at each step of the method.

**Proposition S6** (Subspace Equivalency)

Let  $\mathbf{A}_k$  and  $\mathbf{H}_k$  be the posterior means defined as in Section 2.1 and assume  $\mathbf{A}_0$  and  $\mathbf{H}_0$  are symmetric. Then for  $1 \leq k \leq n$  it holds that

$$\mathbf{A}_k\mathbf{S} = \mathbf{Y} \quad \text{and} \quad \mathbf{H}_k\mathbf{Y} = \mathbf{S}, \tag{S36}$$

i.e.  $\mathbf{A}_k$  and  $\mathbf{H}_k$  act like  $\mathbf{A}$  and  $\mathbf{A}^{-1}$  on the spaces spanned by the actions  $\mathbf{S}$ , respectively the observations  $\mathbf{Y}$ .

*Proof.* Since  $\mathbf{A}_0$  and  $\mathbf{H}_0$  are symmetric so are the expressions  $\Delta_{\mathbf{A}}\mathbf{S}$  and  $\Delta_{\mathbf{H}}^{\top}\mathbf{Y}$ . We have that

$$\begin{aligned} \mathbf{A}_k\mathbf{S} &= (\mathbf{A}_0 + \Delta_{\mathbf{A}}\mathbf{U}_{\mathbf{A}}^{\top} + \mathbf{U}_{\mathbf{A}}\Delta_{\mathbf{A}}^{\top} - \mathbf{U}_{\mathbf{A}}\mathbf{S}^{\top}\Delta_{\mathbf{A}}\mathbf{U}_{\mathbf{A}}^{\top})\mathbf{S} \\ &= \mathbf{A}_0\mathbf{S} + \Delta_{\mathbf{A}}\mathbf{I} + \mathbf{U}_{\mathbf{A}}\Delta_{\mathbf{A}}^{\top}\mathbf{S} - \mathbf{U}_{\mathbf{A}}\mathbf{S}^{\top}\Delta_{\mathbf{A}}\mathbf{I} \\ &= \mathbf{A}_0\mathbf{S} + \mathbf{Y} - \mathbf{A}_0\mathbf{S} \\ &= \mathbf{Y}. \end{aligned}$$

In the case of the inverse model we obtain

$$\begin{aligned} \mathbf{H}_k\mathbf{Y} &= (\mathbf{H}_0 + \Delta_{\mathbf{H}}\mathbf{U}_{\mathbf{H}}^{\top} + \mathbf{U}_{\mathbf{H}}\Delta_{\mathbf{H}}^{\top} - \mathbf{U}_{\mathbf{H}}\mathbf{Y}^{\top}\Delta_{\mathbf{H}}\mathbf{U}_{\mathbf{H}}^{\top})\mathbf{Y} \\ &= \mathbf{H}_0\mathbf{Y} + \Delta_{\mathbf{H}}\mathbf{I} + \mathbf{U}_{\mathbf{H}}\Delta_{\mathbf{H}}^{\top}\mathbf{Y} - \mathbf{U}_{\mathbf{H}}\mathbf{Y}^{\top}\Delta_{\mathbf{H}}\mathbf{I} \\ &= \mathbf{H}_0\mathbf{Y} + \mathbf{S} - \mathbf{H}_0\mathbf{Y} \\ &= \mathbf{S} \end{aligned}$$

□

## S5.1 Conjugate Directions Method

**Theorem 1** (Conjugate Directions Method)

Given a prior  $p(\mathbf{H}) = \mathcal{N}(\mathbf{H}; \mathbf{H}_0, \mathbf{W}_0^{\mathbf{H}} \otimes \mathbf{W}_0^{\mathbf{H}})$  such that  $\mathbf{H}_0, \mathbf{W}_0^{\mathbf{H}} \in \mathbb{R}_{\text{sym}}^{n \times n}$  positive definite, then actions  $\mathbf{s}_i$  of Algorithm 1 are  $\mathbf{A}$ -conjugate, i.e. for  $0 \leq i, j \leq k$  with  $i \neq j$  it holds that  $\mathbf{s}_i^{\top} \mathbf{A} \mathbf{s}_j = 0$ .

*Proof.* Since  $\mathbf{H}_0$  is assumed to be symmetric, the form of the posterior mean in Section 2.1 implies that  $\mathbf{H}_k$  is symmetric for all  $1 \leq k \leq n$ . Now conjugacy is shown by induction. To that end, first consider the base case  $k = 2$ . We have

$$\begin{aligned} \mathbf{s}_2^{\top} \mathbf{A} \mathbf{s}_1 &= -\mathbf{r}_1^{\top} \mathbf{H}_1 \mathbf{A} \mathbf{s}_1 = -(\mathbf{r}_0^{\top} + \alpha_1 \mathbf{y}_1^{\top}) \mathbf{H}_1 \mathbf{A} \mathbf{s}_1 = -\left( \mathbf{r}_0^{\top} \mathbf{H}_1 - \frac{\mathbf{s}_1^{\top} \mathbf{r}_0}{\mathbf{s}_1^{\top} \mathbf{y}_1} \mathbf{y}_1^{\top} \mathbf{H}_1 \right) \mathbf{y}_1 \\ &= -\mathbf{r}_0^{\top} \mathbf{s}_1 + \mathbf{s}_1^{\top} \mathbf{r}_0 = 0 \end{aligned}$$

where we used (S36) and the definition of  $\alpha_i$  in Algorithm 1. Now for the induction step, assume that  $\mathbf{s}_i^{\top} \mathbf{A} \mathbf{s}_j = 0$  for all  $i \neq j$  such that  $1 \leq i, j \leq k$ . We obtain for  $1 \leq j \leq k$  that

$$\begin{aligned} \mathbf{s}_{k+1}^{\top} \mathbf{A} \mathbf{s}_j &= -\mathbf{r}_k^{\top} \mathbf{H}_k \mathbf{A} \mathbf{s}_j = -\left( \sum_{1 \leq l \leq k} \alpha_l \mathbf{y}_l + \mathbf{r}_0 \right)^{\top} \mathbf{H}_k \mathbf{y}_j = -\sum_{1 \leq l \leq k} \alpha_l \mathbf{y}_l^{\top} \mathbf{s}_j - \mathbf{r}_0^{\top} \mathbf{s}_j \\ &= -\alpha_j \mathbf{y}_j^{\top} \mathbf{s}_j - \mathbf{r}_0^{\top} \mathbf{s}_j = \mathbf{s}_j^{\top} \mathbf{r}_{j-1} - \mathbf{r}_0^{\top} \mathbf{s}_j = \mathbf{s}_j^{\top} \left( \sum_{1 \leq l < j} \alpha_l \mathbf{y}_l + \mathbf{r}_0 \right) - \mathbf{r}_0^{\top} \mathbf{s}_j \\ &= \mathbf{s}_j^{\top} \mathbf{r}_0 - \mathbf{r}_0^{\top} \mathbf{s}_j = 0 \end{aligned}$$

where we used the update equation of the residual  $\mathbf{r}_i$  in Algorithm 1, the definition of  $\alpha_i$ , the induction hypothesis and (S36). This proves the statement.  $\square$

## S5.2 Relationship to the Conjugate Gradient Method

**Theorem 2** (Connection to the Conjugate Gradient Method)

Given a scalar prior mean  $\mathbf{A}_0 = \mathbf{H}_0^{-1} = \alpha \mathbf{I}$  with  $\alpha > 0$ , assume (1) and (2) hold, then the iterates  $\mathbf{x}_i$  of Algorithm 1 are identical to the ones produced by the conjugate gradient method.

*Proof.* The proof outlined here is closely related to the proofs connecting Quasi-Newton methods to the conjugate gradient method [47, 11], but makes different assumptions on the prior distribution.

We begin by recognizing that the choice of step length  $\alpha_i$  in Algorithm 1 is identical to the one in the conjugate gradient method [10]. Hence, it suffices to show that  $\mathbf{s}_i \propto \mathbf{s}_i^{\text{CG}}$ . Theorem 1 established that Algorithm 1 is a conjugate directions method. Now by assumption  $\mathbf{A}_0 = \alpha \mathbf{I}$  and  $\mathbf{H}_0 = \mathbf{A}_0^{-1}$ , therefore  $\mathbf{s}_1 = -\alpha \mathbf{I} \mathbf{r}_0 \propto -\mathbf{r}_0 = \mathbf{s}_1^{\text{CG}}$ . It suffices show that  $\mathbf{s}_i$  lies in the Krylov space  $\mathcal{K}_i(\mathbf{A}, \mathbf{r}_0) = \{\mathbf{r}_0, \mathbf{A} \mathbf{r}_0, \dots, \mathbf{A}^{i-1} \mathbf{r}_0\}$  for all  $0 < i \leq n$ . This completes the argument, since  $\mathcal{K}_i(\mathbf{A}, \mathbf{r}_0)$  is an  $i$ -dimensional subspace of  $\mathbb{R}^n$  and thus  $\mathbf{A}$ -conjugacy uniquely determines the search directions up to scaling, as  $\mathbf{A}$  is positive definite.

To complete the proof we proceed as follows. The posterior mean of the inverse model  $\mathbf{H}_{i-1}$  at step  $i - 1$  maps an arbitrary vector  $\mathbf{v} \in \mathbb{R}^n$  to  $\text{span}(\mathbf{H}_0 \mathbf{v}, \mathbf{H}_0 \mathbf{Y}_{1:i-1}, \mathbf{S}_{1:i-1}, \mathbf{W}_0^{\mathbf{H}} \mathbf{Y}_{1:i-1})$ . This follows directly from its form in given in Section 2.1. By assumption  $\mathbf{H}_0 = \mathbf{A}_0^{-1} = \alpha^{-1} \mathbf{I}$ , therefore using (1) and (2) we have  $\text{span}(\mathbf{W}_0^{\mathbf{H}} \mathbf{Y}_{1:i-1}) = \text{span}(\mathbf{Y}_{1:i-1})$ . This implies  $\mathbf{H}_{i-1}$  maps to  $\text{span}(\mathbf{v}, \mathbf{S}_{1:i-1}, \mathbf{Y}_{1:i-1})$  and thus  $\mathbf{s}_i \in \text{span}(\mathbf{r}_{i-1}, \mathbf{S}_{1:i-1}, \mathbf{Y}_{1:i-1})$ . We will now show that  $\text{span}(\mathbf{r}_{i-1}, \mathbf{S}_{1:i-1}, \mathbf{Y}_{1:i-1}) \subset \mathcal{K}_i(\mathbf{A}, \mathbf{r}_0)$  by induction, completing the argument.

