Author Contributions:
Conceptualization: B. Riel, B. Minchew, T. Bischoff
Data curation: B. Riel
Formal analysis: B. Riel
Funding acquisition: B. Riel, B. Minchew
Investigation: B. Riel
Methodology: B. Riel, B. Minchew, T. Bischoff
Software: B. Riel, T. Bischoff

Data-Driven Inference of the Mechanics of Slip Along Glacier Beds Using Physics-Informed Neural Networks: Case Study on Rutford Ice Stream, Antarctica

B. Riel[1] , **B. Minchew[1]** , **and T. Bischoff[2]**

[1]Department of Earth, Atmospheric and Planetary Sciences, Massachusetts Institute of Technology, Cambridge, MA, USA, [2]Climate Modeling Alliance, California Institute of Technology, Pasadena, CA, USA

Abstract Reliable projections of sea-level rise depend on accurate representations of how fast-flowing glaciers slip along their beds. The mechanics of slip are often parameterized as a constitutive relation (or "sliding law") whose proper form remains uncertain. Here, we present a novel deep learning-based framework for learning the time evolution of drag at glacier beds from time-dependent ice velocity and elevation observations. We use a feedforward neural network, informed by the governing equations of ice flow, to infer spatially and temporally varying basal drag and associated uncertainties from data. We test the framework on 1D and 2D ice flow simulation outputs and demonstrate the recovery of the underlying basal mechanics under various levels of observational and modeling uncertainties. We apply this framework to time-dependent velocity data for Rutford Ice Stream, Antarctica, and present evidence that ocean-tide-driven changes in subglacial water pressure drive changes in ice flow over the tidal cycle.

Plain Language Summary The relation between slip of glaciers along their beds and the level of basal drag at the ice-bed interface is a critical component of ice dynamics for fast-flowing glaciers and ice streams. However, uncertainty surrounding the proper form of this relation, often referred to as the sliding law, has hindered efforts to reliably project the contribution of the Greenland and Antarctic ice sheets to future sea-level rise. Here, we utilize the tools of physics-informed deep learning to learn the evolution of drag at glacier beds from time-dependent ice velocity and elevation observations. By training a neural network with both data reconstruction losses and ice physics-based losses, we are able to reconstruct the evolution of drag for glaciers and ice streams undergoing changes in flow speed and surface elevations. Thus, we can investigate the relation between slip and basal drag without specifying the form of the sliding law. We use this approach to present observational evidence that ocean-tide-driven changes in flow speed for Rutford Ice Stream, Antarctica are driven by changes in subglacial water pressure. Ultimately, this approach provides a natural way to integrate our existing knowledge of ice flow physics with remote sensing data in order to improve flow models.

1. Introduction

Fast-flowing outlet glaciers that drain the Greenland and Antarctic Ice Sheets are major contributors to sea-level rise (SLR) (Church et al., 2013; Ritz et al., 2015). While widespread acceleration of these glaciers in response to changing climate conditions has magnified their importance in future projections of SLR, fundamental uncertainties about their long-term dynamical behavior and stability persist (Robel et al., 2019). One of the key sources of uncertainty is the unknown form of the parameterization used to describe how drag at the base of glaciers is related to basal sliding velocity, bed roughness, bed composition, and water pressure (Aschwanden et al., 2019; Ritz et al., 2015). The resistive force provided by basal drag plays a significant role in the evolution of glaciers in response to changes in atmospheric and oceanic conditions. The collection of proposed parameterizations for basal drag are commonly referred to as sliding laws, and the wide range of physical processes governing the interaction between ice, bed materials, and basal hydrology have led to a wide spectrum of proposed sliding laws for quantifying drag dependence on sliding velocity. Despite considerable advances in our understanding of the mechanics of slip along glacier beds, no consensus has emerged as to which sliding law offers the best balance of model simplicity and model fidelity, though one model has emerged as a candidate for a universal sliding law that is applicable to glaciers with rigid and deforming beds (Joughin et al., 2019; Schoof, 2005; Zoet & Iverson, 2020). Since the physical

Validation: B. Riel, B. Minchew
Visualization: B. Riel
Writing – original draft: B. Riel
Writing – review & editing: B. Riel,
B. Minchew, T. Bischoff

processes meant to be represented by the sliding law are not directly observable outside of laboratory settings, inference of the form of the the sliding law and the value of its parameters requires inverse modeling using observations of ice surface velocity and elevation coupled with an accurate physical model of ice flow (Bondzio et al., 2017; Gillet-Chaulet et al., 2016; Joughin et al., 2012; Shapero et al., 2016).

Fortunately, the Earth science community has seen a sharp rise in remote sensing data availability over the past two decades. This rise is due to an ever-increasing number of spaceborne and airborne Earth-observing platforms in combination with increased computational capabilities and data-providing services that operationally produce analysis-ready data sets. The glaciology community has benefited enormously from continent-wide observations of ice surface velocities and elevation over much of the Greenland and Antarctic Ice Sheets (Howat et al., 2019; Joughin et al., 2011; Porter et al., 2018; Rignot et al., 2011). However, within the context of modeling of basal drag, use of these observations in modeling efforts have traditionally involved assimilating instantaneous or time-averaged velocity observations into ice flow models in order to estimate static distributions of basal drag (Larour et al., 2012; MacAyeal, 1993; Morlighem et al., 2010; Shapero et al., 2016). A few studies have expanded upon this approach by estimating basal drag at different time epochs, which can be used to attribute changes in drag to known changes in environmental factors like surface meltwater (e.g., Minchew et al., 2016) or to better constrain parameters in sliding laws (e.g., Gillet-Chaulet et al., 2016; Habermann et al., 2013). The increasing availability of time-dependent velocity fields, which in many places capture the evolution of glacier velocities on sub-monthly timescales in high-spatial resolution, could potentially provide much finer resolution on the time-evolution of basal drag in order to obtain the underlying space- and time-varying functional form of the sliding law.

Within the past decade, machine learning algorithms have exploded in popularity due to their ability to discover patterns and relationships in large volumes of data which are used to inform numerous predictive and analytical tasks (LeCun et al., 2015). In particular, the recent success of deep learning has been attributed to the ability to learn hierarchical, abstract features in unstructured data which can interact in highly non-linear ways (Bengio et al., 2013). The coincident increase in computing power, in large part from the increased utility of graphics processing units, has led to rapid development of specialized network architectures able to learn patterns from video streams, images, and word sequences. Recently, many studies have demonstrated the potential for deep learning algorithms to be integrated with scientific knowledge in order to bridge theoretic gaps, discover new and robust patterns in scientific observations, and predict the evolution of dynamical systems (Karpatne et al., 2017; Raissi, 2018; Reichstein et al., 2019). This type of "theory-guided" learning combines the robustness provided by decades of theoretical and experimental work with the pattern recognition and representation power of deep learning.

In this work, we develop a hybrid modeling framework that can exploit contemporary remote sensing data by incorporating well-known ice dynamics and constitutive laws with a deep neural network model representing the unknown sliding law. In developing this framework, our goal is to demonstrate a general approach for inferring various components of a glacier system from large volumes of data without the need to call a sophisticated numerical solver for a discretized ice flow model. We further discuss how we can pose the learning problem in a probabilistic manner that partially allows for the quantification of uncertainties due to both data errors and uncertainties in the governing equations of ice flow. Since the focus of this work is on learning a spatiotemporal representation for basal drag, we apply our method to several one- and two-dimensional flowline simulations that are representative of real-world basal sliding scenarios that would be challenging to analyze with traditional inverse modeling approaches. Finally, we apply our methods to real velocity data over Rutford Ice Stream in West Antarctica and present observational evidence for the role of subglacial hydrology in propagating tidally driven variations in ice flow roughly 100 km inland.

