



Inferring the basal sliding coefficient field for the Stokes ice sheet model under rheological uncertainty

Olaekan Babaniyi¹, Ruanui Nicholson², Umberto Villa³, and Noémi Petra⁴

¹School of Mathematical Sciences, Rochester Institute of Technology, Rochester, NY 14623, USA

²Department of Engineering Science, University of Auckland, Auckland 1010, New Zealand

³Electrical & Systems Engineering, Washington University in St. Louis, St. Louis, MO 63130, USA

⁴Department of Applied Mathematics, University of California, Merced, Merced, CA 95343, USA

Correspondence: Noémi Petra (npetra@ucmerced.edu)

Received: 11 August 2020 – Discussion started: 31 August 2020

Revised: 9 February 2021 – Accepted: 9 February 2021 – Published: 9 April 2021

Abstract. We consider the problem of inferring the basal sliding coefficient field for an uncertain Stokes ice sheet forward model from synthetic surface velocity measurements. The uncertainty in the forward model stems from unknown (or uncertain) auxiliary parameters (e.g., rheology parameters). This inverse problem is posed within the Bayesian framework, which provides a systematic means of quantifying uncertainty in the solution. To account for the associated model uncertainty (error), we employ the Bayesian approximation error (BAE) approach to approximately pre-marginalize simultaneously over both the noise in measurements and uncertainty in the forward model. We also carry out approximative posterior uncertainty quantification based on a linearization of the parameter-to-observable map centered at the maximum a posteriori (MAP) basal sliding coefficient estimate, i.e., by taking the Laplace approximation. The MAP estimate is found by minimizing the negative log posterior using an inexact Newton conjugate gradient method. The gradient and Hessian actions to vectors are efficiently computed using adjoints. Sampling from the approximate covariance is made tractable by invoking a low-rank approximation of the data misfit component of the Hessian. We study the performance of the BAE approach in the context of three numerical examples in two and three dimensions. For each example, the basal sliding coefficient field is the parameter of primary interest which we seek to infer, and the rheology parameters (e.g., the flow rate factor or the Glen's flow law exponent coefficient field) represent so-called nuisance (secondary uncertain) parameters. Our results indicate that accounting for model uncertainty stemming from the pres-

ence of nuisance parameters is crucial. Namely our findings suggest that using nominal values for these parameters, as is often done in practice, without taking into account the resulting modeling error, can lead to overconfident and heavily biased results. We also show that the BAE approach can be used to account for the additional model uncertainty at no additional cost at the online stage.

1 Introduction

Inferring the basal sliding coefficient field using both the linear and nonlinear Stokes ice sheet model from noisy surface velocity measurements has received considerable attention in recent years (see, for example, Truffer, 2004, Raymond and Gudmundsson, 2009, Pollard and DeConto, 2012, Isaac et al., 2015b, Morlighem et al., 2013, Zhao et al., 2018a, b, Giudici et al., 2014, Petra et al., 2012, 2014, and Isaac et al., 2015a). The standard approach to this problem invariably assumes that the other parameters of the ice, such as those controlling the rheology, are known precisely. This is particularly common, for example, in the case of the so-called flow rate factor and the Glen's flow law exponent, in which nominal values such as $A = 10^{-16} \text{ Pa}^{-n} \text{ a}^{-1}$ and $n = 3$, respectively, are prescribed; we refer, e.g., to Isaac et al. (2015a), Petra et al. (2014), Raymond and Gudmundsson (2009), Truffer (2004), Morlighem et al. (2013), Zhu et al. (2016), Zhao et al. (2018b), Giudici et al. (2014), and Pollard and DeConto (2012). The inference problem is made significantly more challenging (both theoretically and numerically) by al-

lowing the rheology parameters to be uncertain and spatially varying. One possible approach to solve the problem is to infer both the basal sliding coefficient and the rheology parameters. However, this considerably increases both the ill-posedness of the inverse problem and the associated computational costs. For most ice sheet inverse problems considered in the literature, the field of interest is the basal sliding parameter, which arguably presents the largest uncertainty in determining the ice flow rate.

It is well documented that, in practice, the rheology parameters of ice sheets are not known exactly (e.g., Bons et al., 2018; Marshall, 2005; Gillet-Chaulet et al., 2011, 2012; Cuffey and Paterson, 2010; Brondex et al., 2019; Raymond and Gudmundsson, 2011). Compounding this issue is the fact that measured ice velocities can be heavily influenced by rheology parameters (Schlegel et al., 2015; Bulthuis et al., 2019). This fact was demonstrated in Petra et al. (2012), in which the authors used the same Stokes ice sheet model as in the current paper to reconstruct reasonable estimates of the Glen's flow law exponent from noisy surface velocity measurements, suggesting that the surface measurements are indeed sensitive to changes in the Glen's flow law exponent field. Despite these findings, it is standard in the literature to assume that rheology – among other – parameters of the ice are known a priori (see, for instance, Bons et al., 2018, Marshall, 2005, Gillet-Chaulet et al., 2011, 2012, Cuffey and Paterson, 2010, Brondex et al., 2019, and Van der Veen, 2013).

In this paper, we treat the rheology parameters (specifically the Glen's flow law exponent and the flow rate factor fields) as auxiliary (nuisance) parameters, i.e., parameters which are not of primary interest. However, fixing these auxiliary parameters at incorrect, though possibly well-justified, values often induces so-called modeling errors. It is well understood, though, that the solutions to inverse problems are generally sensitive to modeling errors which – if not properly accounted for – can lead to inaccurate, nonphysical, and in some cases meaningless solutions of the inverse problem (Brynjarsdóttir and O'Hagan, 2014; Giudici et al., 2014; Kaipio and Somersalo, 2007, 2005). From a statistical viewpoint, fixing auxiliary parameters to nominal values suggests that these parameters are known exactly and hence neglects all associated uncertainties. This in turn often results in biased and overconfident estimates for the parameters of interest (see, for example, Kaipio and Somersalo, 2007, Kaipio and Kolehmainen, 2013, and Nicholson et al., 2018 and the references therein).

We carry out estimation of the basal sliding coefficient within the Bayesian framework (Kaipio and Somersalo, 2005; Stuart, 2010), which is particularly well suited to incorporating various sources and types of uncertainties, including those resulting from model errors (Tarantola, 2005; Kaipio and Somersalo, 2005, 2007). Moreover, to ensure the work here is readily transferable to inference problems in large-scale ice flow problems, such as those discussed in Isaac et al. (2015a), we make use of the computational

framework proposed in Bui-Thanh et al. (2013) and Petra et al. (2014) for handling infinite-dimensional Bayesian inverse problems (Stuart, 2010). This approach, combined with adjoint-based means to compute the derivative information needed by the optimization solver, ensures mesh independence and computational efficiency.

To account for the uncertainty in the rheology parameters, we utilize the Bayesian approximation error (BAE) approach (Kaipio and Somersalo, 2005, 2007; Kaipio and Kolehmainen, 2013), which, broadly speaking, lumps all modeling and measurement uncertainties into a single additive *total error* term. The total error can then be approximately marginalized over in a similar manner to how standard additive errors are dealt with (Kaipio and Kolehmainen, 2013). The BAE approach is particularly attractive computationally as

- (a) the approximate marginalization can be carried out prior to data acquisition, i.e., *pre-marginalization*, and
- (b) the equations to be solved in the adjoint-state approach maintain the same general form (Nicholson et al., 2018).

The BAE approach has been used in a variety of settings (see, for example, Kaipio and Kolehmainen, 2013, Arridge et al., 2006, Castello and Kaipio, 2019, and Lamien et al., 2019, among others, and the references therein). A particularly relevant, and recent, example is the application of the approach to the so-called Robin inverse problem encountered, for instance, in corrosion detection (Nicholson et al., 2018). There the parameter of interest is also a Robin-type boundary condition on an inaccessible part of the domain, while the nuisance parameter is the (electrical or thermal) conductivity of the domain.

To study the performance of the BAE approach, we formulate and solve three ice sheet flow model problems involving synthetic data. Our results suggest that simply setting rheology parameters to nominal values can result in severely misleading estimates of the basal sliding coefficient field and associated posterior uncertainty if the additional uncertainty in the rheology parameters is not accounted for. In comparison, we show that incorporating the additional modeling uncertainties using the BAE approach leads to sensible estimates of the basal sliding coefficient and reasonable posterior uncertainty at no additional online cost. We place particular emphasis on the feasibility of the posterior uncertainty estimates, in particular, on how well the true parameter is contained within the posterior distribution.

1.1 Contributions

In previous work, we addressed the problem of inferring the basal sliding coefficient field from surface velocity measurements in the context of ice sheet flow in a deterministic, moderate scale, synthetic observational data setting in Petra et al. (2012), in a Bayesian inference and infinite-dimensional setting in Petra et al. (2014), and more recently in a large-scale,

real data setting in Isaac et al. (2015a). Here the goal is to extend this inversion framework to account for additional uncertainties in the ice sheet model. The main contributions of this paper are as follows. Firstly, we show that setting rheology parameters to values commonly found for ice sheet models in the literature can lead to erroneous posterior estimates of the basal sliding coefficient if the underlying uncertainty in the rheology parameters is not accounted for. Secondly, we show that this situation can be remedied by employing the BAE approach to premarginalize over rheology uncertainties. Thirdly, we show that this approach requires no additional computational resources or time at the *online stage* as all computations required for premarginalization are carried out prior to the acquisition of data.

1.2 Organization of paper

The paper is organized as follows. In Sect. 2, we outline the forward nonlinear Stokes flow equations for ice sheet problems, while in Sect. 3 we briefly review the Bayesian framework for inverse problems, the computation of the maximum a posteriori estimate, and the approximate posterior covariance. In Sect. 4, we show how to apply the BAE approach to premarginalize over auxiliary parameters. In Sects. 5 and 6, we formulate and solve three ice sheet inverse problems and study the performance of the proposed method. Finally, Sect. 7 provides concluding remarks.

2 Forward ice sheet flow model

In this section, we describe the forward ice sheet flow problem that is used for the inference of the basal sliding coefficient field under uncertain rheology. As in Petra et al. (2012, 2014) and Isaac et al. (2015a), we model the flow of ice as an isothermal, viscous, shear-thinning, incompressible fluid via the balance of mass and linear momentum (Hutter, 1983; Marshall, 2005; Paterson, 1994), namely

$$\nabla \cdot \mathbf{u} = 0 \text{ in } \Omega, \tag{1a}$$

$$-\nabla \cdot \boldsymbol{\sigma}_u = \rho \mathbf{g} \text{ in } \Omega, \tag{1b}$$

where \mathbf{u} denotes the velocity field, $\boldsymbol{\sigma}_u$ the stress tensor, ρ the density of the ice, and \mathbf{g} gravity. The stress, $\boldsymbol{\sigma}_u$, can be decomposed as $\boldsymbol{\sigma}_u = \boldsymbol{\tau}_u - \mathbf{I}p$, where $\boldsymbol{\tau}_u$ is the deviatoric stress tensor, p the pressure, and \mathbf{I} the identity tensor. The domain considered in this paper is $\Omega = [0, L]^{d-1} \times [0, H]$, for $d = 2$ or $d = 3$. We employ the Glen’s flow law (Glen, 1955) which relates the stress and strain rate tensors by

$$\boldsymbol{\tau}_u = 2\eta(\mathbf{u})\dot{\boldsymbol{\epsilon}}_u \text{ with } \eta(\mathbf{u}) = \frac{1}{2}A^{-\frac{1}{n}}\dot{\boldsymbol{\epsilon}}_{II}^{\frac{1-n}{2n}}, \tag{1c}$$

where η is the effective viscosity, A is the flow rate factor, and $\dot{\boldsymbol{\epsilon}}_u = \frac{1}{2}(\nabla \mathbf{u} + \nabla \mathbf{u}^T)$ and $\dot{\boldsymbol{\epsilon}}_{II} = \frac{1}{2}\text{tr}(\dot{\boldsymbol{\epsilon}}_u^2)$ are the strain rate tensor and its second invariant. Above, $n = n(\mathbf{x})$ is the spatially varying Glen’s flow law exponent, which satisfies $n(\mathbf{x}) \geq 1$




Figure 1. Schematic of a two-dimensional slab of ice (used in examples 1 and 2). The schematic can also be thought of as a cross section through the three-dimensional slab of ice used for Example 3. The blue circles show representative (random) measurement locations but do not necessarily coincide with the actual measurement locations used in the examples. θ is the slope of the ice slab.

for all $\mathbf{x} \in \Omega$ to ensure the ice is a shear-thinning fluid (Glen, 1955).

In line with Petra et al. (2012), the top boundary Γ_t is equipped with a traction-free boundary condition, all lateral boundaries Γ_p are equipped with periodic boundary conditions, and on the basal surface Γ_b we apply a no flow condition for the normal component of \mathbf{u} along with a linear sliding law for the tangential components. That is, the boundary conditions are given by

$$\boldsymbol{\sigma}_u \mathbf{n} = \mathbf{0} \text{ on } \Gamma_t, \tag{1d}$$

$$\mathbf{u}|_{\Gamma_l} = \mathbf{u}|_{\Gamma_r} \text{ and } \boldsymbol{\sigma}_u \mathbf{n}|_{\Gamma_l} = \boldsymbol{\sigma}_u \mathbf{n}|_{\Gamma_r} \text{ on } \Gamma_p, \tag{1e}$$

$$\mathbf{u} \cdot \mathbf{n} = 0 \text{ on } \Gamma_b, \tag{1f}$$

$$\mathbf{T}\boldsymbol{\sigma}_u \mathbf{n} + \exp(\beta)\mathbf{T}\mathbf{u} = \mathbf{0} \text{ on } \Gamma_b, \tag{1g}$$

where $\beta(\mathbf{x})$ is the log basal sliding coefficient field¹, \mathbf{n} is the outward normal unit vector, and $\mathbf{T} := \mathbf{I} - \mathbf{n}\mathbf{n}^T$ is the projection onto the tangential plane. Above, we generically used Γ_l and Γ_r to denote pairs of opposing boundaries on Γ_p on which periodic conditions are imposed. We note that β generally represents a combination of complex phenomena (see, for example, Schoof, 2005, 2010, and Perego et al., 2014). Furthermore, the methods and results discussed in the current paper do not rely on the particular top and lateral boundary conditions specified. As such, alternative boundary conditions could also be imposed on Γ_t and Γ_p . For a simple illustration of the problem setup (shown in two dimensions), see Fig. 1.