We begin with the base case. Since  $\mathbf{H}_0$  is assumed to be scalar, we have  $\mathbf{s}_1 \propto \mathbf{r}_0 \in \mathcal{K}_0(\mathbf{A}, \mathbf{r}_0)$  and therefore  $\mathbf{y}_1 = \mathbf{A} \mathbf{s}_1$  and  $\mathbf{r}_1 = \mathbf{r}_0 + \alpha_1 \mathbf{y}_1$  are in  $\mathcal{K}_1(\mathbf{A}, \mathbf{r}_0)$ . For the induction step assume  $\text{span}(\mathbf{r}_{i-1}, \mathbf{S}_{1:i-1}, \mathbf{Y}_{1:i-1}) \subset \mathcal{K}_i(\mathbf{A}, \mathbf{r}_0)$ . The definition of the policy of Algorithm 1 gives

$$\mathbf{s}_i = -\mathbb{E}[\mathbf{H}] \mathbf{r}_{i-1} \propto \mathbf{H}_{i-1} \mathbf{r}_{i-1} \in \text{span}(\mathbf{r}_{i-1}, \mathbf{S}_{1:i-1}, \mathbf{Y}_{1:i-1}) \subset \mathcal{K}_i(\mathbf{A}, \mathbf{r}_0),$$

where we used the induction hypothesis. This implies that  $\mathbf{y}_i = \mathbf{A} \mathbf{s}_i \in \mathcal{K}_{i+1}(\mathbf{A}, \mathbf{r}_0)$  and  $\mathbf{r}_i = \mathbf{r}_{i-1} + \alpha_i \mathbf{y}_i \in \mathcal{K}_{i+1}(\mathbf{A}, \mathbf{r}_0)$  by the definition of the Krylov space. Therefore,  $\text{span}(\mathbf{r}_i, \mathbf{S}_{1:i}, \mathbf{Y}_{1:i}) \subset \mathcal{K}_{i+1}(\mathbf{A}, \mathbf{r}_0)$ . This completes the proof.  $\square$

## S6 Prior Covariance Class: Proofs for Section 3

### S6.1 Hereditary Positive-Definiteness

**Proposition 1** (Hereditary Positive Definiteness [30, 18])

Let  $\mathbf{A}_0 \in \mathbb{R}_{\text{sym}}^{n \times n}$  be positive definite. Assume the actions  $\mathbf{S}$  are  $\mathbf{A}$ -conjugate and  $\mathbf{W}_0^{\mathbf{A}} \mathbf{S} = \mathbf{Y}$ , then for  $i \in \{0, \dots, k-1\}$  it holds that  $\mathbf{A}_{i+1}$  is symmetric positive definite.

*Proof.* This is shown in Hennig and Kiefel [18]. We give an identical proof in our notation as a reference to the reader. By Theorem 7.5 in Dennis and Moré [30] it holds that if  $\mathbf{A}_i$  is positive definite and  $\mathbf{s}_{i+1}^{\top} \mathbf{W}_i^{\mathbf{A}} \mathbf{s}_{i+1} \neq 0$ , then  $\mathbf{A}_{i+1}$  is positive definite if and only if  $\det(\mathbf{A}_{i+1}) > 0$ . By the matrix determinant lemma and the recursive formulation of the posterior we have

$$\det(\mathbf{A}_{i+1}) = \det(\mathbf{A}_i) \left( \frac{1}{(\mathbf{s}_{i+1}^{\top} \mathbf{W}_i^{\mathbf{A}} \mathbf{s}_{i+1})^2} \left( (\mathbf{y}_{i+1}^{\top} \mathbf{A}_i^{-1} \mathbf{W}_i^{\mathbf{A}} \mathbf{s}_{i+1})^2 - (\mathbf{y}_{i+1}^{\top} \mathbf{A}_i^{-1} \mathbf{y}_{i+1})(\mathbf{s}_{i+1}^{\top} \mathbf{W}_i^{\mathbf{A}} \mathbf{A}_i^{-1} \mathbf{W}_i^{\mathbf{A}} \mathbf{s}_{i+1}) + (\mathbf{s}_{i+1}^{\top} \mathbf{W}_i^{\mathbf{A}} \mathbf{A}_i^{-1} \mathbf{W}_i^{\mathbf{A}} \mathbf{s}_{i+1})(\mathbf{y}_{i+1}^{\top} \mathbf{s}_{i+1}) \right) \right)$$

Hence it suffices to show that

$$0 < (\mathbf{y}_{i+1}^{\top} \mathbf{A}_i^{-1} \mathbf{W}_i^{\mathbf{A}} \mathbf{s}_{i+1})^2 - (\mathbf{y}_{i+1}^{\top} \mathbf{A}_i^{-1} \mathbf{y}_{i+1})(\mathbf{s}_{i+1}^{\top} \mathbf{W}_i^{\mathbf{A}} \mathbf{A}_i^{-1} \mathbf{W}_i^{\mathbf{A}} \mathbf{s}_{i+1}) + (\mathbf{s}_{i+1}^{\top} \mathbf{W}_i^{\mathbf{A}} \mathbf{A}_i^{-1} \mathbf{W}_i^{\mathbf{A}} \mathbf{s}_{i+1})(\mathbf{y}_{i+1}^{\top} \mathbf{s}_{i+1}),$$

which simplifies to

$$\mathbf{y}_{i+1}^{\top} \mathbf{A}_i^{-1} \mathbf{y}_{i+1} - \frac{(\mathbf{y}_{i+1}^{\top} \mathbf{A}_i^{-1} \mathbf{W}_i^{\mathbf{A}} \mathbf{s}_{i+1})^2}{\mathbf{s}_{i+1}^{\top} \mathbf{W}_i^{\mathbf{A}} \mathbf{A}_i^{-1} \mathbf{W}_i^{\mathbf{A}} \mathbf{s}_{i+1}} < \mathbf{y}_{i+1}^{\top} \mathbf{s}_{i+1}$$

Now by  $\mathbf{W}_0^{\mathbf{A}} \mathbf{S} = \mathbf{Y}$ , we have  $\mathbf{W}_i^{\mathbf{A}} \mathbf{s}_{i+1} = \mathbf{W}_0^{\mathbf{A}} \mathbf{s}_{i+1} = \mathbf{y}_{i+1}$  and the above reduces to

$$0 < \mathbf{s}_{i+1}^{\top} \mathbf{A} \mathbf{s}_{i+1},$$

which is fulfilled by the assumption that  $\mathbf{A}$  is positive definite. Thus  $\mathbf{A}_{i+1}$  is positive definite. Symmetry follows immediately from the form of the posterior mean.  $\square$

### S6.2 Posterior Correspondence

**Definition 1**

Let  $\mathbf{A}_i$  and  $\mathbf{H}_i$  be the means of  $\mathbf{A}$  and  $\mathbf{H}$  at step  $i$ . We say a prior induces *posterior correspondence* if

$$\mathbf{A}_i^{-1} = \mathbf{H}_i \tag{S37}$$

for all steps  $0 \leq i \leq k$  of the solver. If only

$$\mathbf{A}_i^{-1} \mathbf{Y} = \mathbf{H}_i \mathbf{Y}, \tag{S38}$$

we say that *weak posterior correspondence* holds.

#### S6.2.1 Matrix-variate Normal Prior

We begin by establishing posterior correspondence in the case of general matrix-variate normal priors, i.e. the inference setting detailed in Corollary S2. We begin by proving a general non-constructive condition and close with a sufficient condition for correspondence with limits the possible choices of covariance factors to a specific class.

**Lemma S1** (General Correspondence)

Let  $1 \leq k \leq n$ ,  $\mathbf{W}_0^{\mathbf{A}}, \mathbf{W}_0^{\mathbf{H}}$  symmetric positive-definite and assume  $\mathbf{A}_0^{-1} = \mathbf{H}_0$ , then (S37) holds if and only if

$$0 = (\mathbf{A} \mathbf{S} - \mathbf{A}_0 \mathbf{S}) \left[ (\mathbf{S}^{\top} \mathbf{W}_0^{\mathbf{A}} \mathbf{A}_0^{-1} \mathbf{A} \mathbf{S})^{-1} \mathbf{S}^{\top} \mathbf{W}_0^{\mathbf{A}} \mathbf{A}_0^{-1} - (\mathbf{S}^{\top} \mathbf{A}^{\top} \mathbf{W}_0^{\mathbf{H}} \mathbf{A} \mathbf{S})^{-1} \mathbf{S}^{\top} \mathbf{A}^{\top} \mathbf{W}_0^{\mathbf{H}} \right]. \tag{S39}$$