2. Methods

2.1. Ice Flow Governing Equations

The flow of ice is well-approximated by incompressible Stokes flow, which describes the motion of a viscous fluid where inertial forces are negligibly small relative to viscous forces. In Stokes flow, the momentum equations (stress balance) reduce to gravitational body forces resisted by stresses induced through ice deformation and shear stresses at the interfaces between ice and the bed and sidewalls. For many fast-flowing

outlet glaciers and ice streams, flow is dominated by basal sliding where sliding velocity is comparable to surface velocity, and forward motion due to vertical shearing is negligible. In this case, the full three-dimensional Stokes equations can be reduced by neglecting certain components of the stress divergence and averaging the resulting momentum balance over depth (see Appendix A). This approximation, commonly referred to as the Shallow Ice Shelf/Stream Approximation (SSA), leads to the following two-dimensional relation in a Cartesian coordinate system with z defined parallel to the gravity vector:

$$\frac{\partial}{\partial x}\left(2\eta h\left(2\frac{\partial u}{\partial x} + \frac{\partial v}{\partial y}\right)\right) + \frac{\partial}{\partial y}\left(\eta h\left(\frac{\partial u}{\partial y} + \frac{\partial v}{\partial x}\right)\right) - \tau_{bx} = \rho_i gh\frac{\partial s}{\partial x}, \tag{1a}$$

$$\frac{\partial}{\partial y}\left(2\eta h\left(2\frac{\partial v}{\partial y} + \frac{\partial u}{\partial x}\right)\right) + \frac{\partial}{\partial x}\left(\eta h\left(\frac{\partial u}{\partial y} + \frac{\partial v}{\partial x}\right)\right) - \tau_{by} = \rho_i gh\frac{\partial s}{\partial y}, \tag{1b}$$

where u and v are the horizontal velocity components of the velocity vector, $\boldsymbol{u}$, along the x- and y-directions, respectively, and taken to be constant with depth; h is the ice thickness; τ_{bx} and τ_{by} represent the x- and y-components of basal drag; s is the ice surface elevation; η is the effective dynamic viscosity of ice; ρ_i is the mass density of ice; and g is the gravitational acceleration. Basal drag is modeled with a sliding law using a power-law relationship (Weertman, 1957):

$$\tau_{bx} = c_b \|\boldsymbol{u}\|^{\frac{1}{m}} \frac{u}{\|\boldsymbol{u}\|}, \tag{2a}$$

$$\tau_{by} = c_b \|\boldsymbol{u}\|^{\frac{1}{m}} \frac{v}{\|\boldsymbol{u}\|}, \tag{2b}$$

where $\|\boldsymbol{u}\| = \sqrt{u^2 + v^2}$, c_b acts as a drag coefficient, and m is a scalar. Thus, the basal drag magnitude is a (potentially nonlinear) function of the basal velocity, which is assumed to be equal to the surface/depth-averaged velocity. The range of possible values for m is wide and determines whether sliding at the bed is rate-weakening ($m < 0$, basal drag decreases with sliding velocity), rate-strengthening ($m > 0$, basal drag increases with sliding velocity), or rate-independent ($m \to \infty$). The mode of sliding can have strong implications on how stress perturbations at the termini of glaciers propagate upstream (Gudmundsson, 2011). Recent laboratory work by Zoet and Iverson (2020) has shown that these sliding modes can be represented as a spectrum of sliding behavior corresponding to rate-strengthening sliding over rigid beds at velocities below a certain threshold, above which till deformation dominates and basal drag is largely rate-independent.

In this work, we use ice flow simulations to generate synthetic data for testing our proposed learning framework, and we simulate flow in both 1D and 2D-planview (horizontal) in order to demonstrate our methods on systems of increasing complexity. For both classes of simulations, we model tidewater glaciers where their termini end at the ocean but are grounded throughout the entire modeling domain. In 1D, the particular model form we use in this work reduces the 2D momentum balance equations (Equations 1a and 1b) by assuming that lateral shear stresses are negligible, which is appropriate for ice streams that are much wider than they are thick (Schoof, 2007). Thus, Equations 1a and 1b reduce to:

$$2\frac{\partial}{\partial x}\left[h\eta\frac{\partial u}{\partial x}\right] - \tau_b = \rho_i gh\frac{\partial s}{\partial x}, \tag{3}$$

with x defined as parallel to flow. We induce velocity variations by periodically varying the longitudinal stress conditions at the terminus, which approximates the periodic rising and falling of ocean levels due to tides (Appendix A). Thus, the final simulation outputs we use as inputs and data for the machine learning models are the time-dependent velocity components and ice thickness.

We use a basal drag sliding relationship where the prefactor c_b can vary in both space and time and the exponent m can vary in space such that:

$$\tau_b\left(x, t\right) = c_b\left(x, t\right)|u|^{\frac{1}{m(x)} - 1} u. \tag{4}$$

The spatial variation of c_b and m can represent changes in bed roughness and composition, ice cavity density, and basal water pressure, among other factors. Thus, c_b acts as an aggregate parameterization of multiple basal processes rather than a specific measurement of bed friction, and its units will vary in space as m varies. The temporal variation of c_b can represent local changes in basal water pressure due to an evolving

subglacial hydrological system, which has been shown to be an important process in many fast-flowing tidewater glaciers around the globe (Schoof, 2010; I. Hewitt, 2013). The non-uniqueness of the sliding parameters c_b and m for a given value of basal drag generally requires a priori information about one of the parameters in order to constrain the other. In many modeling studies, a spatially uniform value of $m = 3$ is often assumed to model sliding over a rigid bed, thus reducing the inverse problem to spatial estimation of c_b. Therefore, simulations generated with both m and c_b variations are useful test cases to assess the feasibility of joint inference of both parameters using time-dependent velocity and elevation fields.

2.2. Learning Basal Drag Function With Hybrid Modeling

Large uncertainties on the form of the sliding law motivate a generic representation of basal drag as a function of sliding velocity as well as a function of space and time in the case of spatially varying till properties or subglacial hydrology. To that end, let us consider the following:

$$\hat{\boldsymbol{\tau}}_b = f\left(\boldsymbol{u}, h, \boldsymbol{x}, t\right), \tag{5}$$

where basal drag predictions $\hat{\boldsymbol{\tau}}_b$ are generated by a generic nonlinear function of sliding velocity, ice geometry (thickness), spatial coordinate, and time. Since basal drag is not directly observable, we must combine quantities that are readily observable (e.g., ice velocity and surface elevation) within the physical modeling framework provided by the momentum balance equations (Equations 1a and 1b). An advantage of the vertically integrated form of the momentum balance used here is that drag can be directly estimated by re-arranging terms, for example, for the 1D case (Equation 3):

$$\hat{\tau}_b(x, t) = 2\frac{\partial}{\partial x}\left(h(x, t)\eta(x, t)\frac{\partial u(x, t)}{\partial x}\right) - \rho_i gh(x, t)\frac{\partial s(x, t)}{\partial x}. \tag{6}$$

This method (also referred to as the force balance method (Van der Veen, 2013)) allows for quantification of spatial and temporal variations of drag if time-dependent measurements of surface velocity and ice geometry are available for a given glacier, in addition to knowledge of ice rheology (Cuffey & Paterson, 2010; Enderlin et al., 2018; Van der Veen, 2013).