- *The weak form of the Stokes equation.* In what follows, let us introduce the weak form of Eq. (1), as it is the starting point for both the finite element discretization of

¹The “exp” function is used to ensure the basal sliding coefficient remains positive. For simplicity, in what follows, we will refer to β as the basal sliding coefficient.

the forward model and the computation of the gradient and action of the Hessian required for the solution of the inverse problem using the adjoint-state method (see, for example, Isaac et al., 2015a). Multiplying the nonlinear Stokes system (1) with arbitrary test functions $\tilde{\mathbf{u}}$ and \tilde{p} and using integration by parts over Ω (Gockenbach, 2006; Elman et al., 2005), the weak form of Eq. (1) is given by the following. Find $(\mathbf{u}, p) \in \mathcal{W} = \mathcal{V} \times \mathcal{Q}$ such that

$$\int_{\Omega} 2\eta(\mathbf{u})\dot{\boldsymbol{\epsilon}}_{\mathbf{u}} : \dot{\boldsymbol{\epsilon}}_{\tilde{\mathbf{u}}} \, dx - \int_{\Omega} (p\nabla \cdot \tilde{\mathbf{u}} + \tilde{p}\nabla \cdot \mathbf{u}) \, dx + \int_{\Gamma_b} \exp(\beta)\mathbf{T}\mathbf{u} \cdot \mathbf{T}\tilde{\mathbf{u}} \, ds = \int_{\Omega} \rho\mathbf{g} \cdot \tilde{\mathbf{u}} \, dx \tag{2}$$

for all $(\tilde{\mathbf{u}}, \tilde{p}) \in \mathcal{W}$. In line with Elman et al. (2005) and Isaac et al. (2015b), we set $\mathcal{V} := \{\mathbf{u} \in (H^1(\Omega))^d : \mathbf{u}|_{\Gamma_1} = \mathbf{u}|_{\Gamma_r}, \mathbf{u} \cdot \mathbf{n}|_{\Gamma_b} = 0\}$ and $\mathcal{Q} := (L^2(\Omega))^d$, for $d = 2$ or $d = 3$.

- *Discretization.* To guarantee the *inf-sup* stability (well-posedness) of the discretized forward problem, we discretize the velocity and pressure using Taylor–Hood finite elements, i.e., quadratic elements for each velocity component and linear elements for pressure (see, for example, Elman et al., 2005). The basal sliding coefficient field is discretized using continuous linear Lagrange basis functions $\{\phi_j(s)\}_{j=1}^m$, i.e., $\beta_h(s) = \sum_{j=1}^m \beta_j \phi_j(s)$, where $s \in \Gamma_b$. In what follows, we denote by $\boldsymbol{\beta} = (\beta_1, \beta_2, \dots, \beta_m) \in \mathbb{R}^m$ the discrete basal sliding coefficient field.

3 Inferring the basal sliding coefficient field

In this section, we summarize the Bayesian inference framework, which will be used in combination with the Bayesian approximation error approach, to account for uncertainties in rheology parameters. To allow for the systematic incorporation of uncertainties, we consider the inverse problem in the Bayesian framework (Tarantola, 2005; Kaipio and Somersalo, 2005). In this framework, the solution of the underlying statistical inverse problem is given by the posterior probability density. For nonlinear inverse problems with expensive forward models and high-dimensional parameters (as is the case for ice sheet inverse problems), fully characterizing the posterior is typically not tractable. Consequently, we compute the Laplace approximation of the posterior, which requires only the maximum a posteriori (MAP) estimate, i.e., the basal sliding coefficient which maximizes the posterior density and the approximate posterior covariance.

We use Bayes’ theorem to write the solution of the Bayesian inverse problem as the posterior measure, which describes the probability law of the parameter conditioned

on measurements (Tarantola, 2005; Stuart, 2010). The Formulation of the posterior relies on both the prior density and the likelihood function, which we outline below. We note that initially we pose the problem in an infinite-dimensional setting, which is particularly well suited to large-scale problems (e.g., Bui-Thanh et al., 2012; Isaac et al., 2015a) as it ensures the discretization invariance and well-posedness of the Bayesian inverse problem (Stuart, 2010).

3.1 The prior

We postulate a Gaussian prior density on the (spatially varying) basal sliding coefficient, i.e., $\beta \sim \mathcal{N}(\beta_*, \mathcal{C}_\beta)$, with covariance operator \mathcal{C}_β , and mean value $\beta_* \in \mathcal{E}$, where \mathcal{E} is defined as the range of $\mathcal{C}_\beta^{\frac{1}{2}}$ (see, for example, Stuart, 2010 and Bui-Thanh et al., 2013 for more details). To ensure the inverse problem is well posed in infinite dimensions, we use a squared inverse elliptic operator to define the prior covariance operator² (e.g., Flath et al., 2011; Bui-Thanh et al., 2013; Petra et al., 2014).

More specifically, we take $\mathcal{C}_\beta = \mathcal{A}^{-2}$, where \mathcal{A} is the second order elliptic differential operator defined by

$$\mathcal{A}\beta := -\nabla \cdot (\gamma_\beta \nabla \beta) + \delta_\beta \beta \text{ on } \Gamma_b, \tag{3}$$

where the strictly positive parameters γ_β (m^2) and δ_β (adi-dimensional) control the correlation length and the marginal variance. Specifically, the correlation length (defined as the distance for which the two points have a correlation coefficient of 0.1) is proportional to $\sqrt{\gamma_\beta/\delta_\beta}$ (m), while the variance is proportional to $\delta_\beta^{-2}(\gamma_\beta/\delta_\beta)^{\frac{d-1}{2}}$ (see, for example, Kristenko et al., 2019 and the references therein). This choice of prior covariance operator is particularly well suited to large-scale problems, as discretization of \mathcal{A} (using a finite element discretization) is sparse (see, for example, Lindgren et al., 2011 and Osborn et al., 2017). As discussed in Kristenko et al. (2019), Daon and Stadler (2018), and Roininen et al. (2014), suitable boundary conditions need to be stipulated to reduce boundary artifacts. In this work, we choose to equip \mathcal{A} with periodic boundary conditions on $\partial\Gamma_b$, which parallels the periodic boundary condition (1e) of the forward model. We note that the discrete representation of the prior covariance operator, denoted $\boldsymbol{\Gamma}_{\text{pr}}$, is defined as follows (e.g., Bui-Thanh et al., 2013; Petra et al., 2014; Villa et al., 2021):

$$\left[\boldsymbol{\Gamma}_{\text{pr}}^{-1} \right]_{ij} = \int_{\Gamma_b} \phi_i(s) \mathcal{A}^2 \phi_j(s) ds \quad i, j \in \{1, 2, \dots, m\}. \tag{4}$$

²For cases in which β has only one spatial dimension, an inverse elliptic operator, i.e., without the squaring, also results in a valid covariance operator (see Petra et al., 2014). However, in the current paper we consider cases in which β is one-dimensional and two-dimensional, and thus for ease of exposition, and in the interest of space, we limit the choice of the prior covariance operator to the squared inverse elliptic operator.

Therefore, the discrete parameter β follows a Gaussian distribution $\mathcal{N}(\beta_*, \Gamma_{pr})$, with prior mean $\beta_* \in \mathbb{R}^m$ and covariance Γ_{pr} . That is, the prior density of β is given by

$$\pi_{\text{prior}}(\beta) \propto \exp \left\{ -\frac{1}{2} \|\beta - \beta_*\|_{\Gamma_{pr}^{-1}}^2 \right\}, \quad (5)$$

where $\|\cdot\|_{\Gamma_{pr}^{-1}}$ denotes the Γ_{pr}^{-1} weighted l_2 inner product.

3.2 The data likelihood

We assume the velocity measurements, denoted d , are corrupted by additive noise and are related to the basal sliding coefficient through

$$d = \mathcal{F}(\beta) + e, \quad (6)$$

where $\mathcal{F} : L^2(\Omega) \rightarrow \mathbb{R}^q$ is called the parameter-to-observable map, and $e \in \mathbb{R}^q$ denotes the noise in the measurements.

As is somewhat common in the literature (e.g., Raymond and Gudmundsson, 2009; Petra et al., 2012, 2014), we take the data to consist of (noisy) point-wise observations of each component of the velocity field on the top surface³. In discrete settings, we compute $\mathcal{F}(\beta)$ by first solving the Stokes equation (Eq. 1) and then applying a linear observation operator that extracts the velocity at the measurement locations. We assume the noise, e , is independent of the basal sliding coefficient, has zero mean, and is Gaussian, i.e., $e \sim \mathcal{N}(\mathbf{0}, \Gamma_e)$. The likelihood is then of the following form (e.g., Tarantola, 2005; Kaipio and Somersalo, 2005):

$$\pi_{\text{like}}(d|\beta) \propto \exp \left\{ -\frac{1}{2} \|\mathcal{F}(\beta) - d\|_{\Gamma_e^{-1}}^2 \right\}. \quad (7)$$

3.3 The posterior

By applying Bayes' theorem, the posterior density of β is proportional to the product of the prior density (Eq. 5) and the data likelihood (Eq. 7). This is given by

$$\pi_{\text{post}}(\beta|d) \propto \exp \left\{ -\frac{1}{2} \|\mathcal{F}(\beta) - d\|_{\Gamma_e^{-1}}^2 - \frac{1}{2} \|\beta - \beta_*\|_{\Gamma_{pr}^{-1}}^2 \right\}. \quad (8)$$

The corresponding MAP estimate is then defined as

$$\beta_{\text{MAP}} := \underset{\beta \in \mathbb{R}^m}{\operatorname{argmin}} \frac{1}{2} \|\mathcal{F}(\beta) - d\|_{\Gamma_e^{-1}}^2 + \frac{1}{2} \|\beta - \beta_*\|_{\Gamma_{pr}^{-1}}^2. \quad (9)$$

We note that the problem of finding the MAP estimate, defined in Eq. (9), reduces to a deterministic inverse problem. To solve this problem we use an inexact Newton conjugate gradient (CG) method, as in Petra et al. (2012). To derive

³Vertical velocity measurements may not always be available; however, as shown in Raymond and Gudmundsson (2009), these measurements are fairly insignificant. Furthermore, the assumed noise level in the current paper is larger than the vertical velocities.

the required first (i.e., gradient) and second (i.e., Hessian) derivative information needed by Newton's method, we use an adjoint-based method and refer the reader to Petra et al. (2012) for the derivation and expressions of the required derivatives.

3.4 Quantifying posterior uncertainty

To (approximately) quantify the resulting uncertainty in the inferred basal sliding parameter, we invoke a local Gaussian approximation of the posterior (i.e., the Laplace approximation). That is, the solution to the Bayesian inverse problem is now given by a Gaussian distribution with mean β_{MAP} and covariance Γ_{po} given by the inverse of the (Gauss–Newton) Hessian of the negative log-posterior, evaluated at the MAP estimate. More specifically, we make the approximation, $\beta|d \sim \mathcal{N}(\beta_{\text{MAP}}, \Gamma_{po})$, with β_{MAP} given by Eq. (9), and

$$\Gamma_{po} = \mathbf{H}(\beta_{\text{MAP}})^{-1} = (\overline{\mathbf{H}}(\beta_{\text{MAP}}) + \Gamma_{pr}^{-1})^{-1} = (\mathbf{F}^T(\beta_{\text{MAP}})\Gamma_e^{-1}\mathbf{F}(\beta_{\text{MAP}}) + \Gamma_{pr}^{-1})^{-1}, \quad (10)$$

where $\overline{\mathbf{H}}(\beta)$ is the Gauss–Newton Hessian of the data misfit term (i.e., the negative log-likelihood), and \mathbf{F} is the Jacobian matrix of the parameter-to-observable map, \mathcal{F} (e.g., Bui-Thanh et al., 2013).

The construction of the posterior covariance matrix (i.e., the inverse of the Hessian) is prohibitive for large-scale problems since its dimension is equal to the dimension of the parameter. To make operations with the posterior covariance matrix tractable, we exploit the fact that the eigenvalues of $\overline{\mathbf{H}}(\beta_{\text{MAP}})$ collapse to zero rapidly since the data contain limited information about the (infinite-dimensional) parameter field. Thus a low-rank approximation of the data misfit component of the Hessian $\overline{\mathbf{H}}$ can be constructed as in Isaac et al. (2015a) by solving the generalized eigenvalue problem

$$\overline{\mathbf{H}}\mathbf{V}_r = \Gamma_{pr}^{-1}\mathbf{V}_r\mathbf{\Lambda}_r, \quad (11)$$

where $\mathbf{\Lambda}_r = \operatorname{diag}(\lambda_1, \lambda_2, \dots, \lambda_r) \in \mathbb{R}^{r \times r}$ is a diagonal matrix collecting the r largest generalized eigenvalues, λ_i , and $\mathbf{V}_r = [\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_r] \in \mathbb{R}^{m \times r}$ is the matrix collecting the corresponding Γ_{pr}^{-1} -orthonormal eigenvectors, \mathbf{v}_i . Above, the truncation index r is chosen such that the remaining eigenvalues, λ_i , for $i = r + 1, \dots, m$, are sufficiently smaller than 1 – often chosen such that $\lambda_i \leq c$, for some $0.01 \leq c \leq 1$ (e.g., Isaac et al., 2015a; Flath et al., 2011).