*Proof.* By the matrix inversion lemma we have

$$\begin{aligned}
0 &= \mathbf{A}_k^{-1} - \mathbf{H}_k \\
&= (\mathbf{A}_0 + (\mathbf{Y} - \mathbf{A}_0\mathbf{S})(\mathbf{S}^\top \mathbf{W}_0^{\mathbf{A}} \mathbf{S})^{-1} \mathbf{S}^\top \mathbf{W}_0^{\mathbf{A}})^{-1} - \mathbf{H}_0 - (\mathbf{S} - \mathbf{H}_0\mathbf{Y})(\mathbf{Y}^\top \mathbf{W}_0^{\mathbf{H}} \mathbf{Y})^{-1} \mathbf{Y}^\top \mathbf{W}_0^{\mathbf{H}} \\
&= \mathbf{A}_0^{-1} - \mathbf{A}_0^{-1}(\mathbf{Y} - \mathbf{A}_0\mathbf{S})(\mathbf{S}^\top \mathbf{W}_0^{\mathbf{A}} \mathbf{S} + \mathbf{S}^\top \mathbf{W}_0^{\mathbf{A}} \mathbf{A}_0^{-1}(\mathbf{Y} - \mathbf{A}_0\mathbf{S}))^{-1} \mathbf{S}^\top \mathbf{W}_0^{\mathbf{A}} \mathbf{A}_0^{-1} \\
&\quad - \mathbf{A}_0^{-1} - \mathbf{A}_0^{-1}(\mathbf{A}_0\mathbf{S} - \mathbf{Y})(\mathbf{Y}^\top \mathbf{W}_0^{\mathbf{H}} \mathbf{Y})^{-1} \mathbf{Y}^\top \mathbf{W}_0^{\mathbf{H}} \\
&= -\mathbf{A}_0^{-1}(\mathbf{Y} - \mathbf{A}_0\mathbf{S}) [(\mathbf{S}^\top \mathbf{W}_0^{\mathbf{A}} \mathbf{A}_0^{-1} \mathbf{Y})^{-1} \mathbf{S}^\top \mathbf{W}_0^{\mathbf{A}} \mathbf{A}_0^{-1} - (\mathbf{Y}^\top \mathbf{W}_0^{\mathbf{H}} \mathbf{Y})^{-1} \mathbf{Y}^\top \mathbf{W}_0^{\mathbf{H}}],
\end{aligned}$$

where we used the assumption  $\mathbf{H}_0 = \mathbf{A}_0^{-1}$ . Left-multiplying with  $-\mathbf{A}_0$  and using  $\mathbf{Y} = \mathbf{A}\mathbf{S}$  completes the proof.  $\square$

**Corollary S4** (Correspondence at Convergence)

Let  $k = n$ ,  $\mathbf{H}_0 = \mathbf{A}_0^{-1}$  and assume  $\mathbf{S}$  has full rank, i.e. the linear solver has performed  $n$  linearly independent actions, then (S37) holds for any symmetric positive-definite choice of  $\mathbf{W}_0^{\mathbf{A}}$  and  $\mathbf{W}_0^{\mathbf{H}}$ .

*Proof.* By assumption,  $\mathbf{S}^\top \mathbf{W}_0^{\mathbf{A}} \mathbf{A}_0^{-1}$  and  $\mathbf{S}^\top \mathbf{A}^\top \mathbf{W}_0^{\mathbf{H}}$  are invertible. Then by Lemma S1 the correspondence condition (S37) holds.  $\square$

**Theorem S5** (Sufficient Condition for Correspondence)

Let  $1 \leq k \leq n$  arbitrary and assume  $\mathbf{H}_0 = \mathbf{A}_0^{-1}$ . Assume  $\mathbf{W}_0^{\mathbf{A}}, \mathbf{A}_0, \mathbf{W}_0^{\mathbf{H}}$  satisfy

$$0 = \mathbf{S}^\top (\mathbf{W}_0^{\mathbf{A}} \mathbf{A}_0^{-1} - \mathbf{A}^\top \mathbf{W}_0^{\mathbf{H}}) \quad (\text{S40})$$

or equivalently let  $\mathbf{B}_{\langle \mathbf{S} \rangle^\perp} \in \mathbb{R}^{n \times k}$  be a basis of the orthogonal space  $\langle \mathbf{S} \rangle^\perp$  spanned by the actions. For  $\Phi \in \mathbb{R}^{(n-k) \times n}$  arbitrary, if

$$\mathbf{W}_0^{\mathbf{H}} = \mathbf{A}^{-\top} (\mathbf{W}_0^{\mathbf{A}} \mathbf{A}_0^{-1} - \mathbf{B}_{\langle \mathbf{S} \rangle^\perp} \Phi) \quad (\text{S41})$$

and the commutation relations

$$[\mathbf{A}_0, \mathbf{A}] = \mathbf{0} \quad (\text{S42})$$

$$[\mathbf{W}_0^{\mathbf{A}}, \mathbf{A}] = \mathbf{0} \quad (\text{S43})$$

$$[\mathbf{B}_{\langle \mathbf{S} \rangle^\perp} \Phi, \mathbf{A}] = \mathbf{0} \quad (\text{S44})$$

are fulfilled, then  $\mathbf{W}_0^{\mathbf{H}}$  is symmetric and (S37) holds.

*Proof.* By assumption  $\mathbf{W}_0^{\mathbf{A}}$  is symmetric positive-definite and (S40) is equivalent to  $\mathbf{S}^\top \mathbf{W}_0^{\mathbf{A}} \mathbf{A}_0^{-1} = \mathbf{S}^\top \mathbf{A}^\top \mathbf{W}_0^{\mathbf{H}}$ , which implies (S39). Now, assumption (S40) is equivalent to columns of the difference  $\mathbf{W}_0^{\mathbf{A}} \mathbf{A}_0^{-1} - \mathbf{A}^\top \mathbf{W}_0^{\mathbf{H}}$  lying in  $L$ , i.e. we can choose a basis  $\mathbf{B}_{\langle \mathbf{S} \rangle^\perp}$  and coefficient matrix  $\Phi$  such that

$$\mathbf{W}_0^{\mathbf{A}} \mathbf{A}_0^{-1} - \mathbf{A}^\top \mathbf{W}_0^{\mathbf{H}} = \mathbf{B}_{\langle \mathbf{S} \rangle^\perp} \Phi.$$

Rearranging the above gives (S41). With the commutation relations and

$$[\mathbf{A}, \mathbf{B}] = \mathbf{0} \iff [\mathbf{A}^{-1}, \mathbf{B}] = \mathbf{0} \iff [\mathbf{A}, \mathbf{B}^{-1}] = \mathbf{0} \iff [\mathbf{A}^{-1}, \mathbf{B}^{-1}] = \mathbf{0}$$

it holds that

$$(\mathbf{W}_0^{\mathbf{H}})^\top = \mathbf{W}_0^{\mathbf{A}} \mathbf{A}_0^{-1} \mathbf{A}^{-1} - \mathbf{B}_{\langle \mathbf{S} \rangle^\perp} \Phi \mathbf{A}^{-1} = \mathbf{A}^{-\top} \mathbf{W}_0^{\mathbf{A}} \mathbf{A}_0^{-1} - \mathbf{A}^{-\top} \mathbf{B}_{\langle \mathbf{S} \rangle^\perp} \Phi = \mathbf{W}_0^{\mathbf{H}}$$

hence  $\mathbf{W}_0^{\mathbf{H}}$  is symmetric. Finally, by Lemma S1 posterior mean correspondence (S37) holds.  $\square$

If we want to ensure correspondence for all iterations, (S44) is trivially satisfied. The question now becomes what form can  $\mathbf{A}_0$  and  $\mathbf{W}_0^{\mathbf{A}}$  take in order to ensure symmetric  $\mathbf{W}_0^{\mathbf{H}}$ . This comes down to finding matrices which commute with  $\mathbf{A}$ .

**Lemma S2** (Commuting Matrices of a Symmetric Matrix)

Let  $r \in \mathbb{N}$ ,  $\mathbf{M} \in \mathbb{R}^{n \times n}$  and  $\mathbf{A} \in \mathbb{R}^{n \times n}$  symmetric. Assume  $\mathbf{M}$  has the form

$$\mathbf{M} = \mathfrak{p}_r(\mathbf{A}) = \sum_{i=0}^r c_i \mathbf{A}^i$$

for a set of coefficients  $c_i \in \mathbb{R}$ , then  $\mathbf{M}$  and  $\mathbf{A}$  commute. If  $\mathbf{A}$  has  $n$  distinct eigenvalues,  $\mathbf{M}$  is diagonalizable and  $[\mathbf{M}, \mathbf{A}] = \mathbf{0}$ , then

$$\mathbf{M} = \mathfrak{p}_{n-1}(\mathbf{A}),$$

i.e.  $\mathbf{M}$  is a polynomial in  $\mathbf{A}$  of degree at most  $n - 1$ .

*Proof.* The first result follows immediately since

$$\mathbf{W}_0^{\mathbf{A}} \mathbf{A} = \mathfrak{p}_r(\mathbf{A}) \mathbf{A} = \sum_{i=0}^r c_i \mathbf{A}^{i+1} = \mathbf{A} \mathfrak{p}_r(\mathbf{A}) = \mathbf{A} \mathbf{W}_0^{\mathbf{A}}.$$

Assume now that  $\mathbf{A}$  has  $n$  distinct eigenvalues  $\lambda_0, \dots, \lambda_{n-1}$ ,  $\mathbf{M}$  is diagonalizable and  $\mathbf{M}$  and  $\mathbf{A}$  commute. Now, if and only if  $[\mathbf{A}, \mathbf{M}] = \mathbf{0}$ , then  $\mathbf{A}$  and  $\mathbf{M}$  are simultaneously diagonalizable by Theorem 5.2 in Conrad [48], i.e. we can find a common basis in which both  $\mathbf{A}$  and  $\mathbf{M}$  are represented by diagonal matrices. Hence, the set of matrices commuting with  $\mathbf{A}$  forms an  $n$ -dimensional subspace  $\mathcal{U}_n \subset \mathbb{R}^{n \times n}$ . Now, by the first part of this proof  $\{\mathbf{I}, \mathbf{A}, \dots, \mathbf{A}^{n-1}\} \subset \mathcal{U}_n$ . It remains to be shown, that this set forms a basis of  $\mathcal{U}_n$ . By isomorphism of finite dimensional vector spaces this is equivalent to proving that

$$\{\mathbf{b}_0, \mathbf{b}_1, \dots, \mathbf{b}_{n-1}\} := \left\{ \begin{pmatrix} 1 \\ \vdots \\ 1 \end{pmatrix}, \begin{pmatrix} \lambda_0 \\ \vdots \\ \lambda_{n-1} \end{pmatrix}, \dots, \begin{pmatrix} \lambda_0^{n-1} \\ \vdots \\ \lambda_{n-1}^{n-1} \end{pmatrix} \right\}$$

forms a basis of  $\mathbb{R}^n$ . It suffices to show that all  $\mathbf{b}_i$  are independent. Assume the contrary, then  $\sum_{i=0}^{n-1} \alpha_i \mathbf{b}_i = \mathbf{0}$  for some  $\alpha_0, \dots, \alpha_{n-1} \in \mathbb{R}$ , such that not all  $\alpha_i = 0$ . This implies that the polynomial  $\sum_{i=0}^{n-1} \alpha_i x^i$  has  $n$  zeros  $\lambda_0, \dots, \lambda_{n-1}$ . This contradicts the fundamental theorem of algebra, concluding the proof.  $\square$

The above suggests that tractable choices of  $\mathbf{A}_0$  and  $\mathbf{W}_0^{\mathbf{A}}$  for the non-symmetric matrix-variate prior, which imply symmetric  $\mathbf{W}_0^{\mathbf{H}}$ , are of polynomial form in  $\mathbf{A}$ .