A key requirement of the above formulation is the availability of first- and second-order spatial derivatives of velocity, surface elevation, and ice thickness. While these derivatives may be computed locally with finite differences or some other interpolation kernel, data that are scattered in space or contain large gaps in spatial coverage are likely to incur large gradient errors. An alternative approach is to construct a continuous, twice differentiable hypersurface through the observations where the hypersurface can be constructed from a physical model (e.g., ice flow model) or empirically (e.g., with a spline or polynomial fit). In this work, we follow the latter approach by modeling the velocity and elevation observations with a feedforward, dense neural network (multi-layer perceptron), $f_{\boldsymbol{\theta}}$, defined such that:

$$[\hat{\boldsymbol{u}}, \hat{h}] = f_{\boldsymbol{\theta}}\left(\boldsymbol{x}, t\right), \tag{7}$$

where the network is parameterized by $\boldsymbol{\theta}$, the weights and biases of the hidden layers. Mathematically, feedforward neural networks perform a sequence of operations for each hidden layer: (a) multiplication of vector of layer inputs with layer weights (represented as a two-dimensional tensor); and (b) application of a nonlinear activation function to the sum of the tensor-vector product and hidden layer bias vector. Layer outputs are then used as inputs to the next layer and the process is repeated. By tuning the weights $\boldsymbol{\theta}$ of the hidden layers, the neural network is "trained" to map inputs to output variables as an approximation to $f_{\boldsymbol{\theta}}\left(\boldsymbol{x}, t\right)$. Estimation of the optimal $\boldsymbol{\theta}$ is performed through minimization of an appropriate cost/loss function that generally includes the misfit between the neural network predictions and observed values at the input coordinates (see LeCun et al. (2015) for an in-depth review of neural networks and deep learning).

The utility of neural networks as universal function approximators (first formulated for a single network by Cybenko (1989) and extended to finite-width multi-layer neural networks (e.g., Delalleau & Bengio, 2011; Lu et al., 2017; Bölcskei et al., 2019)) make them well-suited to represent scattered, time-dependent surface observations with potentially complex spatiotemporal patterns. More importantly, we can evaluate derivatives of $\hat{\boldsymbol{u}}$ and $\hat{h}$ at arbitrary space and time coordinates at machine precision using automatic differentiation (Baydin et al., 2017; Raissi, 2018). Similar to a spline or polynomial fit, the neural network learns a smooth hypersurface between scattered observations in data space and can return the hypersurface value and slope

at any given point. The smoothness of this surface will depend on the network capacity (i.e., layer size and depth) and is only achieved when the activation function is twice differentiable (e.g., tanh or sigmoid functions; see Appendix B for additional network and training details). These smoothed predictions and their gradients can then be used to generate time-dependent predictions of basal drag from an appropriate momentum balance, such as Equation 6.

For surface observations with minimal noise levels and glacier geometries well-suited to the SSA model, the neural network weights $\boldsymbol{\theta}$ can be estimated by minimizing a standard mean square error (MSE) loss function over training data:

$$\mathcal{L}_{mse}\left(\boldsymbol{\theta}\right) = \frac{1}{M}\sum_{k=1}^{M}\left[\|\boldsymbol{u}^k - \hat{\boldsymbol{u}}^k\|^2 + (h^k - \hat{h}^k)^2\right], \tag{8}$$

where $f_{\boldsymbol{\theta}}^k = f_{\boldsymbol{\theta}}\left(\boldsymbol{x}^k, t^k\right)$, $\|\cdot\|$ is a standard Euclidean norm, and M data points are used for training $f_{\boldsymbol{\theta}}$. For the cases investigated in this work, M is on the order of tens of thousands of data points. After training, one can then generate predictions of $\hat{\boldsymbol{u}}$, $\hat{h}$, and $\hat{\boldsymbol{\tau}}_b$ over the entire modeling domain.

Complication, however, arises because surface observations are generally noisy, with noise characteristics that vary for different data sources. This observation noise is amplified by the need for first- and second-order gradients, which are further amplified by nonlinearities in the ice dynamics. In most real-world cases, the intrinsic spatial wavelength of observation noise is considerably smaller than the wavelength we expect ice dynamics to be sensitive to since viscous ice flow effectively acts as a low pass filter to any spatial variations in bed topography (De Rydt et al., 2013; Gudmundsson, 2003; Habermann et al., 2012). Consequently, application of the momentum balance to noisy surface observations will lead to large, un-physical variations in inferred basal drag. A typical strategy for mitigating observation noise is to apply some form of spatial smoothing to velocity and topography data prior to application of the momentum balance. The smoothing operation generally requires proper selection of a window size that is tightly controlled by the type of operation or filter applied, as well as the form of the momentum balance used to infer drag (Brinkerhoff & Johnson, 2015; Kamb & Echelmeyer, 1986; McCormack et al., 2019). Additional checks on physical consistency of the smoothed variables are also necessary, such as the sign of the inferred drag since drag is expected to resist ice flow (never drive flow) such that drag is negative by the convention in Equation 6. However, integration of such physical consistency checks into the smoothing operations is not straightforward and likely requires some form of constrained optimization (e.g., Morlighem et al., 2011; Brinkerhoff & Johnson, 2015) or inverse methods applied to an ice flow model.

Since the primary goal of this study is to infer physically consistent, time-dependent basal drag, we address the challenges of observation noise by augmenting the simple MSE loss function with physics-based loss functions that encode prior knowledge and any constraints on the drag. To that end, we first project the basal drag $\hat{\boldsymbol{\tau}}_b$ to the along-flow direction using the predicted velocity $\hat{\boldsymbol{u}}$:

$$\widehat{\tau_b} = \hat{\boldsymbol{\tau}}_b \cdot \frac{\hat{\boldsymbol{u}}}{\|\hat{\boldsymbol{u}}\|}. \tag{9}$$

We then construct loss functions penalizing the spatial roughness and sign of the predicted along-flow drag:

$$\mathcal{L}_{ph}\left(\boldsymbol{\theta}\right) = \frac{1}{P}\sum_{k=1}^{P}\left[\lambda \cdot \left(\frac{\partial^2 \hat{\tau}_b^k}{\partial x^{k^2}} + \frac{\partial^2 \hat{\tau}_b^k}{\partial y^{k^2}}\right)^2 + \alpha \cdot \text{ReLU}\left(\hat{\tau}_b^k\right)\right], \tag{10}$$

where the first term implements Laplacian smoothing, the second term penalizes positive basal drag via the rectified linear unit (ReLU) function (ReLU$(x) = \max(x,0)$), λ and α are scalars controlling the relative strengths of the losses, and P is the number of examples used evaluating these losses. Choice of the penalty parameters λ and α will generally be controlled by data quality (noise level, spatiotemporal coverage, accuracy of bed topography, etc.) and a priori uncertainties on the parameterization of ice flow, such as uncertainty in ice rheology variations or spatial smoothness of drag. Here, we set $\alpha = 1$ for all cases presented in this work and allow λ to be selected using standard model selection techniques like cross-validation or an L-curve (Figure S5).