Substituting $\overline{\mathbf{H}} \approx \Gamma_{pr}^{-1}\mathbf{V}_r\mathbf{\Lambda}_r\mathbf{V}_r^T\Gamma_{pr}^{-1}$ into Eq. (10) and using the Sherman–Morrison–Woodbury identity (Golub and Van Loan, 1996) and after a few algebraic manipulations (e.g., Isaac et al., 2015a), we obtain the following low-rank-based approximation of the posterior covariance (under the Laplace approximation):

$$\Gamma_{po} \approx \Gamma_{pr} - \mathbf{V}_r\mathbf{d}_r\mathbf{V}_r^T, \quad (12)$$

where $\mathbf{d}_r = \operatorname{diag}(\lambda_1/(\lambda_1 + 1), \lambda_2/(\lambda_2 + 1), \dots, \lambda_r/(\lambda_r + 1)) \in \mathbb{R}^{r \times r}$.

4 Premarginalization over auxiliary parameters and the Bayesian approximation error approach

The Bayesian approximation error (BAE) approach (Kaipio and Somersalo, 2007, 2005; Kaipio and Kolehmainen, 2013) can be used to approximately premarginalize over auxiliary parameters. The BAE approach essentially combines all uncertainties, including those generated by fixing uncertain parameters, into a single additive total error term. The total error term can then be premarginalized over, i.e., marginalized over before the acquisition of data. We next outline the process.

We denote by a the auxiliary parameters which in the current study are defined over the entire computational domain, Ω , and are assumed to be Gaussian distributed with covariance operator $C_a = \mathcal{L}^{-2}$, where \mathcal{L} is defined by

$$\mathcal{L}a := -\nabla \cdot (\gamma_a \nabla a) + \delta_a a, \text{ in } \Omega, \tag{13}$$

and mean value a_* . In line with the forward problem, \mathcal{L} is equipped with periodic boundary conditions on the lateral boundaries of Ω , while on the top and bottom boundaries we enforce Robin boundary conditions. Note that explicit knowledge of the distribution of a is not needed; we only require the ability to sample realizations of a . In what follows, we denote by \mathbf{a} any (possibly more than one) discretized auxiliary (uncertain) parameter, such as the rheology parameters.

Next, we let $\tilde{\mathcal{F}}(\boldsymbol{\beta}, \mathbf{a})$ denote an accurate parameter-to-observable mapping so that the relationship between the parameters and the measured data is given by

$$\mathbf{d} = \tilde{\mathcal{F}}(\boldsymbol{\beta}, \mathbf{a}) + \mathbf{e}. \tag{14}$$

Then, with the aim of avoiding so-called joint inversion, i.e., estimating $\boldsymbol{\beta}$ and \mathbf{a} simultaneously, we introduce the approximate parameter-to-observable mapping: $\mathcal{F}(\boldsymbol{\beta}) = \tilde{\mathcal{F}}(\boldsymbol{\beta}, \mathbf{a}_*)$. That is, the approximate parameter-to-observable map is the accurate parameter-to-observable map but with the auxiliary parameters fixed to the associated mean value, i.e., $\mathbf{a} = \mathbf{a}_*$. Fixing \mathbf{a} to some other nominal value is also possible.

The goal is then to carry out the estimation of $\boldsymbol{\beta}$ using only the approximate parameter-to-observable map, $\mathcal{F}(\boldsymbol{\beta})$, while taking into account the (statistics of) the discrepancy between the models. To this end, Eq. (14) is reformulated as

$$\mathbf{d} = \tilde{\mathcal{F}}(\boldsymbol{\beta}, \mathbf{a}) + \mathbf{e} = \mathcal{F}(\boldsymbol{\beta}) + \mathbf{e} + \boldsymbol{\varepsilon} = \mathcal{F}(\boldsymbol{\beta}) + \mathbf{v}, \tag{15}$$

where $\boldsymbol{\varepsilon} = \tilde{\mathcal{F}}(\boldsymbol{\beta}, \mathbf{a}) - \mathcal{F}(\boldsymbol{\beta})$ is known as the approximation error and $\mathbf{v} = \mathbf{e} + \boldsymbol{\varepsilon}$ as the total error (e.g., Nicholson et al., 2018; Tarvainen et al., 2010). Next, the approximation error is approximated as a Gaussian with mean $\boldsymbol{\varepsilon}_*$ and covariance $\boldsymbol{\Gamma}_\varepsilon$, i.e., $\boldsymbol{\varepsilon} \sim \mathcal{N}(\boldsymbol{\varepsilon}_*, \boldsymbol{\Gamma}_\varepsilon)$. Though, formally, the approximation error depends on the parameters, i.e., $\boldsymbol{\varepsilon} = \boldsymbol{\varepsilon}(\boldsymbol{\beta}, \mathbf{a})$, a further approximation, termed the *enhanced error model* or the *composite error model* approximation, is often made, which approximates $\boldsymbol{\varepsilon}$ as independent of all parameters (Kaipio and

Kolehmainen, 2013). This leads to the total errors being distributed as $\mathbf{v} \sim \mathcal{N}(\mathbf{v}_*, \boldsymbol{\Gamma}_v) = \mathcal{N}(\boldsymbol{\varepsilon}_*, \boldsymbol{\Gamma}_\varepsilon + \boldsymbol{\Gamma}_e)$.

Use of the BAE approach results in an updated posterior density for $\boldsymbol{\beta}$:

$$\pi_{\text{post}}^{\text{BAE}}(\boldsymbol{\beta}) \propto \exp \left\{ -\frac{1}{2} \|\mathcal{F}(\boldsymbol{\beta}) - \mathbf{d} + \mathbf{v}_*\|_{\boldsymbol{\Gamma}_v^{-1}}^2 - \frac{1}{2} \|\boldsymbol{\beta} - \boldsymbol{\beta}_*\|_{\boldsymbol{\Gamma}_{\text{pr}}^{-1}}^2 \right\}, \tag{16}$$

which is obtained by explicit marginalization over \mathbf{v} (Kaipio and Kolehmainen, 2013). The updated MAP estimate is then

$$\boldsymbol{\beta}_{\text{MAP}}^{\text{BAE}} := \underset{\boldsymbol{\beta} \in \mathbb{R}^m}{\text{argmin}} \left\{ \frac{1}{2} \|\mathcal{F}(\boldsymbol{\beta}) - \mathbf{d} + \mathbf{v}_*\|_{\boldsymbol{\Gamma}_v^{-1}}^2 + \frac{1}{2} \|\boldsymbol{\beta} - \boldsymbol{\beta}_*\|_{\boldsymbol{\Gamma}_{\text{pr}}^{-1}}^2 \right\}. \tag{17}$$

This updated expression for the MAP estimate is only a slight modification of the original MAP estimate given in Eq. (9); thus reformulating the corresponding adjoint, incremental forward, and incremental adjoint equations is essentially trivial. Lastly, the updated posterior covariance matrix (under the Laplace approximation) is now given by

$$\boldsymbol{\Gamma}_{\text{po}}^{\text{BAE}} = (\mathbf{F}^T(\boldsymbol{\beta}_{\text{MAP}})\boldsymbol{\Gamma}_v^{-1}\mathbf{F}(\boldsymbol{\beta}_{\text{MAP}}) + \boldsymbol{\Gamma}_{\text{pr}}^{-1})^{-1}. \tag{18}$$

4.1 Computing the approximation error statistics

In the current paper, all parameters are taken to have Gaussian (prior) distributions, i.e., $z \sim \mathcal{N}(z_*, C_z)$, with $z = (\boldsymbol{\beta}, \mathbf{a})$. We also assume $\boldsymbol{\beta}$ and \mathbf{a} are independent; thus specifying $\boldsymbol{\beta}_*$, \mathbf{a}_* , C_β , and C_a fully describes the prior density.

Unlike the statistics of the parameters and the measurement noise, the mean ($\boldsymbol{\varepsilon}_*$) and covariance ($\boldsymbol{\Gamma}_\varepsilon$) of the approximation errors must in general be estimated based on (Monte Carlo) samples. That is,

$$\boldsymbol{\varepsilon}_* = \frac{1}{N} \sum_{\ell=1}^N \boldsymbol{\varepsilon}^{(\ell)}, \text{ and } \boldsymbol{\Gamma}_\varepsilon = \frac{1}{N-1} \mathbf{e}\mathbf{e}^T, \tag{19}$$

with $N \in \mathbb{N}$ the number of samples, $\boldsymbol{\varepsilon}^{(\ell)} = \tilde{\mathcal{F}}(\boldsymbol{\beta}^{(\ell)}, \mathbf{a}^{(\ell)}) - \mathcal{F}(\boldsymbol{\beta}^{(\ell)})$ for $\ell = 1, 2, \dots, N$, where $\boldsymbol{\beta}^{(\ell)}$ and $\mathbf{a}^{(\ell)}$ are samples drawn from the joint prior density, and $\mathbf{e} = [\boldsymbol{\varepsilon}^{(1)} - \boldsymbol{\varepsilon}_*, \boldsymbol{\varepsilon}^{(2)} - \boldsymbol{\varepsilon}_*, \dots, \boldsymbol{\varepsilon}^{(N)} - \boldsymbol{\varepsilon}_*]$. The samples, $\boldsymbol{\beta}^{(\ell)}$ and $\mathbf{a}^{(\ell)}$, are generated efficiently as in Villa et al. (2021, Eq. 30).

It is worth noting that all sampling and computations of the approximation errors and the associated statistics can be carried out prior to the acquisition of any data and is thus often termed *offline computations* (Kaipio and Kolehmainen, 2013). Furthermore, though the computational cost per sample of $\boldsymbol{\varepsilon}$ in the current paper is two forward (nonlinear) Stokes solves, the sampling procedure is *embarrassingly parallel*, i.e., each sample can be carried out independently, and in practice, only a fairly small number of samples is required.

We conclude this section by giving several rules of thumb relating to the use of the BAE approach (for more details see Kaipio and Kolehmainen, 2013). Firstly, the total number of samples required to accurately construct the statistics of the approximation errors is generally (often substantially) less than $N = 1000$. Secondly, two measures have been developed to identify when neglecting the approximation errors can result in misleading, and potentially infeasible, results (Kaipio and Kolehmainen, 2013). Specifically, if either

$$\text{trace}(\Gamma_e) < \|\epsilon_*\|^2 + \text{trace}(\Gamma_\epsilon) \tag{20}$$

holds or if for any $w \in \mathbb{R}^q$

$$w^T \Gamma_e w < (w^T \epsilon_*)^2 + w^T \Gamma_\epsilon w \tag{21}$$

holds, then the approximation errors are said to *dominate* the noise, and ignoring them often gives erroneous results. Intuitively, if the approximation errors dominate the noise, then ignoring them often results in overconfidence in the approximate forward model, in turn leading to overly confident and biased posterior estimates.

5 Numerical examples

In this section, we outline three numerical examples to assess the applicability, performance, and robustness of the BAE approach to account for uncertain rheology parameters. Several additional examples are provided in the accompanying Supplement, as detailed in Sect. 6.4. In all cases, the parameter of interest is the basal sliding coefficient, β . Any other unknown/uncertain parameters are (approximately) premarginalized over using the BAE approach, as outlined in Sect. 4.

The forward problems considered here are based on the models used in the Ice Sheet Model Intercomparison Project for Higher-Order ice sheet Models (ISMIP-HOM) benchmark study carried out in Pattyn et al. (2008). Accordingly, all problems are considered in box-like geometries, i.e., $\Omega = [0, L]^{d-1} \times [0, H]$, for $d = 2$ (in examples 1 and 2) or $d = 3$ (in Example 3). Furthermore, in all model problems we take the ice slab to be set on an incline plane with slope $\theta = 0.1^\circ$, the density of the ice to be $\rho = 910 \text{ kg m}^{-d}$, and the gravitational acceleration constant to be $g = 9.81 \text{ m s}^{-2}$. For all examples, we set the length at $L = 10^4 \text{ m}$, while for examples 1 and 2 we set $H = 250 \text{ m}$, and for Example 3, we set $H = 10^3 \text{ m}$. In Fig. 1, we show a two-dimensional schematic of the model problem setup.

The true basal sliding coefficient fields used for each example are based on those in Petra et al. (2012). Specifically, letting $\omega = 2\pi/L$, for examples 1 and 2 (posed in two dimensions) we set

$$\beta(s) = 7 + \sin(\omega s), \quad \forall s \in \Gamma_b, \tag{22}$$

as shown in Fig. 7, while in Example 3 (posed in three dimensions) we set

$$\beta(s) = 7 + 3 \sin(\omega s_1) \sin(\omega s_2), \quad \forall s \in \Gamma_b, \tag{23}$$

as shown in Fig. 13.

For all numerical experiments, we use synthetic measurements; these are randomly placed noisy point-wise measurements of each component of the velocity on the top surface of the domain, i.e., at points on Γ_t . Examples 1 and 2 are carried out based on $q = 80$ measurement locations, while for Example 3 we use $q = 100$ measurement locations. These measurements are obtained by adding zero mean white noise to the solution of the forward problem. Thus the additive noise is of the form $e \sim \mathcal{N}(\mathbf{0}, \Gamma_e)$ with covariance matrix $\Gamma_e = \delta_e^2 \mathbf{I}$. We take δ_e to satisfy $\delta_e = (\frac{1}{100}) \times (\max(\mathcal{B}u(\beta_{\text{true}})) - \min(\mathcal{B}u(\beta_{\text{true}})))$, i.e., the noise level is 1 % of the range of the noiseless synthetic measurements. The precise noise level is problem specific; however, when using GPS techniques and InSAR velocity measurements, a 1 % noise level is realistic (see, for example, Martin and Monnier, 2014 and the references therein).