**Example S1** (Posterior Correspondence Covariance Class)

Tractable choices of the prior parameters in the  $\mathbf{A}$  view, which satisfy posterior correspondence and the commutation relations are for example

$$\mathbf{A}_0 = c_0 \mathbf{I} \quad \text{and} \quad \mathbf{W}_0^{\mathbf{A}} = \sum_{i=1}^{n-1} c_i \mathbf{A}^i,$$

where  $\mathbf{H}_0 = \mathbf{A}_0^{-1}$  with  $c_i \in \mathbb{R}$ . Motivated by  $\text{tr}(\mathbf{A}) \stackrel{!}{=} \text{tr}(\mathbf{A}_0)$  an initial choice could be  $c_0 = n^{-1} \text{tr}(\mathbf{A})$ .

Finally, note that in practice we do not actually require  $\mathbf{W}_0^{\mathbf{A}}$ . We only ever need access to  $\mathbf{W}_0^{\mathbf{A}} \mathbf{S}$ .

### S6.2.2 Symmetric Matrix-variate Normal Prior

We now turn to the symmetric model, which we assumed throughout the paper, given in Corollary S3. We prove Theorem 3, the main result of this section demonstrating *weak posterior correspondence* for the symmetric Kronecker covariance, by employing the matrix inversion lemma for the posterior mean  $\mathbf{A}_k$ . We begin by establishing a set of technical lemmata first, which mainly expand terms appearing during matrix block inversion.

**Lemma S3** (Symmetric Posterior Inverse)

Under the assumptions of Corollary S3, the inverse of the posterior mean is given by

$$\mathbf{A}_k^{-1} = \mathbf{A}_0^{-1} - \mathbf{A}_0^{-1} [\mathbf{U}_{\mathbf{A}} \quad \mathbf{V}_{\mathbf{A}}] \begin{bmatrix} \mathbf{U}_{\mathbf{A}}^{\top} \mathbf{A}_0^{-1} \mathbf{U}_{\mathbf{A}} & \mathbf{I} + \mathbf{U}_{\mathbf{A}}^{\top} \mathbf{A}_0^{-1} \mathbf{V}_{\mathbf{A}} \\ \mathbf{I} + \mathbf{V}_{\mathbf{A}}^{\top} \mathbf{A}_0^{-1} \mathbf{U}_{\mathbf{A}} & \mathbf{V}_{\mathbf{A}}^{\top} \mathbf{A}_0^{-1} \mathbf{V}_{\mathbf{A}} \end{bmatrix}^{-1} \begin{bmatrix} \mathbf{U}_{\mathbf{A}}^{\top} \\ \mathbf{V}_{\mathbf{A}} \end{bmatrix} \mathbf{A}_0^{-1}$$

where

$$\begin{aligned} \mathbf{U}_{\mathbf{A}} &:= \mathbf{W}_0^{\mathbf{A}} \mathbf{S} (\mathbf{S}^{\top} \mathbf{W}_0^{\mathbf{A}} \mathbf{S})^{-1} \in \mathbb{R}^{n \times k}, \\ \mathbf{V}_{\mathbf{A}} &:= (\mathbf{I} - \frac{1}{2} \mathbf{U}_{\mathbf{A}} \mathbf{S}^{\top}) (\mathbf{Y} - \mathbf{A}_0 \mathbf{S}) = (\mathbf{I} - \frac{1}{2} \mathbf{U}_{\mathbf{A}} \mathbf{S}^{\top}) \mathbf{\Delta}_{\mathbf{A}} \in \mathbb{R}^{n \times k}. \end{aligned}$$



*Proof.* We rewrite the rank-2 update in Section 2.1 as follows

$$A_k = A_0 + U_{\mathbf{A}}V_{\mathbf{A}}^T + V_{\mathbf{A}}U_{\mathbf{A}}^T = A_0 + [U_{\mathbf{A}} \quad V_{\mathbf{A}}] \begin{bmatrix} \mathbf{0} & I \\ I & \mathbf{0} \end{bmatrix} \begin{bmatrix} U_{\mathbf{A}}^T \\ V_{\mathbf{A}}^T \end{bmatrix}.$$

Then the statement follows directly from the matrix inversion lemma.  $\square$

Next, we expand the terms inside the blocks of the matrix to be inverted in Lemma S3. This leads to the following lemma.

**Lemma S4**

Given the assumptions of Corollary S3, let  $W_0^{\mathbf{A}}$  and  $A_0$  be symmetric and assume (2) and (1) hold. Define

$$\begin{aligned} \Lambda &= S^T W_0^{\mathbf{A}} S \\ \Pi &= S^T W_0^{\mathbf{A}} A_0^{-1} \Delta_{\mathbf{A}}, \end{aligned}$$

then  $\Lambda \in \mathbb{R}^{m \times m}$  and  $\Lambda + \Pi \in \mathbb{R}^{m \times m}$  are symmetric and invertible and we obtain

$$\Lambda + \Pi = S^T W_0^{\mathbf{A}} A_0^{-1} A S = S^T A A_0^{-1} A S = S^T A W_0^{\mathbf{H}} A S \quad (\text{S45})$$

$$\Pi = \Delta_{\mathbf{A}}^T A_0^{-1} A S \quad (\text{S46})$$

$$U_{\mathbf{A}}^T A_0^{-1} \Delta_{\mathbf{A}} = \Lambda^{-1} \Pi \quad (\text{S47})$$

$$\Delta_{\mathbf{A}}^T S = S^T \Delta_{\mathbf{A}} \quad (\text{S48})$$

$$U_{\mathbf{A}} = A S \Lambda^{-1} \quad (\text{S49})$$

$$U_{\mathbf{A}}^T A_0^{-1} U_{\mathbf{A}} = \Lambda^{-1} (\Lambda + \Pi) \Lambda^{-1} \quad (\text{S50})$$

$$I + U_{\mathbf{A}}^T A_0^{-1} V_{\mathbf{A}} = \Lambda^{-1} (\Lambda + \Pi) \left( I - \frac{1}{2} \Lambda^{-1} S^T \Delta_{\mathbf{A}} \right) \quad (\text{S51})$$

$$I + V_{\mathbf{A}}^T A_0^{-1} U_{\mathbf{A}} = \left( I - \frac{1}{2} \Delta_{\mathbf{A}}^T S \Lambda^{-1} \right) (\Lambda + \Pi) \Lambda^{-1} \quad (\text{S52})$$

$$V_{\mathbf{A}}^T A_0^{-1} V_{\mathbf{A}} = \Pi - \frac{1}{2} \left( (\Lambda + \Pi) \Lambda^{-1} S^T \Delta_{\mathbf{A}} + \Delta_{\mathbf{A}}^T S \Lambda^{-1} (\Lambda + \Pi) \right) \quad (\text{S53})$$

$$+ \frac{1}{4} \Delta_{\mathbf{A}}^T S \Lambda^{-1} (\Lambda + \Pi) \Lambda^{-1} S^T \Delta_{\mathbf{A}} \quad (\text{S54})$$

*Proof.* We begin by proving that  $\Lambda$  and  $\Lambda + \Pi$  are symmetric and invertible. We have by Sylvester's rank inequality that  $\Lambda$  is invertible. For symmetric  $W_0^{\mathbf{A}}$ ,  $\Lambda$  is symmetric by definition. We have that

$$\begin{aligned} \Lambda + \Pi &= S^T W_0^{\mathbf{A}} S + S^T W_0^{\mathbf{A}} A_0^{-1} (A S - A_0 S) = S^T W_0^{\mathbf{A}} A_0^{-1} A S = S^T A A_0^{-1} A S \\ &= S^T W_0^{\mathbf{A}} A_0^{-1} A S = S^T A W_0^{\mathbf{H}} A S \end{aligned}$$

Thus, by Sylvester's rank inequality  $\Lambda + \Pi$  is invertible. Given symmetric  $A_0$ , it is symmetric. Further, it holds that

$$\begin{aligned}
\Pi &= \Lambda + \Pi - \Lambda = S^\top A A_0^{-1} A S - S^\top A S = \Delta_A^\top A_0^{-1} A S \\
U_A^\top A_0^{-1} \Delta_A &= (S^\top W_0^A S)^{-1} S^\top W_0^A A_0^{-1} \Delta_A = \Lambda^{-1} \Pi \\
\Delta_A^\top S &= (A S - A_0 S)^\top S = S^\top A S - S^\top A_0 S \\
U_A &= W_0^A S (S^\top W_0^A S)^{-1} = A S \Lambda^{-1} \\
U_A^\top A_0^{-1} U_A &= \Lambda^{-1} S^\top A A_0^{-1} A S \Lambda^{-1} = \Lambda^{-1} (\Lambda + \Pi) \Lambda^{-1} \\
I + U_A^\top A_0^{-1} V_A &= I + \Lambda^{-1} S^\top A A_0^{-1} (I - \frac{1}{2} U_A S^\top) \Delta_A = I + \Lambda^{-1} S^\top A A_0^{-1} (I - \frac{1}{2} A S \Lambda^{-1} S^\top) \Delta_A \\
&= I + \Lambda^{-1} S^\top A A_0^{-1} (A S - A_0 S) - \frac{1}{2} \Lambda^{-1} S^\top A A_0^{-1} A S \Lambda^{-1} S^\top \Delta_A \\
&= \Lambda^{-1} (\Lambda + \Pi) - \frac{1}{2} \Lambda^{-1} (\Lambda + \Pi) \Lambda^{-1} S^\top \Delta_A = \Lambda^{-1} (\Lambda + \Pi) (I - \frac{1}{2} \Lambda^{-1} S^\top \Delta_A) \\
I + V_A^\top A_0^{-1} U_A &= (I + U_A^\top A_0^{-1} V_A)^\top = (\Lambda^{-1} (\Lambda + \Pi) (I - \frac{1}{2} \Lambda^{-1} S^\top \Delta_A))^\top \\
&= (I - \frac{1}{2} \Delta_A^\top S \Lambda^{-1}) (\Lambda + \Pi) \Lambda^{-1},
\end{aligned}$$