The smoothness penalty is analogous to a wide variety of regularization techniques used in inverse methods to encourage a preferred distribution of model parameters (e.g., Gillet-Chaulet et al., 2016; Habermann

et al., 2013; MacAyeal, 1993; Morlighem et al., 2010). Here, we use Laplacian smoothing (second-order gradients) to explicitly encourage a smooth spatial distribution of inferred basal drag. Another commonly used metric is the squared-norm of the first-order gradients of $\hat{\tau}_b$ (e.g., Gillet-Chaulet et al., 2016), which we found to give qualitatively similar results as the Laplacian for most cases, but in certain cases we found the former tended to encourage sharper spatial variations in drag, especially when observational noise is higher. We thus use the Laplacian smoothing penalty in this work but note that reconstruction of sharper variations in drag may be a desired feature in certain scenarios, in which case we would simply replace the second-order derivatives with first-order derivatives. For the sign penalty, the ReLU function allows for penalization of positive drag values scaled by their magnitude. This approach effectively casts the sign penalty as an asymmetric shrinkage function that encourages drag to be closer to zero (from the positive direction) but does not explicitly enforce the sign. We found this approach to result in more stable training convergence by allowing for sufficient updates to the neural network weights in the early stages of training, that is, by allowing the optimizer to temporarily permit velocity predictions that lead to drag with the wrong sign. More generally, we also found that permitting small positive drag values provided a convenient way to diagnose areas where the SSA is poorly suited or bed slopes were locally steeper. The ReLU-based sign penalty used here is similar in principle to a Lagrangian relaxation of a constrained optimization problem with a constant multiplier α but is more compatible with the first-order gradient-based optimizers commonly available with deep learning frameworks.

An important feature of the physics-based loss functions is that the number of examples P for $\mathcal{L}_{ph}(\boldsymbol{\theta})$ is not necessarily equal to the number of examples M used for $\mathcal{L}_{mse}(\boldsymbol{\theta})$. For the latter, the M examples are determined by the number and location of ice surface observations, whereas the spatial and temporal coordinates for the P examples can be specified anywhere within the training domain. In other words, the data for computing the physics-based loss are synthetically generated by the neural network $f_{\boldsymbol{\theta}}$ during training. This feature is a well-known benefit of physics-informed neural networks (PINNs) in that even in the small-data regime, the additional physics-based penalties can sufficiently prevent overfitting of the data by large neural network representations for $f_{\boldsymbol{\theta}}$ by allowing for generation of additional synthetic training data (Raissi et al., 2019). In our case, we specify a set of space and time coordinates $\boldsymbol{x}$ and t that are randomly distributed within the training domain and are independent of the coordinates corresponding to the observations. At these coordinates, we use $f_{\boldsymbol{\theta}}$ to generate predictions of $\hat{\boldsymbol{u}}$ and $\hat{h}$ and their spatial gradients in order to then predict $\hat{\tau}_b$ for computing the losses in Equation 10.

While inferred, time-dependent values of basal drag are the primary outputs of the learning framework, equally important are estimates of uncertainties associated with those drag values. Drag uncertainty can stem from observation noise and epistemic uncertainty derived from an uncertain momentum balance and ice rheology. As discussed previously, uncertainty stemming from observation noise can itself be partitioned into measurement noise (noise intrinsic to the data source) and noise of spatial gradients due to incorrect smoothing. We adopt a simple strategy of reformulating $f_{\boldsymbol{\theta}}$ to output standard deviations for the predictions $\hat{\boldsymbol{u}}$ and $\hat{h}$ in addition to their mean values. These outputs are then used to parameterize Gaussian probability distributions (independent for $\hat{h}$ and each component in $\hat{\boldsymbol{u}}$) which can be used to replace the MSE loss function in Equation 8 with a negative log-likelihood function:

$$\mathcal{L}_{nll}(\boldsymbol{\theta}) = \frac{1}{M} \sum_{k=1}^{M} \left[-\log p_{\hat{\boldsymbol{u}}^k}(\boldsymbol{u}^k) - \log p_{\hat{h}^k}(h^k) \right], \tag{11}$$

where $p_{\hat{\boldsymbol{u}}}$ and $p_{\hat{h}}$ are the likelihood functions for $\hat{\boldsymbol{u}}$ and $\hat{h}$, respectively. A more complex probability distribution for the likelihoods, for example, multivariate Gaussians, may more accurately model dependencies between surface variables and could potentially capture epistemic uncertainties by incorporating model uncertainties in the covariance matrix (e.g., Duputel et al., 2014), although strong non-linearities in all but the most simple ice flow models would likely limit the utility of Gaussian-based error models. Since multivariate distributions would introduce more computational complexity for neural network training, we use the simpler, independent Gaussian likelihoods here to estimate the variances of the surface variables. Uncertainties for the time-dependent drag, $\hat{\tau}_b$, can then be estimated via Monte Carlo sampling by passing independent samples of the surface variables through the SSA force balance. We note that in this formulation, estimates of the propagated uncertainties for basal drag will mark the upper bound since multivariate
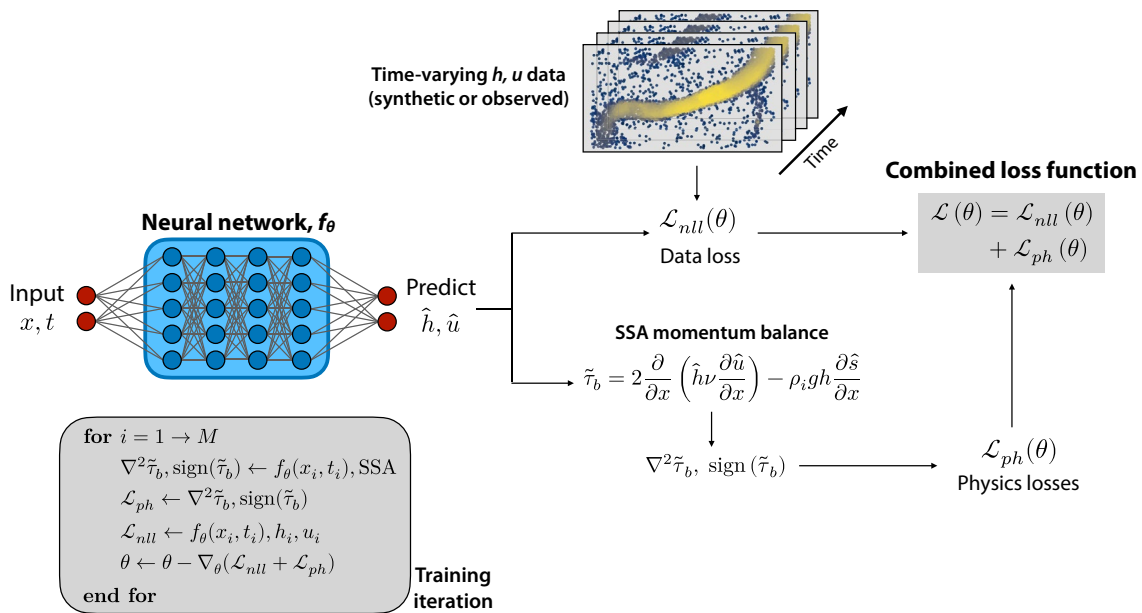
Figure 1. Diagram of neural network architecture and construction of loss function for training. Scattered spatial and time coordinates are input into network f_θ, which generates predictions of ice surface velocity and thickness, $\hat{u}$ and $\hat{h}$, at those coordinates. The misfit between these predictions and time-dependent velocity and thickness data (observed or generated from an ice flow simulator) is used to compute the loss function $\mathcal{L}_{nll}$. Simultaneously, the velocity and thickness predictions by f_θ (potentially at an independent set of space and time coordinates) are used to estimate basal drag from ice flow momentum balance equations (1D momentum balance used here as an example). Measures of spatial roughness and sign inconsistency for basal drag are then used to compute the loss function $\mathcal{L}_{ph}$. The optimal values for the neural network weights θ are obtained by minimizing the sum of the two loss functions through a gradient descent procedure. A simplified algorithm for a single training iteration which processes M training examples is shown in the gray inset.