For all examples considered here, the prior mean for the basal sliding coefficient, β , is set at $\beta_* = 7$. On the other hand, the prior covariance operator, C_β , is identical for examples 1 and 2, while for Example 3 different controlling parameters are used (details are provided in Table 1). Along with the true basal sliding coefficient used in examples 1 and 2, we also show the prior distribution and three samples drawn from the prior in Fig. 7. In Fig. 12, we show four samples from the prior used in Example 3.

5.1 Example problems

We now give the specific details of each model problem and make apparent which parameters we treat as auxiliary parameters and subsequently premarginalize over. Key details about each model problem are summarized in Table 1.

Example 1: uncertain flow rate factor in the two-dimensional linear Stokes ice sheet model

The first example is carried out assuming a linearized (Stokes) ice sheet model in two-dimensions. Specifically, we set $n = 1$ in Eq. (1), resulting in the effective viscosity being given by $\eta(x) = \frac{1}{2}A(x)^{-1}$. The flow rate factor, A , is taken to be unknown and spatially varying, as is often the case in reality. We represent the flow rate factor as $A = A_0 \exp(-na(x))$, with $A_0 = 2.140373 \times 10^{-7} \text{ Pa}^{-1} \text{ a}^{-1}$, $n = 1$ is the Glen’s flow law exponent, and the *pre-factor*, $\exp(-na(x))$, takes the role of the auxiliary parameter, which will subsequently be premarginalized over using the BAE approach. The pre-factor accounts for several physical and computational phenomena, such as the Arrhenius relationship between $A(x)$ and the ice temperature (e.g., Cuffey and Paterson, 2010; Zhu et al., 2016) and the use of the so-




Figure 2. Samples of the flow rate pre-factor for Example 1. Left column: samples for Example 1a. Right column: samples for Example 1b. The samples in the top row are taken as the true flow rate pre-factors. Note that the axes have been stretched in the y direction for ease of visualization.

called enhancement factors (Cuffey and Paterson, 2010; Ma et al., 2010). The “exp” function is used to ensure the pre-factor remains positive.

The prior distribution of the flow rate pre-factor is set by taking the prior mean to be $a_* = 0$, while the controlling parameters of the prior covariance operator are given in Table 1. The true pre-factor and three other samples drawn from this prior distribution are shown in Fig. 2 for examples 1a and 1b. As outlined below, the computational meshes used for examples 1a and 1b are different. This leads to differences in the true pre-factor used for both examples. This in turn results in different synthetic data being used for the inversions; however, in both cases, the standard deviation of the noise is $\delta_e \approx 0.07$. In both cases, the flow rate pre-factor is discretized using continuous quadratic Lagrange basis functions.

We use this example to also demonstrate that the proposed approach is independent of the discretization, a critical property to have when aiming to solve large-scale problems. This is done by considering identical problems on two different levels of discretization. Specifically, we consider the problem on two structured meshes having substantially different levels of discretization.

- In the first case, the computational mesh consists of 2000 triangular elements, which results in the discretized velocity and pressure having 8400 degrees of freedom (DOFs), and 1100 DOFs, while the basal sliding coefficient has 100 unknowns, and the flow rate pre-factor has 4200 DOFs.
- In the second case, the mesh is refined, and it consists of 8000 triangular elements, leading to 32 800, and 4200 DOFs for the discretized velocity and pressure, respectively, while the dimensions of the basal sliding coefficient and the flow rate pre-factor are 200 and 16 400, respectively.

Example 2: uncertain Glen’s flow law exponent in the two-dimensional nonlinear Stokes ice sheet model

For the second example, we use the nonlinear Stokes problem (Eq. 1) as the governing equation. We take the Glen’s flow law exponent, $n(\mathbf{x})$, as an uncertain (and unknown) spatially varying auxiliary parameter, i.e., we set $a(\mathbf{x}) = n(\mathbf{x})$, and proceed to approximately premarginalize over it. The prior mean of the Glen’s flow law exponent is set to $a_* = 3$, while the parameters controlling the covariance operator, \mathcal{C}_a , are given in Table 1 and are chosen to ensure that the Glen’s flow law exponents are in line with the literature. As noted, a shear-thinning rheology is generally used when modeling ice sheets, and we thus enforce $1 \leq n(\mathbf{x})$. In the current paper, this is done by *rejection sampling*, which corresponds to constraining the function space in which n lies (see Dashti and Stuart, 2016, Eq. 10.10 for details), though other methods could also be used, such as reparameterizing the Glen’s flow law exponent.

We also use this example to study the effect of larger modeling errors (i.e., *excessive errors*). That is, we consider the case when the variance of the approximation errors is so large that essentially all information in the data is washed out. As we shall see, however, the resulting uncertainty estimates are still feasible. To induce larger uncertainties (and resulting approximation errors), we alter the prior covariance operator for the Glen’s flow law exponent, n , to favor more highly oscillatory realizations. We can thus further divide Example 2 into two cases:

- modest approximation errors and
- excessive approximation errors.

The parameters used to control the covariance of the distributions on n are shown in Table 1. The true Glen’s flow law exponents used to generate the data for examples 2a and 2b are drawn from the respective distributions, which, along with several other samples of the Glen’s flow law exponent from each of the distributions, are shown in Fig. 3. In both cases, the Glen’s flow law exponent is discretized using continuous linear Lagrange basis functions, while the computational mesh used is the same as that used in Example 1a. Finally, in Example 2a we have $\delta_e \approx 0.04$, while in Example 2b we have $\delta_e \approx 0.05$.

Example 3: uncertain flow rate factor in the three-dimensional nonlinear Stokes ice sheet model

In this example, we consider a three-dimensional ($d = 3$), nonlinear analogue of Example 1. Specifically, we consider Eq. (1) in three dimensions, with the Glen’s flow law exponent set to $n = 3$. Similar to Example 1, we suppose the flow rate factor is spatially heterogeneous, unknown, and parameterized as $A = A_0 \exp(-na(\mathbf{x}))$. Here the nominal value for the flow rate factor is set to $A_0 = 10^{-16} \text{ Pa}^{-3} \text{ a}^{-1}$, the Glen’s flow law exponent is set to $n = 3$, and the pre-factor,

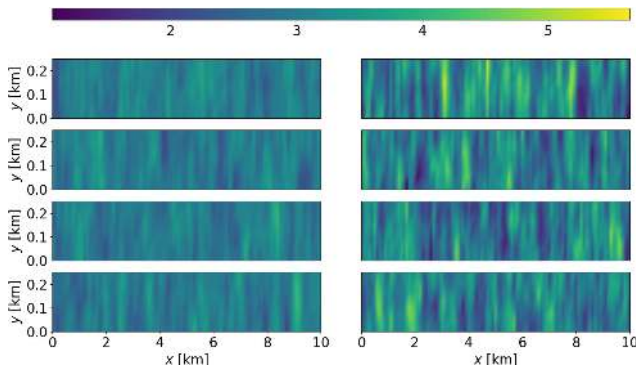


Figure 3. Samples of the Glen’s flow law exponent for Example 2. Left column: samples for Example 2a. Right column: samples for Example 2b. The samples in the top row are taken as the true Glen’s flow law exponents. Note that the axes have been stretched in the y direction for ease of visualization.




Figure 4. Samples of the flow rate pre-factor for Example 3. The top-left sample is taken as the true flow rate pre-factor. Note that the domain has been stretched in the z direction for ease of visualization.

$\exp(-na(\mathbf{x}))$, takes into account several physical and computational phenomena as described previously.

The mean value of the auxiliary parameter is set at $a_* = 0$, while the parameters controlling the distribution of the pre-factor are given in Table 1. These values for the prior covariance operator of a ensure the flow rate values are in line with those presented in the literature (see, for example, Table 3.4 in Cuffey and Paterson, 2010). In Fig. 4, we show the true flow rate pre-factor (a sample from the prior) along with three other samples from the associated prior density. Unlike Example 1, the flow rate pre-factor in this example is discretized using continuous linear Lagrange basis functions. The computational mesh used consists of 19 200 tetrahedral elements, leading to 81 600 DOFs for the velocity, 3600 DOFs for the pressure, 27 200 DOFs for the flow rate pre-factor, and 400 DOFs for the basal sliding coefficient.

5.2 Estimates and approximate posterior covariances

For each of the examples listed above, we compare the estimation results (MAP points and approximate posterior co-

variances) for three different approaches. Within each example, for each of the approaches, the same prior distribution is used for the basal sliding coefficient; thus it is only the associated likelihoods that differ. In our analysis, we place particular emphasis on the feasibility of the posterior estimates, that is, whether or not the computed posterior distributions support the true basal sliding coefficient. The three different approaches considered are as follows.

- (a) *The accurate case (REF).* In this case, any auxiliary parameters are set to their true values; i.e., we use $\tilde{\mathcal{F}}(\boldsymbol{\beta}, a_{\text{true}})$ as the parameter-to-observable map. REF is computed as a benchmark/reference. The resulting likelihood for REF is

$$\pi^{\text{REF}}(\mathbf{d}|\boldsymbol{\beta}) \propto \exp \left\{ -\frac{1}{2} \left\| \tilde{\mathcal{F}}(\boldsymbol{\beta}, a_{\text{true}}) - \mathbf{d} \right\|_{\Gamma_e^{-1}}^2 \right\}, \quad (24)$$

while the accurate MAP estimate and the corresponding posterior covariance matrix are denoted by $\boldsymbol{\beta}_{\text{MAP}}^{\text{REF}}$ and $\Gamma_{\text{po}}^{\text{REF}}$, respectively.

- (b) *The conventional error model approach (CEM).* This approach uses the standard error model (induced by the additive error, \mathbf{e}) while using the approximate model, $\mathcal{F}(\boldsymbol{\beta})$, where the auxiliary parameters are set to some nominal value (such as $a = a_*$). The likelihood is then of the form

$$\pi^{\text{CEM}}(\mathbf{d}|\boldsymbol{\beta}) \propto \exp \left\{ -\frac{1}{2} \left\| \mathcal{F}(\boldsymbol{\beta}) - \mathbf{d} \right\|_{\Gamma_e^{-1}}^2 \right\}. \quad (25)$$

We denote the corresponding MAP estimate and the posterior covariance matrix by $\boldsymbol{\beta}_{\text{MAP}}^{\text{CEM}}$ and $\Gamma_{\text{po}}^{\text{CEM}}$, respectively.

- (c) *The Bayesian approximation error approach (BAE).* This approach also uses the approximate model, $\mathcal{F}(\boldsymbol{\beta})$, but accounts for the approximation errors using the BAE approach outlined in Sect. 4. As given in Eq. (17), the updated likelihood found using the BAE approach is

$$\pi^{\text{BAE}}(\mathbf{d}|\boldsymbol{\beta}) \propto \exp \left\{ -\frac{1}{2} \left\| \mathcal{F}(\boldsymbol{\beta}) - \mathbf{d} + \mathbf{v}_* \right\|_{\Gamma_v^{-1}}^2 \right\}, \quad (26)$$

with the MAP estimate and the posterior covariance matrix denoted by $\boldsymbol{\beta}_{\text{MAP}}^{\text{BAE}}$ and $\Gamma_{\text{po}}^{\text{BAE}}$, respectively.

6 Results

Here we discuss and compare MAP estimates for the basal sliding coefficient and the respective approximate posterior covariance for each example. As alluded to previously, we pay particular attention to the feasibility of the posterior uncertainty estimates when comparing the results. We also examine the spectrum of the prior preconditioned misfit Hessians, which gives further insight into the uncertainty, and

Table 1. Details for each of the examples considered. The first column (Ex.) refers to the example number; the second, third, and fourth columns give details of the forward model used, including which Stokes model is used, the aspect ratio, $\frac{L}{H}$, and the definition of the auxiliary parameter; the fifth, sixth, and seventh columns give the discretization details, including the number of degrees of freedom for the velocities and pressure, the number of degrees of freedom of the unknown parameters, (β, a) DOFs, and the number of measurements, q ; finally, the 8th through 12th columns give details on the prior distributions for the unknowns, including the parameters controlling the prior covariance operator for β , the controlling parameters for the prior covariance operator of a , and the prior mean, a_* . Note that in examples 2a and 2b the prior for the auxiliary parameter is further constrained by enforcing $1 \leq n(x)$, while for all examples the prior mean for the basal sliding coefficient is taken as $\beta_* = 7$. Furthermore, the correlation length for β in examples 1 and 2 is approximately 4900 m, while in Example 3 the correlation length is approximately 3200 m. Finally, in the definition of the auxiliary parameter for Example 3, the Glen’s flow law exponent is $n = 3$.

Ex.	Model details			Discretization details			Prior distribution details				
	Stokes	$\frac{L}{H}$	a	(\mathbf{u}, p) DOFs	(β, a) DOFs	q	γ_β	δ_β	a_*	γ_a	δ_a
1a	Linear 2D	40	$\ln(A_0/A)$	(8400, 1100)	(100, 4200)	80	840	7.0×10^{-5}	0	300	1.5×10^{-4}
1b	Linear 2D	40	$\ln(A_0/A)$	(32 800, 4200)	(200, 16 400)	80	840	7.0×10^{-5}	0	300	1.5×10^{-4}
2a	Nonlinear 2D	40	n	(8400, 1100)	(100, 1100)	80	840	7.0×10^{-5}	3	90	9.0×10^{-3}
2b	Nonlinear 2D	40	n	(8400, 1100)	(100, 1100)	80	840	7.0×10^{-5}	3	41	4.1×10^{-3}
3	Nonlinear 3D	10	$\frac{1}{n} \ln(A_0/A)$	(81 600, 3600)	(400, 27 200)	100	7.5	7.5×10^{-7}	0	12.5	2.5×10^{-6}

the sensitivity of each approach. To conclude the section, we give a brief comparison of the online computational costs (in terms of linearized Stokes partial differential equation, PDE, solves) for computing the MAP estimates.