where we used that  $\Lambda$  and  $\Lambda + \Pi$  are symmetric. Finally, we have that

$$\begin{aligned}
V_A^\top A_0^{-1} V_A &= \Delta_A^\top (I - \frac{1}{2} S U_A^\top) A_0^{-1} (I - \frac{1}{2} U_A S^\top) \Delta_A \\
&= \Delta_A^\top (I - \frac{1}{2} S \Lambda^{-1} S^\top A) A_0^{-1} (I - \frac{1}{2} A S \Lambda^{-1} S^\top) \Delta_A \\
&= \Delta_A^\top A_0^{-1} (I - \frac{1}{2} A S \Lambda^{-1} S^\top) \Delta_A - \frac{1}{2} \Delta_A^\top S \Lambda^{-1} S^\top A A_0^{-1} (I - \frac{1}{2} A S \Lambda^{-1} S^\top) \Delta_A \\
&= (S^\top A A_0^{-1} - S^\top) (I - \frac{1}{2} A S \Lambda^{-1} S^\top) \Delta_A - \frac{1}{2} \Delta_A^\top S \Lambda^{-1} S^\top A A_0^{-1} \Delta_A \\
&\quad + \frac{1}{4} \Delta_A^\top S \Lambda^{-1} S^\top A A_0^{-1} A S \Lambda^{-1} S^\top \Delta_A \\
&= S^\top A A_0^{-1} \Delta_A - S^\top \Delta_A - \frac{1}{2} S^\top A A_0^{-1} A S \Lambda^{-1} S^\top \Delta_A + \frac{1}{2} S^\top A S \Lambda^{-1} S^\top \Delta_A \\
&\quad - \frac{1}{2} \Delta_A^\top S \Lambda^{-1} S^\top A A_0^{-1} \Delta_A + \frac{1}{4} \Delta_A^\top S \Lambda^{-1} S^\top A A_0^{-1} A S \Lambda^{-1} S^\top \Delta_A \\
&= S^\top A A_0^{-1} A S - S^\top A S - S^\top \Delta_A - \frac{1}{2} (\Lambda + \Pi) \Lambda^{-1} S^\top \Delta_A + \frac{1}{2} S^\top \Delta_A \\
&\quad - \frac{1}{2} \Delta_A^\top S \Lambda^{-1} S^\top A A_0^{-1} \Delta_A + \frac{1}{4} \Delta_A^\top S \Lambda^{-1} S^\top A A_0^{-1} A S \Lambda^{-1} S^\top \Delta_A \\
&= \Pi - \frac{1}{2} S^\top \Delta_A - \frac{1}{2} (\Lambda + \Pi) \Lambda^{-1} S^\top \Delta_A - \frac{1}{2} \Delta_A^\top S \Lambda^{-1} S^\top A A_0^{-1} (A S - A_0 S) \\
&\quad + \frac{1}{4} \Delta_A^\top S \Lambda^{-1} S^\top A A_0^{-1} A S \Lambda^{-1} S^\top \Delta_A \\
&= \Pi - \frac{1}{2} S^\top \Delta_A - \frac{1}{2} (\Lambda + \Pi) \Lambda^{-1} S^\top \Delta_A - \frac{1}{2} \Delta_A^\top S \Lambda^{-1} (\Lambda + \Pi) + \frac{1}{2} \Delta_A^\top S \Lambda^{-1} \Lambda \\
&\quad + \frac{1}{4} \Delta_A^\top S \Lambda^{-1} S^\top A A_0^{-1} A S \Lambda^{-1} S^\top \Delta_A \\
&= \Pi - \frac{1}{2} ((\Lambda + \Pi) \Lambda^{-1} S^\top \Delta_A + \Delta_A^\top S \Lambda^{-1} (\Lambda + \Pi)) + \frac{1}{4} \Delta_A^\top S \Lambda^{-1} (\Lambda + \Pi) \Lambda^{-1} S^\top \Delta_A,
\end{aligned}$$

where we dropped some of the terms temporarily for clarity of exposition.  $\square$

We will now use these intermediate results to perform block inversion on the  $2k \times 2k$  matrix to be inverted in Lemma S3.

**Lemma S5**

Given the assumptions of Corollary S3, additionally assume (1) and (2) hold. Let

$$\mathbf{T} = \begin{bmatrix} \mathbf{T}_{11} & \mathbf{T}_{12} \\ \mathbf{T}_{21} & \mathbf{T}_{22} \end{bmatrix} = \begin{bmatrix} \mathbf{U}_{\mathbf{A}}^{\top} \mathbf{A}_0^{-1} \mathbf{U}_{\mathbf{A}} & \mathbf{I} + \mathbf{U}_{\mathbf{A}}^{\top} \mathbf{A}_0^{-1} \mathbf{V}_{\mathbf{A}} \\ \mathbf{I} + \mathbf{V}_{\mathbf{A}}^{\top} \mathbf{A}_0^{-1} \mathbf{U}_{\mathbf{A}} & \mathbf{V}_{\mathbf{A}}^{\top} \mathbf{A}_0^{-1} \mathbf{V}_{\mathbf{A}} \end{bmatrix}^{-1},$$

then the block matrices  $\mathbf{T}_{ij} \in \mathbb{R}^{m \times m}$  are given by

$$\mathbf{T}_{11} = \mathbf{\Lambda}(\mathbf{\Lambda} + \mathbf{\Pi})^{-1} \mathbf{\Lambda} - \left(\mathbf{I} - \frac{1}{2} \mathbf{S}^{\top} \mathbf{\Delta}_{\mathbf{A}} \mathbf{\Lambda}^{-1}\right) \left(\mathbf{I} - \frac{1}{2} \mathbf{\Lambda}^{-1} \mathbf{\Delta}_{\mathbf{A}}^{\top} \mathbf{S}\right)$$

$$\mathbf{T}_{12} = \left(\mathbf{I} - \frac{1}{2} \mathbf{S}^{\top} \mathbf{\Delta}_{\mathbf{A}} \mathbf{\Lambda}^{-1}\right)$$

$$\mathbf{T}_{21} = \mathbf{T}_{12}^{\top} = \left(\mathbf{I} - \frac{1}{2} \mathbf{\Lambda}^{-1} \mathbf{\Delta}_{\mathbf{A}}^{\top} \mathbf{S}\right)$$

$$\mathbf{T}_{22} = -\mathbf{\Lambda}^{-1}.$$

*Proof.* Let

$$\mathbf{K} = \mathbf{T}^{-1} = \begin{bmatrix} \mathbf{U}_{\mathbf{A}}^{\top} \mathbf{A}_0^{-1} \mathbf{U}_{\mathbf{A}} & \mathbf{I} + \mathbf{U}_{\mathbf{A}}^{\top} \mathbf{A}_0^{-1} \mathbf{V}_{\mathbf{A}} \\ \mathbf{I} + \mathbf{V}_{\mathbf{A}}^{\top} \mathbf{A}_0^{-1} \mathbf{U}_{\mathbf{A}} & \mathbf{V}_{\mathbf{A}}^{\top} \mathbf{A}_0^{-1} \mathbf{V}_{\mathbf{A}} \end{bmatrix},$$

then the inverse of the Schur complement  $\mathbf{D} = \mathbf{K} / (\mathbf{U}_{\mathbf{A}}^{\top} \mathbf{A}_0^{-1} \mathbf{U}_{\mathbf{A}})$  is given by

$$\begin{aligned} \mathbf{D}^{-1} &= (\mathbf{K}_{22} - \mathbf{K}_{21} \mathbf{K}_{11}^{-1} \mathbf{K}_{12})^{-1} \\ &= (\mathbf{V}_{\mathbf{A}}^{\top} \mathbf{A}_0^{-1} \mathbf{V}_{\mathbf{A}} - (\mathbf{I} + \mathbf{V}_{\mathbf{A}}^{\top} \mathbf{A}_0^{-1} \mathbf{U}_{\mathbf{A}}) (\mathbf{U}_{\mathbf{A}}^{\top} \mathbf{A}_0^{-1} \mathbf{U}_{\mathbf{A}})^{-1} (\mathbf{I} + \mathbf{U}_{\mathbf{A}}^{\top} \mathbf{A}_0^{-1} \mathbf{V}_{\mathbf{A}}))^{-1} \\ &= (\mathbf{V}_{\mathbf{A}}^{\top} \mathbf{A}_0^{-1} \mathbf{V}_{\mathbf{A}} - (\mathbf{I} - \frac{1}{2} \mathbf{\Delta}_{\mathbf{A}}^{\top} \mathbf{S} \mathbf{\Lambda}^{-1}) (\mathbf{\Lambda} + \mathbf{\Pi}) (\mathbf{I} - \frac{1}{2} \mathbf{\Lambda}^{-1} \mathbf{S}^{\top} \mathbf{\Delta}_{\mathbf{A}}))^{-1} \\ &= (\mathbf{V}_{\mathbf{A}}^{\top} \mathbf{A}_0^{-1} \mathbf{V}_{\mathbf{A}} - (\mathbf{\Lambda} - \frac{1}{2} \mathbf{\Delta}_{\mathbf{A}}^{\top} \mathbf{S}) \mathbf{\Lambda}^{-1} (\mathbf{\Lambda} + \mathbf{\Pi}) \mathbf{\Lambda}^{-1} (\mathbf{\Lambda} - \frac{1}{2} \mathbf{S}^{\top} \mathbf{\Delta}_{\mathbf{A}}))^{-1} \\ &= (\mathbf{V}_{\mathbf{A}}^{\top} \mathbf{A}_0^{-1} \mathbf{V}_{\mathbf{A}} - (\mathbf{\Lambda} + \mathbf{\Pi}) + \frac{1}{2} (\mathbf{\Delta}_{\mathbf{A}}^{\top} \mathbf{S} \mathbf{\Lambda}^{-1} (\mathbf{\Lambda} + \mathbf{\Pi}) + (\mathbf{\Lambda} + \mathbf{\Pi}) \mathbf{\Lambda}^{-1} \mathbf{S}^{\top} \mathbf{\Delta}_{\mathbf{A}}) \\ &\quad - \frac{1}{4} \mathbf{\Delta}_{\mathbf{A}}^{\top} \mathbf{S} \mathbf{\Lambda}^{-1} (\mathbf{\Lambda} + \mathbf{\Pi}) \mathbf{\Lambda}^{-1} \mathbf{S}^{\top} \mathbf{\Delta}_{\mathbf{A}})^{-1} \\ &= (\mathbf{\Pi} - \mathbf{\Lambda} - \mathbf{\Pi})^{-1} \\ &= -\mathbf{\Lambda}^{-1}, \end{aligned}$$