Gaussians with finite mean and (independent) variances are known to maximize information entropy in the absence of covariance information (Cover & Thomas, 1999).

The final learning objective is to estimate $\hat{\theta}$ that minimizes the following joint loss function incorporating both data and physics-based loss functions (Figure 1):

$$\mathcal{L}(\theta) = \mathcal{L}_{nll}(\theta) + \mathcal{L}_{ph}(\theta). \tag{12}$$

A key point to reiterate is that for the joint loss function, observed data are only used for the loss function $\mathcal{L}_{nll}$ while evaluation points for the physics-based loss function $\mathcal{L}_{ph}$ can be evaluated anywhere within the training domain (see Appendix B for further details on training and neural network architecture).

2.2.1. Comparison With Control Methods for Ice Flow Models

The learning framework applied to time-dependent ice surface velocity and elevation data effectively forms a physics-aware space-time interpolator of the data. The interpolation kernel is provided by the hypersurface learned by the neural network, and the physical constraints are encoded in the loss functions specifying our prior assumptions on the characteristics of the underlying basal drag field. As such, while this approach can be viewed as an analog to a time-dependent inversion of basal drag using control methods applied to an ice flow model (e.g., Goldberg et al., 2015), there are several key differences. Forward runs of transient ice flow models generally require specification of key boundary conditions regarding surface mass balance, grounding line stresses and migration, ice velocities at inflow boundaries, and a functional form for the basal drag (e.g., power-law form in Equation 2). Each of these boundary conditions are time-varying and subject to varying degrees of uncertainties, which can require a significant number of spin-up runs and fine-tuning of model parameters in order to generate velocity and elevation fields that match the observations (Larour et al., 2014). By directly having access to the time-varying surface elevation and velocities from observations, we eliminate the need for evaluation of a forward model (and required boundary conditions) and simply rearrange the SSA momentum balance. We thus decouple the time-evolution of the basal drag from other processes that can influence surface elevations, for example, surface mass balance (Larour et al., 2014), yet we also enforce that the inferred basal drag is fully consistent with the predicted velocity

Figure 2. Experimental setup and initial ice geometry for 1D simulations of a marine-terminating glacier. Simulation outputs are used as synthetic data for testing basal drag recovery with neural network model. (a) Initial grounded ice (light blue) slides on its bed (dark brown) below sea level (dark blue; dashed black line). For perturbation experiments, ocean level $S(t)$ is periodic in time. X-coordinates indicate distance upstream from terminus. The ice surface, $s(X, t)$, varies in space and time while the bed, $b(X)$, varies only in space. Spatial distributions of prefactor values, c_b (b) and exponents, $1 / m$ (c) for two different simulation scenarios with different sliding modes but similar levels of basal stress. Case I (blue lines) corresponds to a constant exponent and slowly varying prefactor while Case II (red lines) corresponds to periodically varying exponent and prefactor profiles. The values for Case II are chosen such that the steady-state basal drag values are roughly equal to the steady-state drag for Case I.

and elevations. Computationally, the neural network model is mesh-free and can be evaluated anywhere within the training domain, which avoids the potentially costly parameterization of basal drag at each time step for each finite element (Goldberg et al., 2015). While recent work has utilized Ensemble Kalman Filters to assimilate surface data in a sequential manner in order to more efficiently infer time-dependent basal drag (Gillet-Chaulet, 2020), the requirements for specification of boundary conditions and grounding line migration still persist. Overall, our more focused objective of reconstructing the time-dependent basal drag allows us to bypass several algorithmic requirements necessary for forward runs of transient ice flow models, which are still necessary for any prognostic evaluation of future ice states.

3. Validation on 1D Ice Flow Simulations

To evaluate inference of basal drag using the neural network model, we apply the learning framework to various sets of synthetic ice velocity and thickness data generated by 1D SSA simulations for both spatially and temporally varying frictional parameters. These simulations are designed to generate observations similar to those observed at real glaciers subject to time-varying stress conditions. The prescription of spatially and temporally varying frictional parameters in the simulator provides mathematically convenient scenarios for testing recovery of the underlying sliding law parameters by our learning framework.

3.1. Spatially Varying Drag

We first generate 1D SSA simulations for two different cases of frictional parameter spatial distributions (Figure 2).In the first case, we prescribe a constant exponent of $m = 3$ and a spatially varying prefactor, c_b, with values that slowly increase with upstream distance to approximate increasing basal drag. In the second case, we prescribe periodic exponent values with values ranging from approximately 1 to 6, which spans the regimes from linear to approximately plastic (rate-independent) sliding. Additionally, we assign values of c_b such that the modeled basal drag is approximately equal to the drag from the first case. In this way, both cases will have similar values of basal drag throughout the simulation but different time-dependent sliding, providing a good test for the recovery of the true frictional parameters from time-dependent observations. For both simulations, we prescribe an ice temperature of $-5°$ C and compute the flow rate parameter using an Arrhenius relation (Cuffey & Paterson, 2010), and we use ice and seawater densities of 917 and 1,024 kg/m^3, respectively. After a spin-up period of 200 years, we force both simulations with periodic variations of the longitudinal stress at the terminus to simulate periodic ocean tides over a five-year duration. We allow the terminus to migrate based on a flotation criterion (e.g., Nick et al., 2009). The resulting velocity

time series show strong periodicity in time while the ice thicknesses and elevations are roughly constant throughout the five-year duration (Figure S1).

For training the network f_{θ} subject to the loss function in Equation 12, we select a spatial subset spanning the minimum terminus position and 50 km upstream of that position to use as training data. We restrict all subsequent analysis of the inference results to this spatial subset. To simulate measurement noise, we add white noise with a standard deviation of 10 m/year and spatially correlated noise generated from a squared exponential covariance function with a lengthscale of 5 km (equivalent to approximately 10 ice thicknesses) and an amplitude of 6 m/yr, resulting in noise levels that are approximately 10% of the mean velocity variation, similar to noise in remote sensing observations (Minchew et al., 2017; Mouginot et al., 2019). Generally, the correlated noise will have a much larger effect on the inferred basal drag since coherent velocity gradients will be mapped to spurious basal drag variations, particularly when the wavelength of the noise is on the order of several ice thicknesses. For the ice thickness data, we similarly add white noise with an amplitude of 10 meters and correlated noise with the same lengthscale and an amplitude of 6 meters, which again results in noise levels that are approximately 10% of the mean thickness. This noise level is similar to bed topography uncertainties estimated using mass conservation techniques for fast-flowing glaciers (Morlighem et al., 2017). Finally, we uniformly sample 50,000 data points within the space-time training volume (about 30% of the simulation outputs) for computing the data loss function, $\mathcal{L}_{nll}$, and we uniformly sample an additional 50,000 space-time coordinates for computing the physics losses, $\mathcal{L}_{ph}$.