To solve the optimization problems, we use an inexact Newton-CG method (see, for example, Petra et al., 2012). In all cases, we start the optimization procedure using the prior mean for the initial estimate of the basal sliding coefficient, while the prior covariance operator is used as the preconditioner. The optimization is carried out using a Gauss–Newton Hessian approximation for the first five iterations and then full Newton combined with an Armijo line search (Nocedal and Wright, 2006). Convergence is established when the gradient has decreased by a factor of 10^6 relative to the norm of the initial gradient.

The numerical results presented in this paper are obtained using hIPPYlib (an inverse problem Python library; Villa et al., 2018, 2021). The hIPPYlib library implements state-of-the-art scalable adjoint-based algorithms for PDE-based deterministic and Bayesian inverse problems. It builds on FEniCS (Dupont et al., 2003; Logg et al., 2012) for the discretization of the PDEs and on PETSc (Balay et al., 2009) for scalable and efficient linear algebra operations and solvers needed for the solution of the PDEs. In line with the finite element discretization used for the weak form of the forward problem (Eq. 2), in what follows, we use Taylor–Hood finite elements for the adjoint, incremental forward, and incremental adjoint equations, as in Petra et al. (2012).

6.1 Example 1

In this example, we consider the case of an uncertain flow rate factor in the two-dimensional linear Stokes ice sheet model and demonstrate the mesh independence of the approach. We begin by discussing the statistics of the ap-

proximation errors which are induced by treating the unknown flow rate factor as a known constant, specifically, $A = 2.140373 \times 10^{-7} \text{ Pa}^{-1} \text{ a}^{-1}$. In Fig. 5, we show the marginal distribution of the approximation errors in the x component (top) and y component (bottom) for Example 1a (left) and Example 1b (right). The approximation errors are similar for the coarser mesh (Example 1a) and the finer mesh (Example 1b), both having fairly constant mean and variance in each component. For both examples, the mean of the approximation errors in the x component of the velocity measurements is non-zero, $\epsilon_* \approx 0.2$, while the standard deviation of the approximation errors is substantially larger than the additive noise ($\delta_\epsilon \approx 0.07$). That is, the approximation errors dominate the additive noise, as explained in Sect. 4.1, and it is likely (and is indeed the fact) that ignoring the approximation errors may lead to infeasible results.

To illustrate the convergence of the approximation errors, in the top row of Fig. 6, we show the spectrum of the approximation errors covariance matrices, $\mathbf{\Gamma}_\epsilon$, for examples 1a and 1b for increasing sample sizes ($N = 62, 250, 500,$ and 1000). From the figure, it is evident that for $N \geq 250$ samples, the spectra essentially coincide and have both converged, thus demonstrating the discretization independence of the approach and that approximately $N \geq 250$ samples are likely sufficient to characterize the approximation error statistics. Note, however, that the results displayed here use $N = 1000$ samples.

To give further insight into the distribution of the approximation errors, we show the (Pearson’s) correlation matrix of the approximation errors for Example 1a (left) and 1b (right) in the bottom row of Fig. 6. Firstly, we notice that the correlation matrices are highly structured (unlike the noise covariance matrix which is diagonal). It is also apparent that the correlation matrices are (visually) identical, further il-

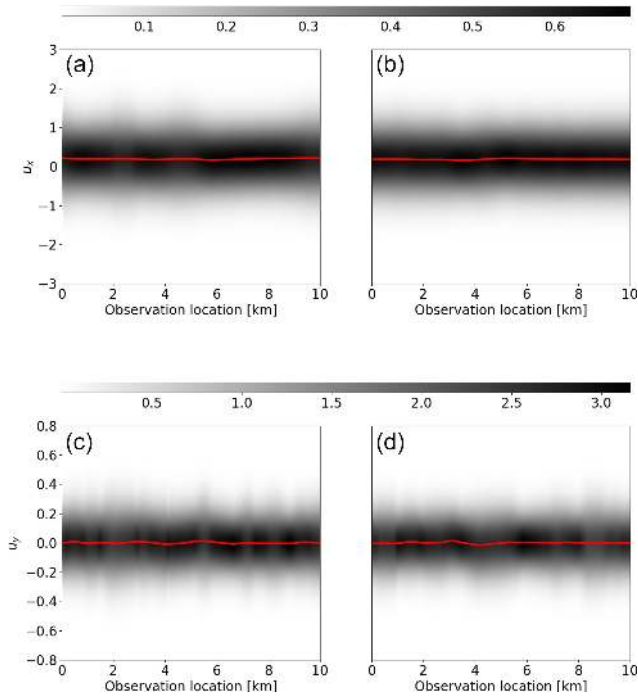


Figure 5. Second order statistics of the approximation errors for Example 1. (a, b) Distribution of the approximation errors in the x direction velocity measurements for Example 1a (a) and 1b (b). (c, d) Distribution of the approximation errors in the y direction velocity measurements for Example 1a (c) and 1b (d). The mean of the approximation errors, ϵ_* , is indicated with a red line, while higher probability density is indicated by darker shading.

illustrating the discretization independence. The 2×2 block structure of the correlation matrices is to be expected since the measurement number indexing used corresponds to measuring the $q = 80$ velocity measurements in the x direction first, followed by the 80 velocity measurements in the y direction. The behavior within the diagonal blocks is also fairly intuitive as periodic boundary conditions are used, while the structure also illustrates that measurements (relatively) far away from each other are fairly uncorrelated. Comparing the diagonal blocks we see that the approximation errors in the x component of the velocity measurements are more highly correlated at greater distances than those of the y component. Finally, the off-diagonal blocks show a nontrivial correlation between the approximation errors in the x and y components. In particular, this figure shows that the approximation errors have a similar structure to the main diagonal but reveals smaller (Pearson’s) correlation coefficients (these range from about -0.5 to 0.5).

In the top row of Fig. 7, we show the marginal prior distributions and the resulting marginal posterior distributions. Also shown are the corresponding MAP estimates, the true basal sliding coefficient, and three draws from each of the distributions. Firstly, the accurate MAP estimate, $\beta_{\text{MAP}}^{\text{REF}}$, is in good agreement with the true basal sliding coefficient,

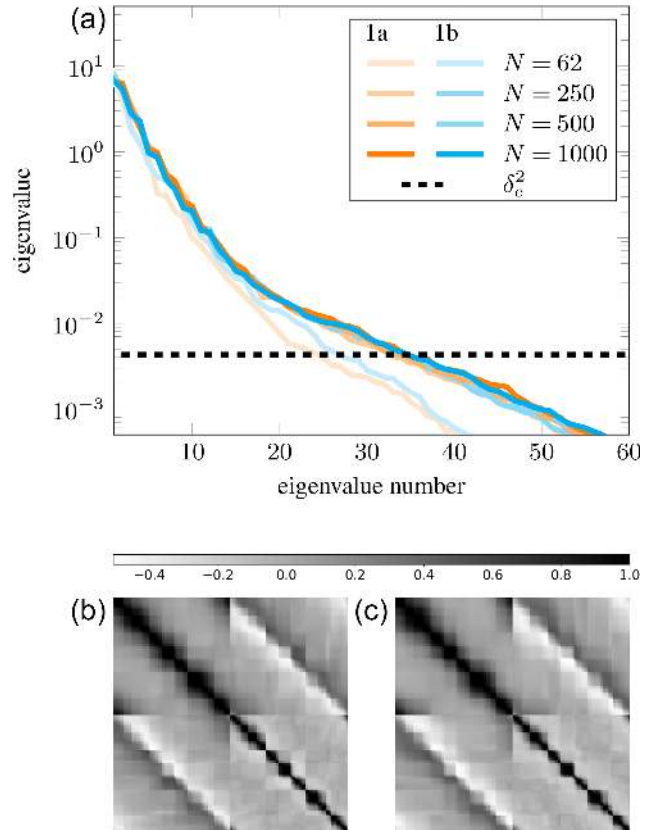


Figure 6. Convergence and (Pearson’s) correlation matrix of the approximation errors for Example 1. (a) Spectrum of Γ_ϵ for various sample sizes, N , for Example 1a (orange) and Example 1b (cyan), along with the noise variance, δ_ϵ^2 . (b, c) (Pearson’s) correlation matrix of the approximation errors for Example 1a (b) and 1b (c).

while the accurate posterior is clearly feasible in the sense that the true basal sliding coefficient is well supported by the Laplace-approximated posterior. On the other hand, the MAP estimate found using the conventional error model, $\beta_{\text{MAP}}^{\text{CEM}}$, differs substantially from the true basal sliding coefficient over most of the domain. Furthermore, the posterior is essentially infeasible, with the actual coefficient having virtually no posterior density. Conversely, the MAP estimate found using the BAE approach, $\beta_{\text{MAP}}^{\text{BAE}}$, is in fairly good agreement with the true coefficient, and the Laplace-approximated posterior supports the truth well. We do see that the marginal posterior standard deviations found using the BAE approach are somewhat larger than those found using the accurate and conventional error approaches. This is typical, and to be expected, as the additional uncertainty in the flow rate pre-factor manifests itself as extra posterior uncertainty.

In the bottom row of Fig. 7, we show the corresponding results for Example 1b. The results are fairly similar to Example 1a when using the accurate approach and when using

the BAE approach⁴ despite the substantial difference in the discretizations used. Lastly, the MAP estimate found using the conventional error model has changed drastically from Example 1a, though the posterior is equally as bad.

6.2 Example 2

In this example, we consider the case of an uncertain Glen's flow law exponent in the two-dimensional nonlinear Stokes ice sheet model and also demonstrate what happens when the approximation errors are, in some sense, too large. The approximation errors here are the result of treating the unknown, and spatially varying, Glen's flow law exponent as a fixed constant, i.e., setting $n = n_0 = 3$. To induce the larger approximation errors for Example 2b compared to Example 2a, we increase the uncertainty in the Glen's flow law exponent by altering the associated prior distribution (see Table 1). The difference in magnitude of the approximation errors is apparent in Fig. 8, where we show the marginal distribution of the approximation errors at the observation locations in the x direction (top) and y direction (bottom) velocities for examples 2a (left) and 2b (right). Note that the variance of the approximation errors for Example 2b is substantially larger than that of Example 2a. Considering that the standard deviation of the added noise for the *small* approximation error case (Example 2a) is $\delta_{e_a} \approx 0.04$ and for the *large* approximation error case (Example 2b) is $\delta_{e_b} \approx 0.05$, the approximation errors in both examples dominate the noise (see Sect. 4.1).

The top row of Fig. 9 shows the spectrum of the covariance matrices of the approximation errors, Γ_ε , for $N = 125, 500, 1000$, and 2000 samples for Example 2a and for $N = 250, 1000, 2000$, and 4000 samples for Example 2b. For Example 2a, it appears $N \approx 500$ is enough samples, though for the results here we used $N = 1000$, while for Example 2b we require $N \approx 2000$ samples. The fact that more samples are required to ensure convergence of the approximation errors in Example 2b follows naturally from the increased uncertainty. It is worth pointing out that Example 2b is used mainly to demonstrate how the BAE approach performs in the presence of *too much* modeling uncertainty; thus for the purposes of the current study we deem taking $N = 2000$ as *tolerable*.

In the bottom row of Fig. 9, we show the (Pearson's) correlation matrices of the approximation errors. The correlation matrices for this example share several of the characteristics seen in the corresponding correlation matrices in Example 1, specifically, the block structure and general behavior within the blocks. Comparing the correlation matrices for examples 2a and 2b, it appears the approximation errors in the x component for Example 2a are more highly correlated at greater distances towards the edges of the computational

domain compared to the approximation errors in the x component of the velocity measurements for Example 2b.

In the top row of Fig. 10, we show the marginal prior and Laplace-approximated posterior distributions, as well as three draws from each of the distributions, the corresponding MAP estimates, and the true basal sliding coefficient for Example 2a. A couple of conclusions can be drawn from this figure. First, the accurate MAP estimate, $\beta_{\text{MAP}}^{\text{REF}}$, closely resembles the true basal sliding coefficient, and the truth is well supported by the accurate posterior distribution. Second, the Laplace-approximated posterior found using the conventional error approach is infeasible for most of the right half of the domain, with the MAP estimate, $\beta_{\text{MAP}}^{\text{CEM}}$, (severely) underestimating the true basal sliding coefficient. Third, the true basal sliding coefficient lies well within the bulk of the (Laplace-approximated) posterior for the BAE approach, with the MAP estimate, $\beta_{\text{MAP}}^{\text{BAE}}$, in fairly good agreement with the true basal sliding coefficient.

In the bottom row of Fig. 10, we show the corresponding results for Example 2b, in which the approximation errors are excessive. Under the Laplace approximation, the accurate posterior, found by using the true Glen's flow law exponent, remains an accurate representation of the truth as in Example 2a. The posterior found using the conventional error model approach has significantly deteriorated, however, with the true basal sliding coefficient even more markedly underestimated and the truth lying well outside the bulk of the Laplace-approximated posterior over almost all of the domain. Conversely, by taking into account the excessive modeling errors in Example 2b, the posterior found using the BAE approach is comparable to the prior, with the corresponding MAP estimate, $\beta_{\text{MAP}}^{\text{BAE}}$, being fairly similar to the prior mean. This demonstrates that when using the BAE approach, as the modeling errors become larger, the corresponding posterior density tends towards the prior, as should be hoped, to avoid overconfidence in biased results.