where we used Lemma S4. By block matrix inversion and again with Lemma S4 we obtain

$$\begin{aligned} \mathbf{T}_{11} &= (\mathbf{U}_{\mathbf{A}}^{\top} \mathbf{A}_0^{-1} \mathbf{U}_{\mathbf{A}})^{-1} + (\mathbf{U}_{\mathbf{A}}^{\top} \mathbf{A}_0^{-1} \mathbf{U}_{\mathbf{A}})^{-1} (\mathbf{I} + \mathbf{U}_{\mathbf{A}}^{\top} \mathbf{A}_0^{-1} \mathbf{V}_{\mathbf{A}}) \mathbf{D}^{-1} (\mathbf{I} + \mathbf{V}_{\mathbf{A}}^{\top} \mathbf{A}_0^{-1} \mathbf{U}_{\mathbf{A}}) (\mathbf{U}_{\mathbf{A}}^{\top} \mathbf{A}_0^{-1} \mathbf{U}_{\mathbf{A}})^{-1} \\ &= \mathbf{\Lambda}(\mathbf{\Lambda} + \mathbf{\Pi})^{-1} \mathbf{\Lambda} + \mathbf{\Lambda} \left(\mathbf{I} - \frac{1}{2} \mathbf{\Lambda}^{-1} \mathbf{S}^{\top} \mathbf{\Delta}_{\mathbf{A}}\right) \mathbf{D}^{-1} \left(\mathbf{I} - \frac{1}{2} \mathbf{\Delta}_{\mathbf{A}}^{\top} \mathbf{S} \mathbf{\Lambda}^{-1}\right) \mathbf{\Lambda} \\ &= \mathbf{\Lambda}(\mathbf{\Lambda} + \mathbf{\Pi})^{-1} \mathbf{\Lambda} + (\mathbf{\Lambda} - \frac{1}{2} \mathbf{S}^{\top} \mathbf{\Delta}_{\mathbf{A}}) \mathbf{D}^{-1} (\mathbf{\Lambda} - \frac{1}{2} \mathbf{\Delta}_{\mathbf{A}}^{\top} \mathbf{S}) \end{aligned}$$

as well as

$$\begin{aligned} \mathbf{T}_{12} &= -(\mathbf{U}_{\mathbf{A}}^{\top} \mathbf{A}_0^{-1} \mathbf{U}_{\mathbf{A}})^{-1} (\mathbf{I} + \mathbf{U}_{\mathbf{A}}^{\top} \mathbf{A}_0^{-1} \mathbf{V}_{\mathbf{A}}) \mathbf{D}^{-1} \\ &= -\mathbf{\Lambda}(\mathbf{\Lambda} + \mathbf{\Pi})^{-1} \mathbf{\Lambda} \mathbf{\Lambda}^{-1} (\mathbf{\Lambda} + \mathbf{\Pi}) \left(\mathbf{I} - \frac{1}{2} \mathbf{\Lambda}^{-1} \mathbf{S}^{\top} \mathbf{\Delta}_{\mathbf{A}}\right) \mathbf{D}^{-1} \\ &= -(\mathbf{\Lambda} - \frac{1}{2} \mathbf{S}^{\top} \mathbf{\Delta}_{\mathbf{A}}) \mathbf{D}^{-1} \end{aligned}$$

$$\mathbf{T}_{21} = \mathbf{T}_{12}^{\top} = -\mathbf{D}^{-\top} (\mathbf{\Lambda} - \frac{1}{2} \mathbf{\Delta}_{\mathbf{A}}^{\top} \mathbf{S})$$

and finally  $\mathbf{T}_{22} = \mathbf{D}^{-1} = -\mathbf{\Lambda}^{-1}$ . □

**Lemma S6**

Given the assumptions of Corollary S3, additionally assume (1) and (2) hold. Let

$$\mathbf{F} = \mathbf{A}_0^{-1} [\mathbf{U}_{\mathbf{A}} \quad \mathbf{V}_{\mathbf{A}}] \begin{bmatrix} \mathbf{T}_{11} & \mathbf{T}_{12} \\ \mathbf{T}_{21} & \mathbf{T}_{22} \end{bmatrix} \begin{bmatrix} \mathbf{U}_{\mathbf{A}}^{\top} \\ \mathbf{V}_{\mathbf{A}}^{\top} \end{bmatrix} \mathbf{A}_0^{-1},$$

where  $T$  is chosen as in Lemma S5, then if  $S^\top AS = I$ , we have

$$F = A_0^{-1}AS(I + \Pi)^{-1}S^\top AA_0^{-1} - SS^\top.$$

*Proof.* By expanding the quadratic and using Lemma S5, we obtain the terms

$$\begin{aligned} F_{11} &:= A_0^{-1}U_A T_{11} U_A^\top A_0^{-1} \\ &= A_0^{-1}U_A \Lambda (\Lambda + \Pi)^{-1} \Lambda U_A^\top A_0^{-1} - A_0^{-1}U_A (I - \frac{1}{2}S^\top \Delta_A \Lambda^{-1}) (I - \frac{1}{2}\Lambda^{-1} \Delta_A^\top S) U_A^\top A_0^{-1} \\ &= A_0^{-1}AS(\Lambda + \Pi)^{-1}S^\top AA_0^{-1} - A_0^{-1}AS\Lambda^{-1}(I - \frac{1}{2}S^\top \Delta_A \Lambda^{-1})(I - \frac{1}{2}\Lambda^{-1} \Delta_A^\top S)\Lambda^{-1}S^\top AA_0^{-1} \\ &= A_0^{-1}AS(\Lambda + \Pi)^{-1}S^\top AA_0^{-1} - A_0^{-1}AS\Lambda^{-2}S^\top AA_0^{-1} \\ &\quad + \frac{1}{2}A_0^{-1}AS\Lambda^{-1}(S^\top \Delta_A \Lambda^{-1} + \Lambda^{-1} \Delta_A^\top S)\Lambda^{-1}S^\top AA_0^{-1} \\ &\quad - \frac{1}{4}A_0^{-1}AS\Lambda^{-1}S^\top \Delta_A \Lambda^{-2} \Delta_A^\top S \Lambda^{-1}S^\top AA_0^{-1} \end{aligned}$$

$$\begin{aligned} F_{12} &:= A_0^{-1}U_A T_{12} V_A^\top A_0^{-1} \\ &= A_0^{-1}U_A (I - \frac{1}{2}S^\top \Delta_A \Lambda^{-1}) V_A^\top A_0^{-1} \\ &= A_0^{-1}AS\Lambda^{-1}(I - \frac{1}{2}S^\top \Delta_A \Lambda^{-1}) \Delta_A^\top (I - \frac{1}{2}S U_A^\top) A_0^{-1} \\ &= A_0^{-1}AS\Lambda^{-1}(I - \frac{1}{2}S^\top \Delta_A \Lambda^{-1}) \Delta_A^\top (I - \frac{1}{2}S \Lambda^{-1} S^\top A) A_0^{-1} \\ &= A_0^{-1}AS\Lambda^{-1} \Delta_A^\top A_0^{-1} - \frac{1}{2}A_0^{-1}AS\Lambda^{-1}(S^\top \Delta_A \Lambda^{-1} \Delta_A^\top + \Delta_A^\top S \Lambda^{-1} S^\top A) A_0^{-1} \\ &\quad + \frac{1}{4}A_0^{-1}AS\Lambda^{-1}S^\top \Delta_A \Lambda^{-1} \Delta_A^\top S \Lambda^{-1}S^\top AA_0^{-1} \end{aligned}$$

$$\begin{aligned} F_{21} &:= F_{12}^\top = A_0^{-1}(I - \frac{1}{2}AS\Lambda^{-1}S^\top) \Delta_A (I - \frac{1}{2}\Lambda^{-1} \Delta_A^\top S) \Lambda^{-1}S A A_0^{-1} \\ &= A_0^{-1} \Delta_A \Lambda^{-1} S^\top A A_0^{-1} - \frac{1}{2}A_0^{-1}(\Delta_A \Lambda^{-1} \Delta_A^\top S + AS\Lambda^{-1}S^\top \Delta_A) \Lambda^{-1}S^\top AA_0^{-1} \\ &\quad + \frac{1}{4}A_0^{-1}AS\Lambda^{-1}S^\top \Delta_A \Lambda^{-1} \Delta_A^\top S \Lambda^{-1}S^\top AA_0^{-1} \end{aligned}$$

$$\begin{aligned} F_{22} &:= A_0^{-1}V_A T_{22} V_A^\top A_0^{-1} \\ &= -A_0^{-1}(I - \frac{1}{2}U_A S^\top) \Delta_A \Lambda^{-1} \Delta_A^\top (I - \frac{1}{2}S U_A^\top) A_0^{-1} \\ &= -A_0^{-1}(I - \frac{1}{2}AS\Lambda^{-1}S^\top) \Delta_A \Lambda^{-1} \Delta_A^\top (I - \frac{1}{2}S \Lambda^{-1} S^\top A) A_0^{-1} \\ &= -A_0^{-1} \Delta_A \Lambda^{-1} \Delta_A^\top A_0^{-1} + \frac{1}{2}A_0^{-1}(AS\Lambda^{-1}S^\top \Delta_A \Lambda^{-1} \Delta_A^\top + \Delta_A \Lambda^{-1} \Delta_A^\top S \Lambda^{-1} S^\top A) A_0^{-1} \\ &\quad - \frac{1}{4}A_0^{-1}AS\Lambda^{-1}S^\top \Delta_A \Lambda^{-1} \Delta_A^\top S \Lambda^{-1}S^\top AA_0^{-1} \end{aligned}$$