After training f_{θ}, we perform Monte Carlo sampling of the learned time-dependent u and h fields in order to generate time-dependent samples of τ_b throughout the entire training domain using the SSA momentum balance. In this manner, we can visually examine the τ_b versus u relationship to infer the underlying physical relationship without needing a closed-form symbolic expression of that relationship. A natural domain for viewing τ_b versus u is in log space where, for the power-law form of the sliding law, the slope corresponds to $1/m$ and the intercept corresponds to $\log(c_b)$. For the forcing environment simulated here (longitudinal stress perturbations applied at the terminus), the sliding parameters are expected to be time-invariant and can be estimated for each point along the glacier. Additionally, since we generate samples of τ_b for any triplet of (u, x, t), we can compute the mean and standard deviation of the $1/m$ and c_b estimates. As a comparison, we also compute basal drag directly from the noisy surface observations using Equation 6 where we first spatially smoothe the u and h fields with a low-pass Butterworth filter with a cutoff period of approximately six ice thicknesses (commensurate with the expected stress coupling lengthscale for glaciers on colder ice sheets (Cuffey & Paterson, 2010)), and then we compute the momentum balance using finite differences for the spatial gradients. This combined approach of smoothing and finite differencing is commonly used when computing the force balance from spatially continuous surface observations.

For both simulation cases, we are able to accurately recover the true sliding law parameters for the region of the glacier with sufficiently large velocity variations (within ≈35 km of the terminus; Figure 3).As the upstream distance increases, the amplitudes of the velocity fluctuations caused by the stress perturbations applied at the terminus attenuate, which ultimately results in increasing uncertainties in both $1/m$ and c_b as the linear fits in log space become more ill-conditioned. For the case where the prescribed prefactor and exponent both vary spatially, we observe larger uncertainties on c_b where $1/m$ is lower (Figure 3), which implies that sliding regimes that are closer to plastic will generally lead to more uncertain prefactors using the log-domain line fit used here.

Another key result is that estimates of basal drag using direct application of Equation 6 lead to highly biased and noisy sliding law parameter estimates, even when significant spatial smoothing is applied to the data prior to application of the momentum balance (Figures 3c–3f). The ability of the neural network framework to accurately recover the true parameter values (to the extent where velocity variations are large enough) indicates similar levels of robustness to noise as traditional inverse modeling schemes that apply some form of regularization on the modeled basal drag (e.g., Habermann et al., 2012; Larour et al., 2012). The ability to quantify uncertainties in predictions of basal drag is an important additional benefit of the probabilistic loss functions used to train the network f_{θ}. The benefit of basal drag uncertainties has been previously demonstrated with Markov Chain Monte Carlo (MCMC) techniques used to sample the posterior distribution of time-invariant basal drag friction coefficients (Petra et al., 2014). There, it was found that certain statistical metrics of the posterior distribution (e.g., eigenvectors of the covariance matrix) could
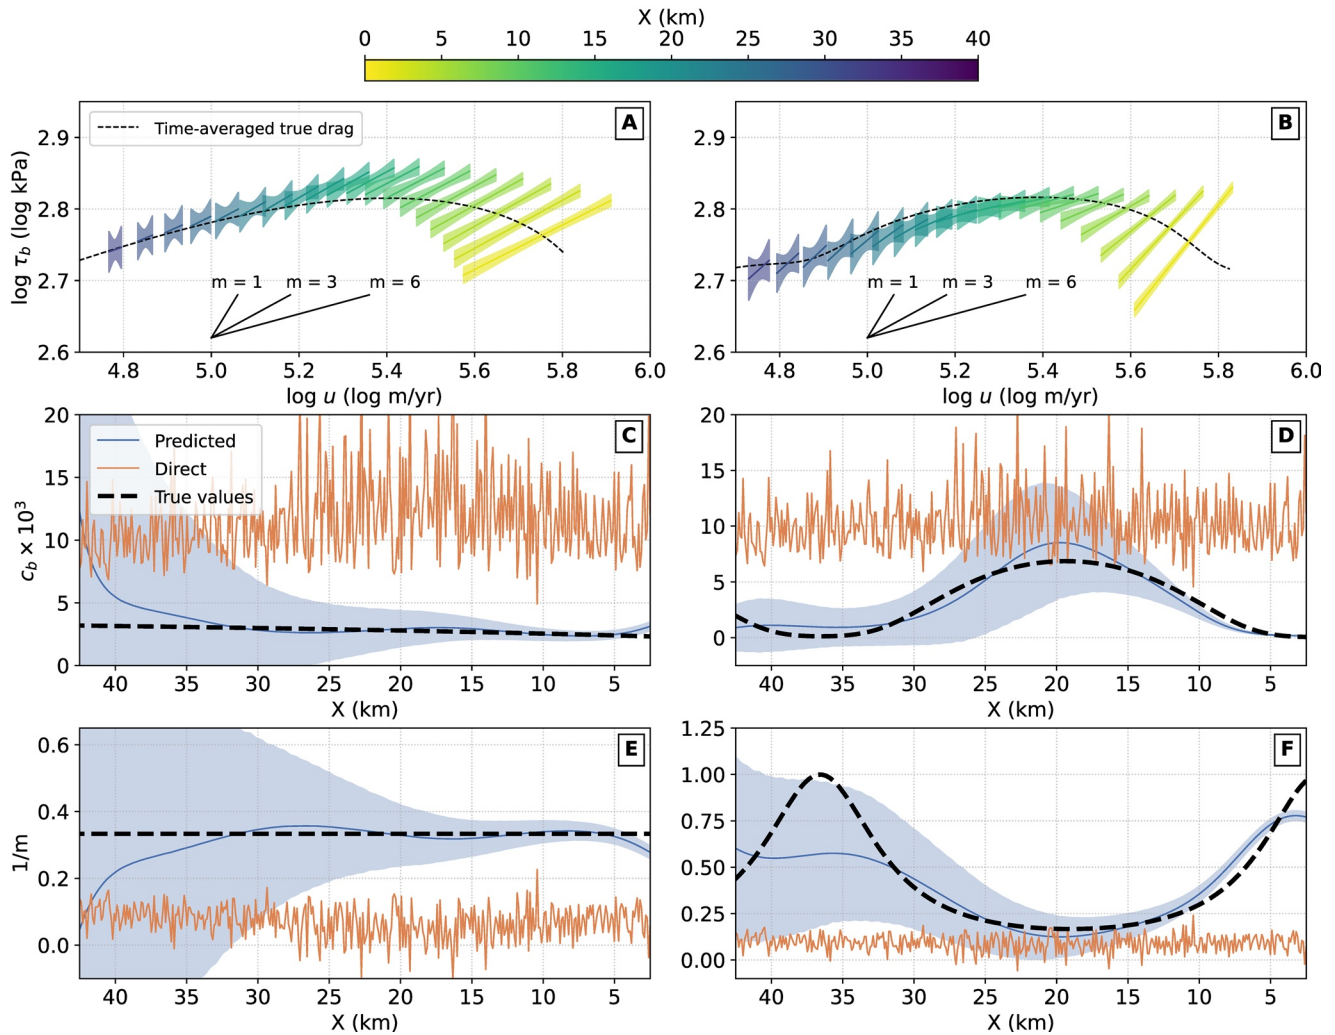
Figure 3. Inferred basal parameter profiles and predicted time-dependent basal drag τ_b versus sliding velocity u for 1D simulations. Left column corresponds to results for the simulations with spatially varying prefactor (case I) while the right column correspond to results for simulations with spatially varying prefactor and exponent (case II). (a) Log-domain plot of basal drag versus sliding velocity for case I. Colors indicate distance upstream of the terminus, where dark lines correspond to a linear fit of stochastic samples and shaded regions correspond to sample standard deviation (3σ). Thin black dashed line indicates the time-averaged, noise-free basal drag from the simulations, and solid black lines at the bottom illustrate theoretical slopes for $m = 1$, 3, and 6. (b) Same as (a) but for case II. (c) and (e) Profiles of estimated prefactor, c_b, and exponent, $1/m$, respectively, for case I. Blue lines correspond to the neural network predictions while blue shaded areas are the prediction uncertainties (3σ). Red lines correspond to direct estimates of basal drag using momentum balance of surface observations. Thick black dashed lines correspond to the true values. (d) and (f) same as (c) and (e) for case II. Sliding law parameter inference is best constrained where velocity variations are high and areas away from edges of training domain.

discriminate between areas where observations provided robust constraints on parameter values versus those where parameter uncertainties were primarily dictated by prior assumptions. While our approach to uncertainty quantification differs from these model-based inference methods (see Section 5.4 for a detailed discussion), we share the common goal of locating areas where basal drag and observations are well-constrained relative to other areas.