6.3 Example 3

In this example, we consider an uncertain flow rate factor in a larger scale, three-dimensional nonlinear Stokes ice sheet model. The approximation errors are the result of setting the unknown flow rate factor to $A = 10^{-16} \text{ Pa}^{-3} \text{ a}^{-1}$. The spectrum for the approximation errors are shown in Fig. 11. The plot indicates that taking $500 < N \leq 1000$ samples would likely suffice to accurately characterize the approximation errors. For the results discussed here, we used $N = 1000$ samples. The average standard deviation of the approximation errors in the x component of the approximation errors is approximately 3.1, while for the y and z components, the average standard deviation of the approximation errors are 0.5 and 0.4, respectively. The standard deviation of the noise, on the other hand, is $\delta_e \approx 0.25$. We thus can expect the resulting estimates found by disregarding the approximation errors to be unreasonable (see Sect. 4.1).

⁴We attribute the differences in the BAE approach to the differences in the true flow rate pre-factor, the noise realization, and the specific samples of ε .




Figure 7. Prior and MAP estimates of the basal sliding parameter for Example 1. Top row: Example 1a prior (far left), accurate/reference (REF) case (center left), conventional error model (CEM) case (center right), and Bayesian approximation error (BAE) case (far right). The bottom row shows the same plots for Example 1b. In each plot, the mean of the distribution (blue line) is shown along with the true basal sliding parameter, β_{true} (red line), three samples from the respective distributions (green lines), the marginal distribution (shaded) with darker shading indicating higher probability, and the ± 2 (approximate) standard deviation intervals (dashed black line).

In Fig. 12, we show four draws from the prior density on the basal sliding coefficient, while in Fig. 13 we show the true basal sliding coefficient (top left) and each of the MAP estimates: $\beta_{\text{MAP}}^{\text{REF}}$ (top right), $\beta_{\text{MAP}}^{\text{CEM}}$ (bottom left), and $\beta_{\text{MAP}}^{\text{BAE}}$ (bottom right). We also show the locations ($y = 2.5$ km and $y = 7.5$ km) of two lines, labeled l_1 and l_2 , for which cross-sectional plots are shown in Fig. 14. It is clear from Figs. 13 and 14 that the reference posterior is completely feasible, and the corresponding MAP estimate is in good agreement with the true basal sliding coefficient. On the other hand, although the MAP estimate found using the conventional error model (CEM) shows similar qualitatively behavior, as seen in Fig. 13, when taking the corresponding posterior density into account, it is clear that the approach is essentially infeasible, with the truth lying well outside the bulk of the posterior across most of the domain (see Fig. 14). Finally, from Figs. 13 and 14 we see that, though not as good as the accurate case, the MAP estimate found using the BAE approach qualitatively remains similar to the truth. Furthermore, the truth is generally very well supported by the BAE posterior under the Laplace approximation.

6.4 Additional examples in the Supplement

To further demonstrate the flexibility and robustness of the proposed approach, we provide several additional numerical examples in the Supplement accompanying this paper. Specifically, we consider three additional cases that are variations of Example 1. Section S.1 of the Supplement demonstrates the robustness of the BAE approach in the case in which the true distribution of the auxiliary parameter is not

known. The results show that as long as the true auxiliary parameter is well supported by the assumed distribution, the BAE approach provides posterior-consistent estimates of the basal sliding coefficient β . Section S.2 compares the CEM and BAE approaches for different assumptions of the mean and marginal variance of the prior distribution. The results show that the qualitative behavior of the two approaches is consistent to that observed for Example 1. Finally, Sect. S.3 compares the BAE approach and a so-called tempering approach, in which a heuristic criterion (such as the L-curve) is used to select an appropriate scaling of the prior or likelihood density. This tempering approach can be understood as varying the regularization parameter in a deterministic setup. This example demonstrates that the BAE approach provides a robust solution to the inverse problem without requiring multiple solutions of the inverse problem for different scaling parameters. In addition, in contrast to the tempering approach, the BAE approach does not rely on (possibly) unreliable heuristic methods to select the scaling parameter.

6.5 Spectra of the data misfit Hessians and computational costs

In this section, we compare the spectra of the data misfit Hessian and the computational cost of the three approaches (accurate, conventional error, and Bayesian approximation error) for each of the three examples. The dominant eigenvalues of the data misfit Hessian, $\bar{\mathbf{H}}$ (see Eq. 10), evaluated at the corresponding MAP estimate, are shown in Fig. 15 for each example. Firstly, we observe that for all three approaches in all three examples, we only need to retain a relatively low

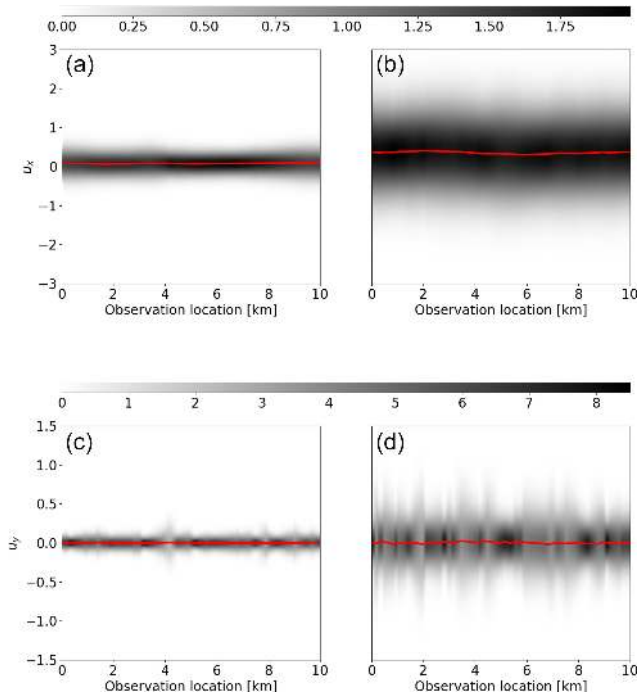


Figure 8. Second order statistics of the approximation errors for Example 2. (a, b) Distribution of the approximation errors in the x direction velocity measurements for Example 2a (a) and 2b (b). (c, d) Distribution of the approximation errors in the y direction velocity measurements for Example 2a (c) and 2b (d). The mean of the approximation errors, ϵ_* , is indicated with a red line, while higher probability density is indicated by darker shading.

number of eigenvalues and eigenvectors to compute a reasonable low-rank approximation of the Laplace posterior covariance matrix. Secondly, we see that the dominant spectrum resulting from using the BAE approach is often lower than that of the reference and conventional error approach cases. This is to be expected since we are accommodating the approximation errors, which naturally lead to an increase in uncertainty. Finally, the dominant spectrum of the misfit Hessian for examples 1 and 2, found using the conventional error approach, further illustrates the fact that ignoring the uncertainty in the auxiliary parameters can lead to overconfidence in erroneous estimates.

The spectrum of the misfit Hessian for Example 3, found using the conventional error approach, seems to be somewhat anomalous in that the spectrum decays faster than that of the misfit Hessian found using the BAE approach. However, this is possibly explained by the fact that the respective misfit Hessians are evaluated at quite different MAP estimates.

Figure 15 shows that the number of eigenvalues required to compute a reasonable low-rank approximation, in the sense of Eq. (12), is considerably lower for the BAE approach in most of the examples. This result suggests that computing the low-rank approximation is cheaper for the BAE approach compared to the other two approaches.

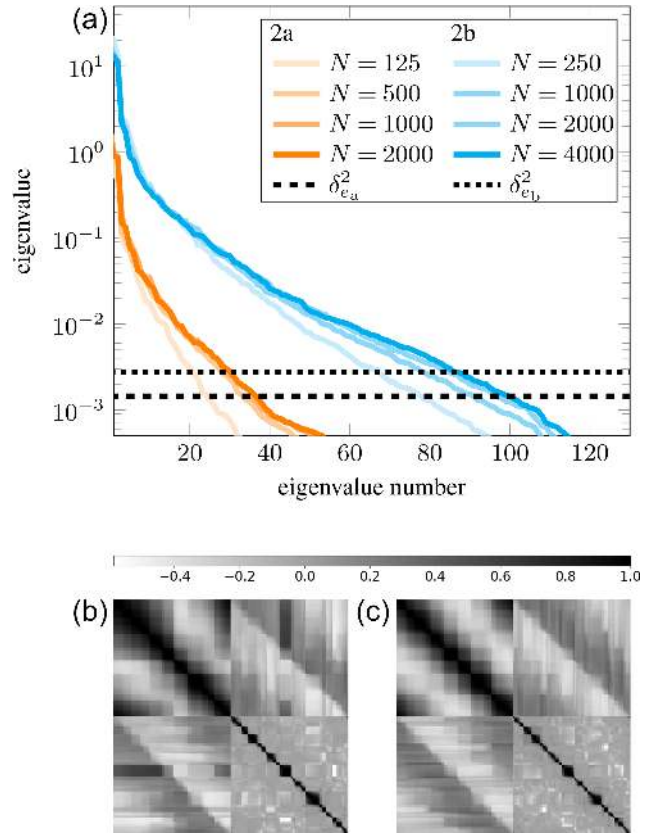


Figure 9. Convergence and (Pearson's) correlation matrix of the approximation errors for Example 2. (a) Spectrum of Γ_ϵ for various sample sizes, N , for Example 2a (orange) and Example 2b (cyan), along with the noise variance for Example 2a, $\delta^2_{\epsilon_a}$, and Example 2b, $\delta^2_{\epsilon_b}$. (b, c) (Pearson's) correlation matrix of the approximation errors for Example 2a (b) and Example 2b (c).

With regard to the computational cost, we consider the number of (linearized) Stokes problem solves required for the optimization algorithm to converge as the unit of cost. As stated in Sect. 3, we use the inexact Newton-CG algorithm with Armijo line search to find the MAP point. At each iteration, inexact Newton-CG requires the following:

- (a) one (or more if required to satisfy the sufficient descent condition) evaluation of the log-likelihood, which involves solving the nonlinear Stokes equations;
- (b) one gradient evaluation, which involves solving an additional linearized Stokes problem, i.e., the adjoint equation; and
- (c) one Newton system solve using the conjugate gradient (CG) method, which at each CG iteration requires solving two linearized Stokes problems, i.e., the incremental forward and adjoint problems.

The total number of linearized Stokes solves required to compute the MAP estimate can then be calculated – per Gauss–

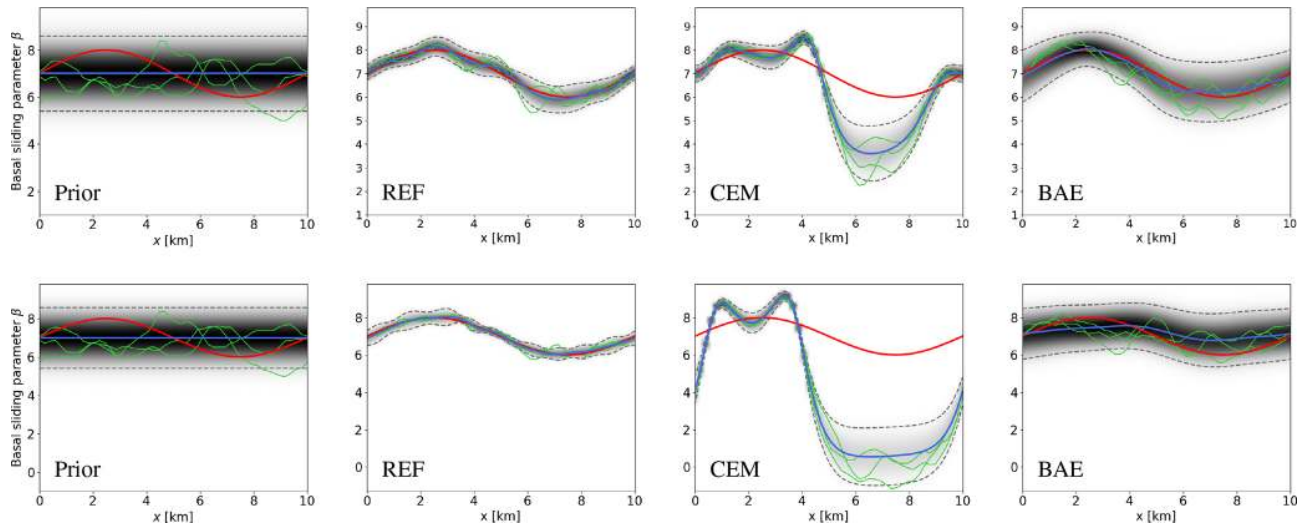


Figure 10. Prior and MAP estimates of the basal sliding parameter for Example 2. Top row: Example 2a prior (far left), accurate/reference (REF) case (center left), conventional error model (CEM) case (center right), and Bayesian approximation error (BAE) case (far right). The bottom row shows the same plots for Example 2b. In each plot, the mean of the distribution (blue line) is shown along with the true basal sliding parameter, β_{true} (red line), three samples from the respective distributions (green lines), the marginal distribution (shaded) with darker shading indicating higher probability, and the ± 2 (approximate) standard deviation intervals (dashed black line).




Figure 11. Spectrum of Γ_ϵ for various sample sizes, N , for Example 3, along with the noise variance, δ_ϵ^2 .