Assuming  $S^\top AS = I$ , it holds that

$$\begin{aligned} F_{11} &= A_0^{-1}AS(I + \Pi)^{-1}S^\top AA_0^{-1} - A_0^{-1}ASS^\top AA_0^{-1} + \frac{1}{2}A_0^{-1}AS(S^\top \Delta_A + \Delta_A^\top S)S^\top AA_0^{-1} \\ &\quad - \frac{1}{4}A_0^{-1}ASS^\top \Delta_A \Delta_A^\top SS^\top AA_0^{-1} \\ F_{12} &= A_0^{-1}ASS^\top AA_0^{-1} - A_0^{-1}ASS^\top - \frac{1}{2}A_0^{-1}AS(S^\top \Delta_A \Delta_A^\top + \Delta_A^\top SS^\top A)A_0^{-1} \\ &\quad + \frac{1}{4}A_0^{-1}ASS^\top \Delta_A \Delta_A^\top SS^\top AA_0^{-1} \\ F_{21} &= A_0^{-1}ASS^\top AA_0^{-1} - SS^\top AA_0^{-1} - \frac{1}{2}A_0^{-1}(\Delta_A \Delta_A^\top S + ASS^\top \Delta_A)S^\top AA_0^{-1} \end{aligned}$$

$$\begin{aligned}
& + \frac{1}{4} \mathbf{A}_0^{-1} \mathbf{A} \mathbf{S} \mathbf{S}^\top \Delta_{\mathbf{A}} \Delta_{\mathbf{A}}^\top \mathbf{S} \mathbf{S}^\top \mathbf{A} \mathbf{A}_0^{-1} \\
\mathbf{F}_{22} & = \mathbf{A}_0^{-1} \Delta_{\mathbf{A}} \mathbf{S}^\top - \mathbf{A}_0^{-1} \Delta_{\mathbf{A}} \mathbf{S}^\top \mathbf{A} \mathbf{A}_0^{-1} + \frac{1}{2} (\mathbf{A} \mathbf{S} \mathbf{S}^\top \Delta_{\mathbf{A}} \Delta_{\mathbf{A}}^\top + \Delta_{\mathbf{A}} \Delta_{\mathbf{A}}^\top \mathbf{S} \mathbf{S}^\top \mathbf{A}) \mathbf{A}_0^{-1} \\
& - \frac{1}{4} \mathbf{A}_0^{-1} \mathbf{A} \mathbf{S} \mathbf{S}^\top \Delta_{\mathbf{A}} \Delta_{\mathbf{A}}^\top \mathbf{S} \mathbf{S}^\top \mathbf{A} \mathbf{A}_0^{-1},
\end{aligned}$$

which leads to

$$\begin{aligned}
\mathbf{F}_{11} + \mathbf{F}_{12} & = \mathbf{A}_0^{-1} \mathbf{A} \mathbf{S} (\mathbf{I} + \mathbf{\Pi})^{-1} \mathbf{S}^\top \mathbf{A} \mathbf{A}_0^{-1} - \mathbf{A}_0^{-1} \mathbf{A} \mathbf{S} \mathbf{S}^\top + \frac{1}{2} \mathbf{A}_0^{-1} \mathbf{A} \mathbf{S} (\mathbf{S}^\top \Delta_{\mathbf{A}} \mathbf{S}^\top \mathbf{A} - \mathbf{S}^\top \Delta_{\mathbf{A}} \Delta_{\mathbf{A}}^\top) \mathbf{A}_0^{-1} \\
\mathbf{F}_{21} + \mathbf{F}_{22} & = \mathbf{A}_0^{-1} \Delta_{\mathbf{A}} \mathbf{S}^\top + \frac{1}{2} \mathbf{A}_0^{-1} \mathbf{A} \mathbf{S} (\mathbf{S}^\top \Delta_{\mathbf{A}} \Delta_{\mathbf{A}}^\top - \mathbf{S}^\top \Delta_{\mathbf{A}} \mathbf{S}^\top \mathbf{A}) \mathbf{A}_0^{-1} \\
& = \mathbf{A}_0^{-1} \mathbf{A} \mathbf{S} \mathbf{S}^\top - \mathbf{S} \mathbf{S}^\top + \frac{1}{2} \mathbf{A}_0^{-1} \mathbf{A} \mathbf{S} (\mathbf{S}^\top \Delta_{\mathbf{A}} \Delta_{\mathbf{A}}^\top - \mathbf{S}^\top \Delta_{\mathbf{A}} \mathbf{S}^\top \mathbf{A}) \mathbf{A}_0^{-1}.
\end{aligned}$$

Finally, adding up the individual terms we obtain

$$\mathbf{F} = \mathbf{F}_{11} + \mathbf{F}_{12} + \mathbf{F}_{21} + \mathbf{F}_{22} = \mathbf{A}_0^{-1} \mathbf{A} \mathbf{S} (\mathbf{I} + \mathbf{\Pi})^{-1} \mathbf{S}^\top \mathbf{A} \mathbf{A}_0^{-1} - \mathbf{S} \mathbf{S}^\top.$$

□

### Theorem 2 (Weak Posterior Correspondence)

Let  $\mathbf{W}_0^{\mathbf{H}} \in \mathbb{R}_{\text{sym}}^{n \times n}$  be positive definite. Assume  $\mathbf{H}_0 = \mathbf{A}_0^{-1}$ , and that  $\mathbf{W}_0^{\mathbf{A}}, \mathbf{A}_0, \mathbf{W}_0^{\mathbf{H}}$  satisfy (1) and (2), then weak posterior correspondence holds for the symmetric Kronecker covariance.

*Proof.* First note that without loss of generality  $\mathbf{S}^\top \mathbf{A} \mathbf{S} = \mathbf{I}$ , i.e. only the direction of the action matters in Algorithm 1 not its magnitude. This can be seen from the forms of  $\mathbf{A}_k$  and  $\mathbf{H}_k$  in Section 2.1. Any positive factor  $\alpha > 0$  of  $\mathbf{s}_k$  cancels in the update expressions. Expanding the right hand side we have using (S36), that  $\mathbf{H}_k \mathbf{Y} = \mathbf{S}$ . Then by Lemma S3, Lemma S6 and  $\mathbf{S}^\top \mathbf{A} \mathbf{S} = \mathbf{I}$ , the left hand side evaluates to

$$\begin{aligned}
\mathbf{A}_k^{-1} \mathbf{Y} & = (\mathbf{A}_0^{-1} - \mathbf{F}) \mathbf{Y} \\
& = (\mathbf{A}_0^{-1} - \mathbf{A}_0^{-1} \mathbf{A} \mathbf{S} (\mathbf{I} + \mathbf{\Pi})^{-1} \mathbf{S}^\top \mathbf{A} \mathbf{A}_0^{-1} + \mathbf{S} \mathbf{S}^\top) \mathbf{A} \mathbf{S} \\
& = \mathbf{A}_0^{-1} \mathbf{A} \mathbf{S} - \mathbf{A}_0^{-1} \mathbf{A} \mathbf{S} + \mathbf{S} \\
& = \mathbf{S} \\
& = \mathbf{H}_k \mathbf{Y}.
\end{aligned}$$

This concludes the proof. □

This theorem shows that for a certain choice of symmetric matrix-variate normal prior the estimated inverse of the matrix  $\mathbf{H}_k$  corresponds to the inverse of the estimated matrix  $\mathbf{A}_k^{-1}$ . It also shows that both act like  $\mathbf{A}^{-1}$  on the space spanned by  $\mathbf{Y}$ , consistent with the interpretation of the two being the best guess for the inverse  $\mathbf{A}^{-1}$ .

## S7 Galerkin's Method for PDEs

In the spirit of applying machine learning in the sciences [35], we briefly outlined an application of Algorithm 1 to the solution of partial differential equations in Section 4. As an example we considered the Dirichlet problem for the Poisson equation given by

$$\begin{cases} -\Delta u(x, y) = f(x, y) & (x, y) \in \text{int } \Omega \\ u(x, y) = u_{\partial\Omega}(x, y) & (x, y) \in \partial\Omega \end{cases} \quad (\text{S55})$$

where  $\Omega$  is a connected open region with sufficiently regular boundary and  $u_{\partial\Omega} : \partial\Omega \rightarrow \mathbb{R}$  defines the boundary conditions. The corresponding weak solution of (S55) is given by  $u \in V$  such that for all test functions  $v \in V$

$$a(u, v) := \int_{\Omega} \nabla u \cdot \nabla v \, dx = \int_{\Omega} f v \, dx =: f(v), \quad (\text{S56})$$

where  $a(\cdot, \cdot)$  is a bilinear form. Next, one derives the *Galerkin equation* by choosing a finite-dimensional subspace  $V_\square \subset V$  and corresponding basis  $e_1^\square, \dots, e_n^\square$ . Then (S56) reduces to finding  $u \in V_\square$  such that for all  $i \in \{1, \dots, n\}$  it holds that  $a(u, e_i^\square) = \sum_{j=1}^n u_j a(e_j^\square, e_i^\square) = f(e_i^\square)$  which is a linear system  $\mathbf{A}u = \mathbf{f}$  with the entries of the Gram matrix given by  $\mathbf{A}_{ij} = a(e_j^\square, e_i^\square)$  and  $f_i = f(e_i^\square)$ .

### S7.1 Operator View

The operator view provides another motivation for placing a distribution over the matrix  $\mathbf{A}$  of a linear system. When approximating the solution to a PDE, as we do here, then solution-based inference for linear systems [21, 13] can be viewed as placing a Gaussian process prior over the solution  $u : \Omega \rightarrow \mathbb{R}$  [49]. The matrix-based approach [11] instead can be interpreted as placing a Gaussian measure [50] on the infinite-dimensional space of the differential operator instead. This induces a Gaussian distribution on the Gram matrix  $\mathbf{A}$  modelling the uncertainty about the actions of the (discretized) differential operator.