3.2. Time Varying Drag

In the previous subsection, the sliding law parameters were simulated to be time-invariant. However, for some glaciers and ice streams, basal drag has been hypothesized to evolve in time, for example, in response to changes in the subglacial hydrological system. As water flows into and out of the hydrological system, the basal water pressure compensates some of the overburden pressure and thus changes the effective pressure

(the difference between overburden and water pressures) at the bed. The overall change in effective pressure will thus affect the magnitude of the basal drag and the corresponding flow of ice (Flowers, 2015; Hewitt, 2013; Iken & Bindschadler, 1986; Minchew et al., 2016; Rosier et al., 2015; Schoof, 2010; Stevens et al., 2018). Here, we implement a simplified model for temporally varying water pressure by representing the prefactor in the power-law sliding law as a Mohr-Coulomb yield criteria (e.g., Tulaczyk et al., 2000) such that

$$c_b\left(x, t\right) = \mu\left(\rho_i g h\left(x, t\right) - p_w\left(x, t\right)\right),\tag{13}$$

where μ is a constant friction coefficient (which is generally a function of the internal friction angle ϕ), and the function $p_w\left(x, t\right)$ represents spatially and temporally varying basal water pressure. We model the basal pressure as a periodic pressure wave that propagates upstream in the following manner:

$$p_w\left(x, t\right) = \overline{p}_w + \hat{p}_w e^{-x/L} \cos\left(-\omega t + \frac{\omega}{v_p} x\right),\tag{14}$$

where $\overline{p}_w$ is a constant water pressure, $\hat{p}_w$ is the amplitude of the pressure oscillations, L is a decay lengthscale (allowing for upstream exponential decay of the pressure perturbation), ω is the angular frequency of the oscillation, and v_p is the wave speed (phase velocity) controlling the upstream propagation speed of the pressure wave. This functional form for the pressure wave approximates diffusive models for subglacial hydrology (Rosier et al., 2015) where the phase velocity and decay lengthscale of the pressure wave is controlled by the conductivity of the hydraulic system. Using the same ice geometry as the previous two 1D simulations, we set values of $\mu = 2 \times 10^{-3}$, $\overline{p}_w = 1000\,\text{kPa}$, $\hat{p}_w = 500\,\text{kPa}$, $L = 45\,\text{km}$, $\omega = 1\,\text{rad/year}$, and $v_p = 0.27\,\text{km/day}$. Furthermore, to investigate the dynamics of a plastically deforming bed subject to hydrological variations, we generate simulations for $m = 3$ and 10, where the latter approximates a plastic bed. In order to maintain a similar velocity range for the two sliding exponents, we increase the friction coefficient for $m = 10$ to $\mu = 5 \times 10^{-3}$ (in order to match the secular velocities for the $m = 3$ outputs) and reduce the pressure wave amplitude to $\hat{p}_w = 200\,\text{kPa}$ since a plastic bed will result in large velocity variations for a given stress perturbation. The resulting velocity fields for both simulations show similar annual variations as the simulations forced by periodic variations at the terminus (Figures 4a and 4b). While the velocity variation amplitudes are similar at the terminus for $m = 3$ and 10, the upstream extent of the variations is larger for the latter case, even with a substantially reduced pressure wave amplitude.

Using the same training procedure for f_θ as the previous two cases (but without noise added to the data in order to highlight the mechanical effects of the pressure wave), we reconstruct the full time history of the modeled basal drag. The stochastic predictions for $\hat{\tau}_b$ demonstrate that, similar to the previous experiments, drag variations are well constrained in the regions where velocity variations are higher (Figures 4c and 4d). Generally, decreases in drag are associated with increases in velocity since the propagating pressure wave is the primary driver of speedups in ice flow. However, this trend changes with upstream distance as the pressure wave amplitude decays and longitudinal stress perturbations become the dominant forcing mechanism. The crossover point at which longitudinal stresses become more important is controlled by the pressure wave decay lengthscale, phase velocity, and sliding law exponent. To illustrate this point further, we fit a temporal function consisting of a linear trend and an annual sinusoid to the $\hat{\tau}_b$ time series at each point while accounting for the uncertainties in $\hat{\tau}_b$. The amplitude and phase delay of the sinusoids, as well as their formal uncertainties, can then be estimated along the glacier (Figures 4c and 4d). For $m = 3$, we observe a significant phase offset between the first 10 km upstream of the terminus and the rest of the ice stream; this phase offset is minimized for $m = 10$. For the sinusoidal amplitudes, recall that the imposed water pressure variations for $m = 3$ were 2.5x larger than those for $m = 10$. However, the recovered maximum drag amplitude is only twice as large for $m = 3$ compared to $m = 10$, which indicates a negative feedback between the reduction in drag from the pressure wave and an increase in drag resulting from the induced speedup (Rosier et al., 2015). This negative feedback also manifests as a sharper drop-off in amplitude with upstream distance for $m = 3$. The amplitude reaches a local minimum at the same location where the drag phase gradients are at their peak. For $m = 10$, as with the phase gradients, the amplitude drop-off is much less pronounced.