Newton iteration – as the sum of the number of iterations required to solve the nonlinear forward problem plus one adjoint solve (to calculate the gradient), along with one incremental forward solve and one incremental adjoint solve (to calculate the action of the Hessian) per CG iteration. To ensure a sufficient decrease in the objective function at each (Gauss–)Newton iteration, the forward problem may be solved multiple times until the Armijo condition is satisfied, thus further increasing the number of linearized Stokes solves.




Figure 12. Four samples from the prior for the basal sliding parameter field for Example 3.

The results shown in Table 2 indicate that in each of the examples considered in this paper, the BAE approach generally requires less than half the number of the linearized Stokes solves that are required for the REF case to converge to the MAP point. Furthermore, the conventional error approach re-




Figure 13. Basal sliding parameter estimates. (a, b) The true basal sliding coefficient (a) and the accurate MAP estimate $\beta_{\text{MAP}}^{\text{REF}}$ (b). (c, d) The conventional error model MAP estimate $\beta_{\text{MAP}}^{\text{CEM}}$ (c) and the Bayesian approximation error MAP estimate $\beta_{\text{MAP}}^{\text{BAE}}$ (d). The dashed black lines are used to show the location of the cross sections ($l_1 = 2.5$ km and $l_2 = 7.5$ km) for Fig. 14.

quires (in some cases, significantly) more iterations, and thus linearized Stokes solves, than the accurate model. This is to be expected as the optimization is hampered by model mismatch. It is also worth noting that the CEM approach requires substantially more backtracking iterations compared to the REF and BAE approaches, which is in line with Nicholson et al. (2018). Furthermore, the number of CG iterations is significantly reduced for the BAE approach when compared to the CEM and REF case.

7 Conclusions

In this paper, we have considered the inference for the basal sliding coefficient field for ice sheet flow problems with uncertain rheology from surface velocity measurements. The rheology parameters of the ice, in particular the flow rate factor and the Glen's flow law exponent, are often uncertain and can, at best, only be estimated in practice. We considered examples in both two and three dimensions and used both the linear and nonlinear Stokes ice sheet model. In each of the cases considered, our goal was to infer the basal sliding coefficient only; as such the unknown rheology parameters were a priori fixed to nominal values and treated as auxiliary parameters. To account for the resulting modeling uncertainties (or errors), we employed the Bayesian approximation er-

Table 2. The cost of solving for the MAP estimates, measured in number of linearized Stokes solves. The first column (Ex.) refers to the example number, and the second column (Est.) refers to which MAP estimate we are solving for, i.e., the reference MAP (REF), the MAP found using the conventional error model (CEM), or the MAP found using the BAE approach (BAE). The third column (#N) gives the number of (Gauss–)Newton iterations, while the fourth column (#CG) reports the total number of CG iterations. The fifth column (#back) reports the number of backtracks needed throughout the (Gauss–)Newton iterations, and the sixth column (#O) gives the total number of objective function evaluations. Finally, the last column (#Stokes) gives the total number of linearized Stokes solves (for forward, adjoint, incremental forward, and incremental adjoint problems). The (Gauss–)Newton iterations are terminated when the norm of the gradient is decreased by a factor of 10^6 , while the CG iterations are terminated in line with the Eisenstat–Walker condition (Eisenstat and Walker, 1996) (to avoid over-solving) and the Steihaug criteria (Steihaug, 1983) (to avoid negative curvature). The results illustrate that the use of the approximation error approach can be carried out at no additional online cost compared to the conventional error approach and reference case.

Ex.	Est.	#O	#N	#CG	#back	#Stokes
1a	REF	10	10	68	0	156
	CEM	19	17	96	2	228
	BAE	7	7	30	0	74
1b	REF	12	11	68	1	159
	CEM	18	16	91	2	216
	BAE	7	7	28	0	70
2a	REF	21	16	66	5	244
	CEM	42	27	99	15	376
	BAE	12	11	25	1	110
2b	REF	12	11	78	1	241
	CEM	35	26	62	9	286
	BAE	8	8	18	0	78
3	REF	18	16	178	2	491
	CEM	37	23	115	14	438
	BAE	14	13	64	1	240

ror (BAE) approach. This approach shifts all uncertainty into a single additive total error term, which is approximated as Gaussian and can be premarginalized over.

The quantification of the resulting uncertainty in the estimated basal sliding coefficient was carried out based on the Laplace approximation to the posterior. In all of the examples considered here, the results suggest that fixing rheology parameters to standard values found in the literature can lead to overly confident and (heavily) biased estimates, with the true basal sliding coefficient generally lying outside the bulk of the posterior density if the uncertainty in the rheology parameters is not accounted for. Conversely, carrying out approximate premarginalization over the unknown rheology parameters, via the BAE approach, leads to feasible estimates for the basal sliding coefficient in all cases considered. To il-

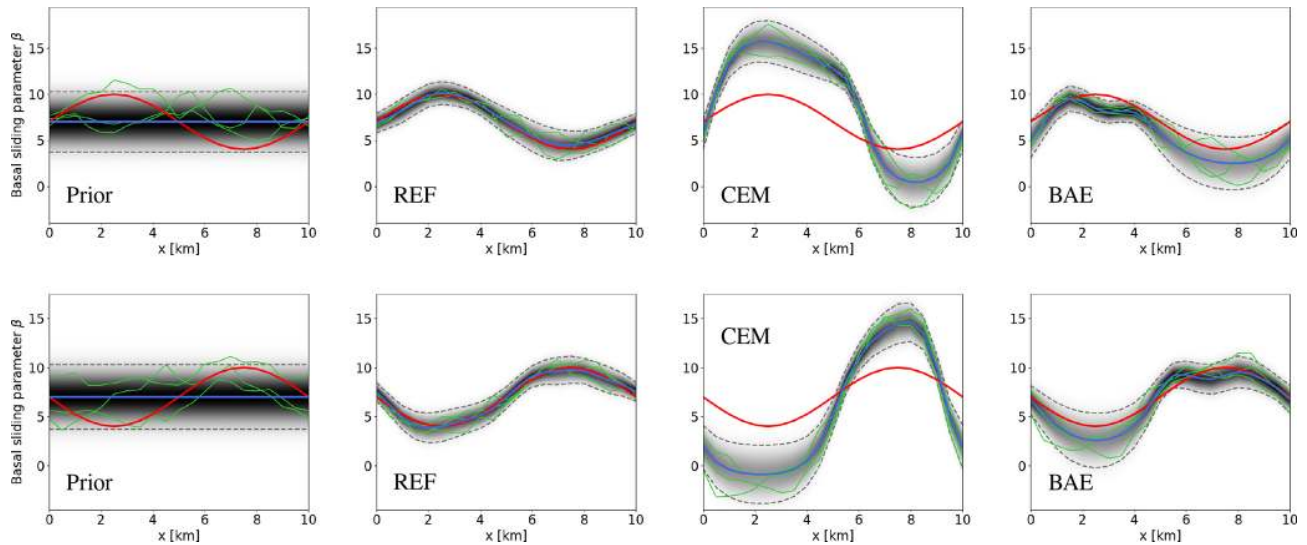


Figure 14. Cross sections of prior and MAP estimates of the basal sliding parameter for Example 3. Top row: cross section along line l_1 ($y = 2.5$ km) of prior (far left), accurate/reference (REF) case (center left), conventional error model (CEM) case (center right), and Bayesian approximation error (BAE) case (far right). The bottom row shows the cross section along line l_2 ($y = 7.5$ km) in the same order. In each plot, the mean of the distribution (blue line) is shown along with the true basal sliding parameter, β_{true} (red line), three samples from the respective distributions (green lines), the marginal distribution (shaded) with darker shading indicating higher probability, and the ± 2 (approximate) standard deviation intervals (dashed black line).

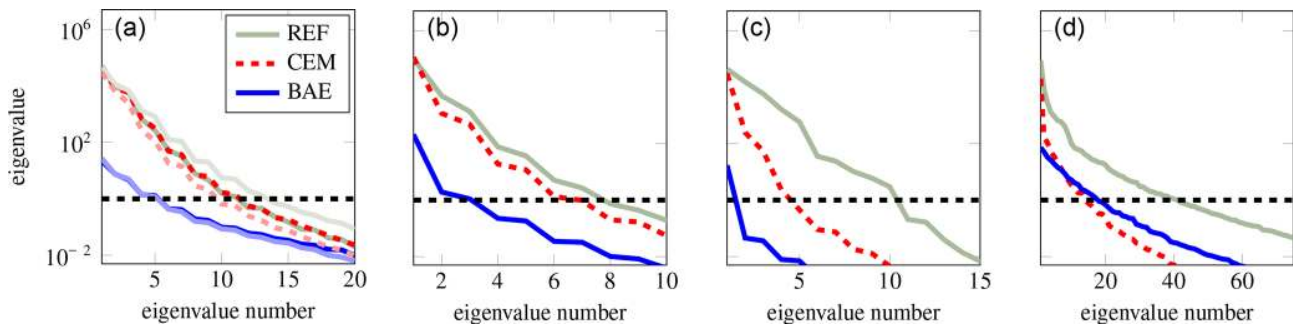


Figure 15. Spectra of the prior-preconditioned Hessian of the data misfit computed using Eq. (11) for Example 1 (a), Example 2a (b), Example 2b (c), and Example 3 (d). The spectra for Example 1a (coarse mesh) are shown in the fainter colors, while the spectra for Example 1b (fine mesh) are shown in the richer colors. The horizontal dashed black line (at $\lambda = 1$) shows the reference value for the truncation of the spectrum of the prior-preconditioned Hessian of the data misfit.

illustrate a limitation of the BAE approach, we included an example in which the modeling errors introduced were, in some sense, *too* large. This case led to a posterior density (found using the BAE approach) which showed very little reduction in variance compared to the prior, though it still contained the truth.

By avoiding the simultaneous estimation of the basal sliding coefficient and rheology parameters (which are spatially varying over the entire domain), the online computational overheads of the estimation problem are substantially reduced. To ensure the work carried out here is applicable to large-scale problems, i.e., scalable, we initially posed the problem in infinite dimensions and then employed the adjoint-state methodology to compute the MAP estimate.

In assessing the applicability and performance of the BAE approach, the current study only considers fairly limited domains, i.e., box-like geometries and idealized boundary conditions. A natural next step for future work is to apply the same framework to more realistic setups and to continental-scale ice flow problems.

Code availability. The numerical results presented in this paper were obtained using hIPPYlib version 3.0 (<https://doi.org/10.5281/zenodo.3634136>, Villa et al., 2018, 2021). The hIPPYlib software is an open-source project released under GPL v2. It is available for download at <https://hippylib.github.io> (last access: 1 April 2021).

Data availability. No external data or geometries/meshes were used in this paper. The observations for the inversion are synthetic measurements obtained via simulation. All the parameters used to obtain the data are reported in the paper.

Supplement. The supplement related to this article is available online at: <https://doi.org/10.5194/tc-15-1731-2021-supplement>.

Author contributions. All authors contributed equally to the design of numerical experiments, the discussion and interpretation of the numerical results, and the writing of the original and revised version of this paper. OB implemented the numerical methods and ran the numerical experiments with help from all co-authors.

Competing interests. The authors declare that they have no conflict of interest.

Acknowledgements. The authors gratefully acknowledge computing time on the Multi-Environment Computer for Exploration and Discovery (MERCED) cluster at UC Merced, which was funded by National Science Foundation grant no. ACI-1429783. We thank the editor, Olivier Gagliardini, as well as the anonymous reviewer and Douglas Brinkerhoff for their valuable comments that contributed to improving the final version of this paper.

Financial support. This research has been supported by the National Science Foundation (grant nos. ACI-1550593, ACI-1550547, and ACI-1429783 and CAREER-1654311).

Review statement. This paper was edited by Olivier Gagliardini and reviewed by Douglas Brinkerhoff and one anonymous referee.