**Definition S4** (Infinite-dimensional Gaussian Measures [50])

Let  $W$  be a topological vector space with Borel probability measure  $\mu$ , then  $\mu$  is Gaussian, iff for each continuous linear functional  $f \in W^*$ , the pushforward  $\mu \circ f^{-1}$  is a Gaussian measure on  $\mathbb{R}$ , i.e.  $f$  is a Gaussian random variable on  $(W, \mathcal{B}_W, \mu)$ .

This definition and further detail on Gaussian measures in infinite-dimensional spaces can be found in the book by Bogachev [50]. We now model the differential operator as a random variable on the space of bounded linear operators and show that this induces a distribution on the Gram matrix arising from discretization via Galerkin's method.

**Theorem S6** (Gaussian Measures on the Space of Bounded Linear Operators)

Let  $V$  be a Hilbert space and let  $W = B(V, V)$  be the space of bounded linear operators from  $V$  to  $V$  with Borel probability measure  $\mu$  and let  $\mathbf{A}$  be a Gaussian random variable on  $(W, \mathcal{B}_W, \mu)$ . Consider the operator equation

$$\mathbf{A}u = \mathbf{f}$$

and let  $a : V \times V \rightarrow \mathbb{R}$ ,  $(u, v) \mapsto \langle \mathbf{A}u, v \rangle_V = \langle \mathbf{f}, v \rangle_V$  be its corresponding bilinear form. Let  $V_\square$  be an  $n$ -dimensional subspace of  $V$ , then the resulting Gram matrix  $\mathbf{A} \in \mathbb{R}^{n \times n}$  is matrix-variate Gaussian.

*Proof.* Since  $V$  is Banach, so is  $W$ . Define the functional  $a_W : W \rightarrow \mathbb{R}$  given by  $a_W(\mathbf{A}, u, v) = a(u, v)$  for fixed  $u, v \in V$ . The map  $a_W(\cdot, u, v)$  is linear by linearity of the inner product and bounded since using the Cauchy-Schwarz inequality, it holds that

$$|a_W(\mathbf{A}, u, v)| = |\langle \mathbf{A}u, v \rangle_V| \leq \|\mathbf{A}u\|_V \|v\|_V \leq \|\mathbf{A}\|_W \|u\|_V \|v\|_V = C \|\mathbf{A}\|_W.$$

Therefore  $a_W(\cdot, u, v) \in W^*$  for all  $u, v \in V$ . By Definition S4 of a Gaussian measure the push forward  $\mu \circ a_W^{-1}$  is a Gaussian measure on  $\mathbb{R}$  for all  $u, v \in V$ , in particular also for a basis  $\{v_i\}_{i=1}^n$  of  $V_\square$ . Therefore the Gram matrix  $\mathbf{A}$  given by  $\mathbf{A}_{ij} = a(v_i, v_j) = a_W(\mathbf{A}, v_i, v_j)$  is matrix-variate Gaussian since its components are Gaussian.  $\square$

**Remark S1**

The Laplacian  $\Delta : H^2(\Omega) \rightarrow L^2(\Omega)$  is a bounded linear operator on the Sobolev space  $H^2(\Omega)$ . Note, that in general differential operators are in fact *not bounded*. Hence, the simple argument above does not generalize to arbitrary differential operators.

**Remark S2**

If the bilinear form  $a$  in addition to being continuous is also weakly coercive, then by the Lax-Milgram theorem the operator equation has a unique solution. A symmetric and weakly coercive operator implies a symmetric positive-definite Gram matrix.

### S7.2 Discretization Refinement

The linear system  $\mathbf{A}u = \mathbf{f}$  arises from discretizing (S55) using Galerkin's method on a given mesh  $\square$  defined via a finite-dimensional subspace  $V_\square \subset V$  such that  $u \in V_\square$ . By solving this problem using a probabilistic linear solver we obtain a posterior distribution over the inverse  $\mathbf{H}$  of the discretized

differential operator  $\mathbf{A}$ . Our goal is to leverage the obtained information about the solution on the coarse mesh to extrapolate to a refined discretization, similar in spirit to multi-grid methods [51]. This approach can be seen as an instance of transfer learning and could be used for adaptive probabilistic mesh refinement strategies based on the uncertainty about the solution in a certain region of the mesh.

Consider a fine mesh  $\mathbb{H}$  given by  $V_{\mathbb{H}}$ , where  $n_{\mathbb{H}} = \dim(V_{\mathbb{H}}) > \dim(V_{\square}) = n_{\square}$  such that  $V_{\square} \subset V_{\mathbb{H}} \subset V$ . We would like to transfer information from solving the problem on the coarse mesh  $V_{\square}$  to the solution of the discretized PDE on the fine mesh  $V_{\mathbb{H}}$ . To do so we compute the predictive distribution on the fine mesh, given the belief over the inverse differential operator on the coarse mesh, i.e.

$$p(\mathbf{H}_{\mathbb{H}}) = \int p(\mathbf{H}_{\mathbb{H}} | \mathbf{H}_{\square}) p(\mathbf{H}_{\square}) d\mathbf{H}_{\square}.$$

Define the *prolongation operator*  $\mathbf{P} : \mathbb{R}^{n_{\square}} \rightarrow \mathbb{R}^{n_{\mathbb{H}}}$  given by  $P_{ij} = \langle \mathbf{e}_i^{\mathbb{H}}, \mathbf{e}_j^{\square} \rangle$  satisfying  $\mathbf{P}^{\top} \mathbf{P} = \mathbf{I} \in \mathbb{R}^{n_{\square} \times n_{\square}}$ , implying it is injective. The distribution over the inverse operator on the fine mesh given the inverse operator on the coarse mesh is given by

$$p(\mathbf{H}_{\mathbb{H}} | \mathbf{H}_{\square}) = \mathcal{N}(\mathbf{H}_{\mathbb{H}}; \mathbf{P} \mathbf{H}_{\square} \mathbf{P}^{\top}, \mathbf{\Lambda}) \quad (\text{S57})$$

where  $\mathbf{\Lambda} \in \mathbb{R}_{\text{sym}}^{n_{\mathbb{H}} \times n_{\mathbb{H}}}$  positive definite models the numerical uncertainty induced by the coarser discretization. This corresponds to the assumption that solving the problem on a coarser grid approximates the solution on a fine grid projected to the coarse grid.

Now assume we have a posterior distribution over the inverse differential operator on the coarse grid from a solve of the coarse problem using Algorithm 1, given by

$$p(\mathbf{H}_{\square}) = \mathcal{N}(\mathbf{H}_{\square}; \mathbf{H}_{\square}^k, \mathbf{W}_{\square}^k \otimes \mathbf{W}_{\square}^k).$$

The projection in (S57) is a linear map, since by the characteristic property of the Kronecker product (S4) we have

$$\text{svec}(\mathbf{P} \mathbf{H}_{\square} \mathbf{P}^{\top}) = \mathbf{Q}(\mathbf{P} \otimes \mathbf{P}) \mathbf{Q}^{\top} \text{svec}(\mathbf{H}_{\square}).$$

Therefore by Theorem S4 the predictive distribution is also closed-form and Gaussian.

**Proposition S7** (Predictive Distribution on Fine Mesh)

Let  $p(\mathbf{H}_{\square}) = \mathcal{N}(\mathbf{H}_{\square}; \mathbf{H}_{\square}^k, \mathbf{W}_{\square}^k \otimes \mathbf{W}_{\square}^k)$  be a prior on  $\mathbf{H}_{\square}$  and assume a likelihood of the form (S57). Then the predictive distribution is given by  $p(\mathbf{H}_{\mathbb{H}}) = \mathcal{N}(\mathbf{H}_{\mathbb{H}}; \mathbf{H}_{\mathbb{H}}^0, \mathbf{\Sigma}_{\mathbb{H}}^0)$ , where

$$\begin{aligned} \mathbf{H}_{\mathbb{H}}^0 &= \mathbf{P} \mathbf{H}_{\square}^k \mathbf{P}^{\top}, \\ \mathbf{\Sigma}_{\mathbb{H}}^0 &= \mathbf{P} \mathbf{W}_{\square}^k \mathbf{P}^{\top} \otimes \mathbf{P} \mathbf{W}_{\square}^k \mathbf{P}^{\top} + \mathbf{\Lambda}. \end{aligned}$$

*Proof.* By Theorem S4 we obtain for the mean and covariance of the predictive distribution

$$\begin{aligned} \mathbf{H}_{\mathbb{H}}^0 &= \mathbf{P} \mathbf{H}_{\square}^k \mathbf{P}^{\top} \\ \mathbf{\Sigma}_{\mathbb{H}}^0 &= \mathbf{Q}(\mathbf{P} \otimes \mathbf{P}) \mathbf{Q}^{\top} (\mathbf{W}_{\square}^k \otimes \mathbf{W}_{\square}^k) \mathbf{Q}(\mathbf{P}^{\top} \otimes \mathbf{P}^{\top}) \mathbf{Q}^{\top} + \mathbf{\Lambda} \\ &= \frac{1}{2} \mathbf{Q}(\mathbf{P} \mathbf{W}_{\square}^k \mathbf{P}^{\top} \otimes \mathbf{P} \mathbf{W}_{\square}^k \mathbf{P}^{\top} + \mathbf{P} \mathbf{W}_{\square}^k \mathbf{P}^{\top} \boxtimes \mathbf{P} \mathbf{W}_{\square}^k \mathbf{P}^{\top}) \mathbf{Q}^{\top} + \mathbf{\Lambda} \\ &= \mathbf{P} \mathbf{W}_{\square}^k \mathbf{P}^{\top} \otimes \mathbf{P} \mathbf{W}_{\square}^k \mathbf{P}^{\top} + \mathbf{\Lambda} \end{aligned}$$

where we used (S34) and the symmetry of  $\mathbf{W}_{\square}^k$ .  $\square$

For general  $\mathbf{\Lambda}$  the covariance of the predictive distribution does not have symmetric Kronecker form, making its use as a prior for a new solve on the fine mesh challenging. We aim to exploit structural assumptions on  $\mathbf{\Lambda}$  and results on nearest Kronecker products to a sum of Kronecker products to remedy this shortcoming in the future.