References

- Arridge, S., Kaipio, J., Kolehmainen, V., Schweiger, M., Somersalo, E., Tarvainen, T., and Vauhkonen, M.: Approximation errors and model reduction with an application in optical diffusion tomography, *Inverse Probl.*, 22, 175–195, <https://doi.org/10.1088/0266-5611/22/1/010>, 2006.
- Balay, S., Buschelman, K., Gropp, W. D., Kaushik, D., Knepley, M. G., McInnes, L. C., Smith, B. F., and Zhang, H.: PETSc Web page, available at: <http://www.mcs.anl.gov/petsc> (last access: 25 March 2021), 2009.
- Bons, P. D., Kleiner, T., Llorens, M.-G., Prior, D. J., Sachau, T., Weikusat, I., and Jansen, D.: Greenland Ice Sheet: Higher non-linearity of ice flow significantly reduces estimated basal motion, *Geophys. Res. Lett.*, 45, 6542–6548, 2018.
- Brondex, J., Gillet-Chaulet, F., and Gagliardini, O.: Sensitivity of centennial mass loss projections of the Amundsen basin to the friction law, *The Cryosphere*, 13, 177–195, <https://doi.org/10.5194/tc-13-177-2019>, 2019.
- Brynjarsdóttir, J. and O’Hagan, A.: Learning about physical parameters: The importance of model discrepancy, *Inverse Probl.*, 30, 114007, <https://doi.org/10.1088/0266-5611/30/11/114007>, 2014.
- Bui-Thanh, T., Burstedde, C., Ghattas, O., Martin, J., Stadler, G., and Wilcox, L. C.: Extreme-scale UQ for Bayesian inverse problems governed by PDEs, in: SC12: Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis, 10–16 November 2012, Salt Lake City, UT, USA, 1–11, <https://doi.org/10.1109/SC.2012.56>, 2012.
- Bui-Thanh, T., Ghattas, O., Martin, J., and Stadler, G.: A Computational Framework for Infinite-Dimensional Bayesian Inverse Problems Part I: The Linearized Case, with Application to Global Seismic Inversion, *SIAM J. Sci. Comput.*, 35, A2494–A2523, 2013.
- Bulthuis, K., Arnst, M., Sun, S., and Pattyn, F.: Uncertainty quantification of the multi-centennial response of the Antarctic ice sheet to climate change, *The Cryosphere*, 13, 1349–1380, <https://doi.org/10.5194/tc-13-1349-2019>, 2019.
- Castello, D. A. and Kaipio, J. P.: Modeling errors due to Timoshenko approximation in damage identification, *Int. J. Numer. Meth. Eng.*, 120, 1148–1162, 2019.
- Cuffey, K. M. and Paterson, W. S. B.: *The physics of glaciers*, Academic Press, Cambridge, MA, 2010.
- Daon, Y. and Stadler, G.: Mitigating the Influence of Boundary Conditions on Covariance Operators Derived from Elliptic PDEs, *Inverse Probl. Imag.*, 12, 1083–1102, 2018.
- Dashti, M. and Stuart, A. M.: The Bayesian approach to inverse problems, in: *Handbook of Uncertainty Quantification*, edited by: Ghanem, R., Higdon, D., and Owhadi, H., Springer International Publishing, Cham, Switzerland, 311–428, https://doi.org/10.1007/978-3-319-12385-1_7, 2016.
- Dupont, T., Hoffman, J., Johnson, C., Kirby, R. C., Larson, M. G., Logg, A., and Scott, L. R.: *The FEniCS project*, Chalmers Finite Element Centre, Chalmers University of Technology, 2003.
- Eisenstat, S. C. and Walker, H. F.: Choosing the forcing terms in an inexact Newton method, *SIAM J. Sci. Comput.*, 17, 16–32, 1996.
- Elman, H. C., Silvester, D. J., and Wathen, A. J.: *Finite Elements and Fast Iterative Solvers with applications in incompressible fluid dynamics*, Oxford University Press, Oxford, 2005.
- Flath, P. H., Wilcox, L. C., Akçelik, V., Hill, J., van Bloemen Waanders, B., and Ghattas, O.: Fast Algorithms for Bayesian Uncertainty Quantification in Large-Scale Linear Inverse Problems Based on Low-Rank Partial Hessian Approximations, *SIAM J. Sci. Comput.*, 33, 407–432, 2011.
- Gillet-Chaulet, F., Hindmarsh, R. C., Corr, H. F., King, E. C., and Jenkins, A.: In-situ quantification of ice rheology and direct measurement of the Raymond Effect at Summit, Greenland using a phase-sensitive radar, *Geophys. Res. Lett.*, 38, 2011.
- Gillet-Chaulet, F., Gagliardini, O., Seddik, H., Nodet, M., Durand, G., Ritz, C., Zwinger, T., Greve, R., and Vaughan, D. G.: Greenland ice sheet contribution to sea-level rise from a new-generation ice-sheet model, *The Cryosphere*, 6, 1561–1576, <https://doi.org/10.5194/tc-6-1561-2012>, 2012.
- Giudici, M., Baratelli, F., Comunian, A., Vassena, C., and Cattaneo, L.: Model calibration for ice sheets and glaciers dynamics: a general theory of inverse problems in glaciology, *The Cryosphere Discuss.*, 8, 5511–5537, <https://doi.org/10.5194/tcd-8-5511-2014>, 2014.

- Glen, J. W.: The creep of polycrystalline ice, *P. Roy. Soc. Lond. A. Mat.*, 228, 519–538, 1955.
- Gockenbach, M. S.: Understanding and implementing the finite element method, Society for Industrial and Applied Mathematics (SIAM), Philadelphia, PA, 2006.
- Golub, G. H. and Van Loan, C. F.: *Matrix Computations*, Johns Hopkins University Press, Baltimore, Maryland, USA, 1996.
- Hutter, K.: *Theoretical Glaciology, Mathematical Approaches to Geophysics*, D. Reidel Publishing Company, Springer Netherlands, Heidelberg, Germany, 1983.
- Isaac, T., Petra, N., Stadler, G., and Ghattas, O.: Scalable and efficient algorithms for the propagation of uncertainty from data through inference to prediction for large-scale problems with application to flow of the Antarctic ice sheet, *J. Comput. Phys.*, 296, 348–368, 2015a.
- Isaac, T., Stadler, G., and Ghattas, O.: Solution of nonlinear Stokes equations discretized by high-order finite elements on nonconforming and anisotropic meshes, with application to ice sheet dynamics, *SIAM J. Sci. Comput.*, 37, B804–B833, 2015b.
- Kaipio, J. and Kolehmainen, V.: *Bayesian Theory and Applications*, chap. Approximate Marginalization Over Modeling Errors and Uncertainties in Inverse Problems, pp. 644–672, Oxford University Press, Oxford, UK, 2013.
- Kaipio, J. and Somersalo, E.: *Statistical and Computational Inverse Problems*, vol. 160 of *Applied Mathematical Sciences*, Springer-Verlag New York, 2005.
- Kaipio, J. and Somersalo, E.: Statistical Inverse Problems: Discretization, Model Reduction and Inverse Crimes, *J. Comput. Appl. Math.*, 198, 493–504, 2007.
- Khristenko, U., Scarabosio, L., Swierczynski, P., Ullmann, E., and Wohlmuth, B.: Analysis of Boundary Effects on PDE-Based Sampling of Whittle–Matérn Random Fields, *SIAM/ASA Journal on Uncertainty Quantification*, 7, 948–974, 2019.
- Lamien, B., Le Maux, D., Courtois, M., Pierre, T., Carin, M., Le Masson, P., Orlande, H. R. B., and Paillard, P.: A Bayesian approach for the estimation of the thermal diffusivity of aerodynamically levitated solid metals at high temperatures, *Int. J. Heat Mass Transf.*, 141, 265–281, 2019.
- Lindgren, F., Rue, H., and Lindström, J.: An explicit link between Gaussian fields and Gaussian Markov random fields: The stochastic partial differential equation approach, *J. R. Stat. Soc. B*, 73, 423–498, 2011.
- Logg, A., Mardal, K.-A., and Wells, G. N.: *Automated Solution of Differential Equations by the Finite Element Method*, vol. 84, in: *Lecture Notes in Computational Science and Engineering*, Springer, Berlin, Germany, 2012.
- Ma, Y., Gagliardini, O., Ritz, C., Gillet-Chaulet, F., Durand, G., and Montagnat, M.: Enhancement factors for grounded ice and ice shelves inferred from an anisotropic ice-flow model, *J. Glaciol.*, 56, 805–812, 2010.
- Marshall, S. J.: Recent advances in understanding ice sheet dynamics, *Earth Planet. Science Lett.*, 240, 191–204, 2005.
- Martin, N. and Monnier, J.: Adjoint accuracy for the full Stokes ice flow model: limits to the transmission of basal friction variability to the surface, *The Cryosphere*, 8, 721–741, <https://doi.org/10.5194/tc-8-721-2014>, 2014.
- Morlighem, M., Seroussi, H., Larour, E., and Rignot, E.: Inversion of basal friction in Antarctica using exact and incomplete adjoints of a higher-order model, *J. Geophys. Res.-Earth*, 118, 1746–1753, 2013.
- Nicholson, R., Petra, N., and Kaipio, J. P.: Estimation of the Robin coefficient field in a Poisson problem with uncertain conductivity field, *Inverse Probl.*, 34, 115005, <https://doi.org/10.1088/1361-6420/aad91e>, 2018.
- Nocedal, J. and Wright, S.: *Numerical optimization*, Springer Science & Business Media, New York, NY, USA, 2006.
- Osborn, S., Vassilevski, P. S., and Villa, U.: A multilevel, hierarchical sampling technique for spatially correlated random fields, *SIAM J. Sci. Comput.*, 39, S543–S562, 2017.
- Paterson, W. S. B.: *The Physics of Glaciers*, 3rd edn., Butterworth Heinemann, Pergamon Oxford, UK, 1994.
- Pattyn, F., Perichon, L., Aschwanden, A., Breuer, B., de Smedt, B., Gagliardini, O., Gudmundsson, G. H., Hindmarsh, R. C. A., Hubbard, A., Johnson, J. V., Kleiner, T., Kononov, Y., Martin, C., Payne, A. J., Pollard, D., Price, S., Rückamp, M., Saito, F., Souček, O., Sugiyama, S., and Zwinger, T.: Benchmark experiments for higher-order and full-Stokes ice sheet models (ISMIP–HOM), *The Cryosphere*, 2, 95–108, <https://doi.org/10.5194/tc-2-95-2008>, 2008.
- Perego, M., Price, S., and Stadler, G.: Optimal initial conditions for coupling ice sheet models to Earth system models, *J. Geophys. Res.-Earth*, 119, 1894–1917, 2014.
- Petra, N., Zhu, H., Stadler, G., Hughes, T. J. R., and Ghattas, O.: An inexact Gauss-Newton method for inversion of basal sliding and rheology parameters in a nonlinear Stokes ice sheet model, *J. Glaciol.*, 58, 889–903, 2012.
- Petra, N., Martin, J., Stadler, G., and Ghattas, O.: A computational framework for infinite-dimensional Bayesian inverse problems: Part II. Stochastic Newton MCMC with application to ice sheet inverse problems, *SIAM J. Sci. Comput.*, 36, A1525–A1555, 2014.
- Pollard, D. and DeConto, R. M.: A simple inverse method for the distribution of basal sliding coefficients under ice sheets, applied to Antarctica, *The Cryosphere*, 6, 953–971, <https://doi.org/10.5194/tc-6-953-2012>, 2012.
- Raymond, M. J. and Gudmundsson, G. H.: Bayesian estimation of basal conditions on Rutford Ice Stream, West Antarctica, from surface data, *J. Glaciol.*, 57, 315–324, 2011.
- Raymond, M. J. and Gudmundsson, G. H.: Estimating basal properties of ice streams from surface measurements: a nonlinear Bayesian inverse approach applied to synthetic data, *The Cryosphere*, 3, 265–278, <https://doi.org/10.5194/tc-3-265-2009>, 2009.
- Roininen, L., Huttunen, J. M. J., and Lasanen, S.: Whittle–Matérn priors for Bayesian statistical inversion with applications in electrical impedance tomography, *Inverse Probl. Imag.*, 8, 561, 2014.
- Schlegel, N.-J., Larour, E., Seroussi, H., Morlighem, M., and Box, J.: Ice discharge uncertainties in Northeast Greenland from boundary conditions and climate forcing of an ice flow model, *J. Geophys. Res.-Earth*, 120, 29–54, 2015.
- Schoof, C.: The effect of cavitation on glacier sliding, *P. R. Soc. A-Math. Phys.*, 461, 609–627, 2005.
- Schoof, C.: Coulomb friction and other sliding laws in a higher-order glacier flow model, *Math. Mod. Meth. Appl. S.*, 20, 157–189, 2010.

- Steihaug, T.: Local and superlinear convergence for truncated iterated projections methods, *Math. Program.*, 27, 176–190, <https://doi.org/10.1007/BF02591944>, 1983.
- Stuart, A. M.: Inverse problems: A Bayesian perspective, *Acta Numer.*, 19, 451–559, 2010.
- Tarantola, A.: *Inverse Problem Theory and Methods for Model Parameter Estimation*, SIAM, Philadelphia, PA, 2005.
- Tarvainen, T., Kolehmainen, V., Pulkkinen, A., Vauhkonen, M., Schweiger, M., Arridge, S. R., and Kaipio, J. P.: Approximation error approach for compensating modelling errors in optical tomography, in: *Biomedical Optics, BSuD48*, 11–14 April 2010, Miami, Florida, USA, Optical Society of America, BSuD48, 2010.
- Truffer, M.: The basal speed of valley glaciers: an inverse approach, *J. Glaciol.*, 50, 236–242, 2004.
- Van der Veen, C. J.: *Fundamentals of glacier dynamics*, CRC press, Boca Raton, Florida, 2013.
- Villa, U., Petra, N., and Ghattas, O.: hIPPYlib: An Extensible Software Framework for Large-Scale Inverse Problems, *J. Open Source Soft.*, 3, 115005, <https://doi.org/10.21105/joss.00940>, 2018 (data available at <https://doi.org/10.5281/zenodo.3634136>).
- Villa, U., Petra, N., and Ghattas, O.: hIPPYlib: An Extensible Software Framework for Large-Scale Inverse Problems Governed by PDEs; Part I: Deterministic Inversion and Linearized Bayesian Inference, *ACM Transactions of Mathematical Software*, 47, 16, <https://doi.org/10.1145/3428447>, 2021.
- Zhao, C., Gladstone, R. M., Warner, R. C., King, M. A., Zwinger, T., and Morlighem, M.: Basal friction of Fleming Glacier, Antarctica – Part 1: Sensitivity of inversion to temperature and bedrock uncertainty, *The Cryosphere*, 12, 2637–2652, <https://doi.org/10.5194/tc-12-2637-2018>, 2018a.
- Zhao, C., Gladstone, R. M., Warner, R. C., King, M. A., Zwinger, T., and Morlighem, M.: Basal friction of Fleming Glacier, Antarctica – Part 2: Evolution from 2008 to 2015, *The Cryosphere*, 12, 2653–2666, <https://doi.org/10.5194/tc-12-2653-2018>, 2018b.
- Zhu, H., Petra, N., Stadler, G., Isaac, T., Hughes, T. J. R., and Ghattas, O.: Inversion of geothermal heat flux in a thermomechanically coupled nonlinear Stokes ice sheet model, *The Cryosphere*, 10, 1477–1494, <https://doi.org/10.5194/tc-10-1477-2016>, 2016.