Another important difference in these pressure wave-driven simulations is that the relationship between τ_b versus u in the log domain exhibits a cycle (Figures 4e and 4f), with elliptical behavior arising from varying
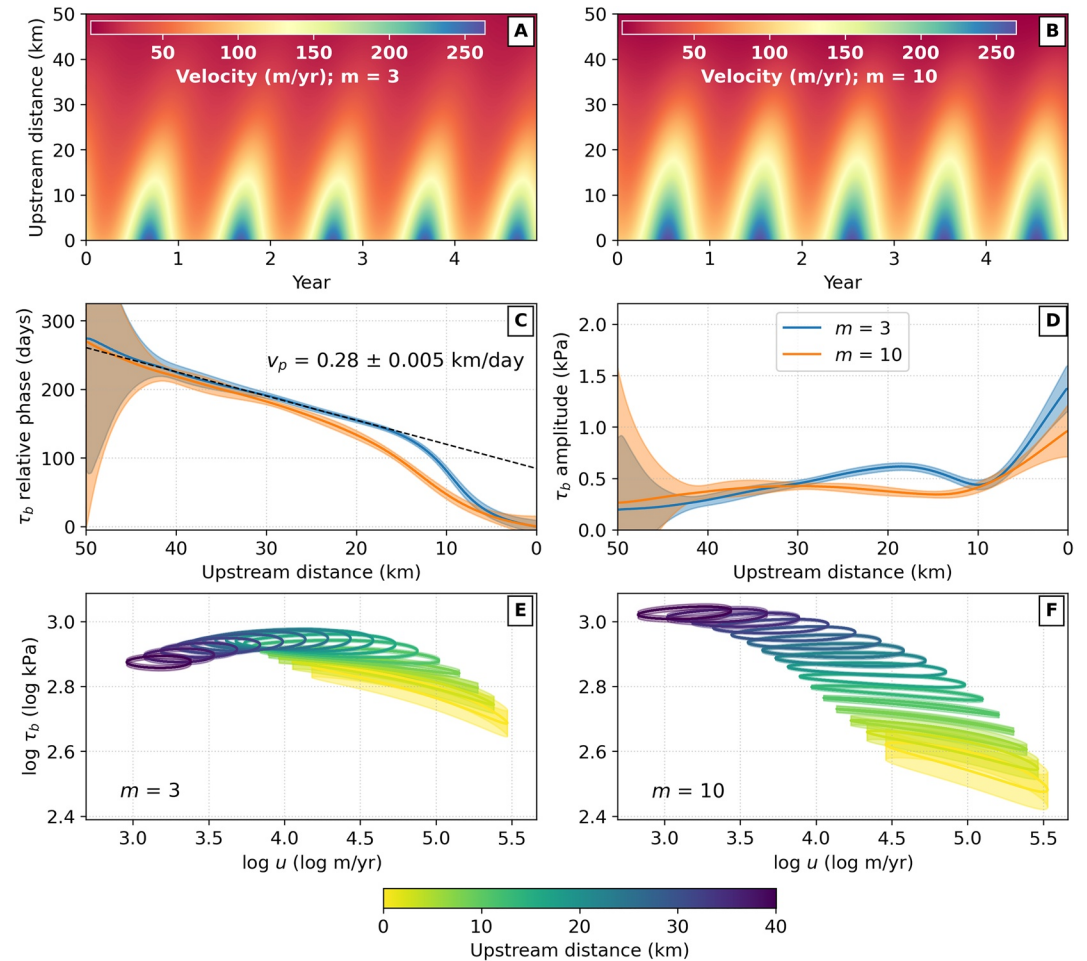
Figure 4. Neural network velocity and basal drag predictions for temporally varying drag simulation. Simulations were conducted with spatially uniform sliding law exponent values of $m = 3$ and $m = 10$. (a) and (b) show spacetime evolution of predicted velocity for $m = 3$ and $m = 10$, respectively. (c) Profiles of estimated basal drag periodic phase delay where blue and orange lines correspond to $m = 3$ and $m = 10$, respectively. Black dashed line corresponds to best fit line to the phase delay for $m = 3$, indicating a phase velocity of 0.28 ± 0.005 km/day. (d) Profiles of estimated basal drag periodic amplitude. (e) Log drag. versus log velocity for select points for $m = 3$ where color indicates distance upstream from the grounding line. Solid lines correspond to mean drag predictions while shaded regions correspond to 3σ uncertainties. (f) Same as (e) but for $m = 10$. The amplitude and phase delay profiles combined with the ellipticity of $\log \tau_b$ versus $\log u$ can be used to infer the propagating pressure wave.

levels of phase lag between the periodic velocity and basal drag signals (Figure S2). The varying phase lag is again a consequence of competing basal drag perturbations from the pressure wave and balancing of longitudinal stress perturbations resulting from the initial speedup where the latter generally propagates upstream with a higher phase velocity.

The similarity in the velocity variations between the pressure wave-forced simulations in this section and the terminus-forced simulations in the previous section obscures the stark differences in the basal drag evolution between the two model classes. While the neural network-based drag reconstruction is well-constrained for both cases, attribution of the dominant forcing mechanism for a given glacier without a priori information is considerably more uncertain. Nevertheless, the strong inverse proportionality between $\hat{\tau}_b$ and u in Figures 4e and 4f for regions closer to the terminus does suggest that glaciers and ice streams exhibiting similar cycles are likely influenced by time-varying effective pressure. Considering that effective pressure changes can be subsumed into a time-varying sliding law prefactor, simultaneous recovery of both the sliding law prefactor and exponent as shown in the previous section is not possible for glaciers influenced by substantial subglacial hydrological effects. In these cases, independent observations of basal water pressure

variations (and thus, prefactor variations) or explicit modeling of subglacial hydrology would be needed to recover values of the exponent. Conversely, if a priori information about the exponent were available (e.g., the bed is well-approximated by plastic deformation), then it is possible to derive estimates of basal water pressure variations from the time-varying drag (Minchew et al., 2016).

4. Rutford Ice Stream, Antarctica

Rutford Ice Stream (RIS) in West Antarctica is a fast-flowing ice stream which flows into the Filchner-Ronne Ice Shelf (Figure 5a). RIS is laterally confined by the Ellsworth Mountains and Fletcher Promontory with an average width of approximately 23 km (Figure 5b), and most (> 95%) of the forward velocity is due to basal sliding (De Rydt et al., 2013; Gudmundsson, 2007; Joughin et al., 2006; Smith et al., 2015). The high width-to-thickness and slip ratios support the use of the SSA approximations for examining basal drag variations at intermediate to long spatial wavelengths (De Rydt et al., 2013; Gudmundsson, 2003). Furthermore, RIS exhibits strong variations in flow velocity due to tidal forcing where non-zero variations are measured almost 100 km away from the grounding line. While variations in vertical velocity are mostly modulated by diurnal and semi-diurnal tides, along-flow variations are observed primarily at the fortnightly M_{sf} (14.77 days) period (Gudmundsson, 2006; Minchew et al., 2017; Murray et al., 2007), which indicates a non-linear response of RIS flow to tidal forcing (Gudmundsson, 2011; Rosier & Gudmundsson, 2016; Rosier et al., 2015). While several recent studies have compared different mechanisms for originating along-flow variations at the M_{sf} frequency on the ice shelf (e.g., Rosier & Gudmundsson, 2020; Robel et al., 2017), our focus in this study is on using the response of ice flow in the grounded ice stream to infer the mechanics of slip at the ice-bed interface. Thus, our analysis focuses on regions of the ice stream greater than 10 km upstream of the grounding line in order to avoid elastic effects due to bending stresses, which are not incorporated into the SSA approximations (Rosier & Gudmundsson, 2016). We emphasize that a rigorous exploration of the ice stream stress response (including the elastic response) to tidal forcing is outside the scope of this work. Rather, our aim is demonstrate the machine learning-based techniques for inferring time-dependent basal drag on high-quality surface observations.

4.1. Data and Learning Objectives

We use existing data sets to constrain the surface velocity fields and ice-stream geometry. The 3D surface velocity fields were derived from 9 months of synthetic aperture radar data collected from multiple viewing angles in order to constrain a parametric surface displacement model consisting of sinusoids corresponding to the primary tidal constituents and a steady-state velocity (Minchew et al., 2017). Since our main focus is on the along-flow velocity variations (where variations at diurnal and semi-diurnal constituents are minimal (Murray et al., 2007)), we use only the steady-state velocity and sinusoid periods at the M_{sf} frequency. Geometric information (surface elevation and ice thickness) were obtained from BedMachine V1 (Morlighem et al., 2020). While our analysis is focused on the regions of the ice stream greater than 10 km upstream of the grounding line, our model domain spans from 150 km upstream of the grounding line to regions of the ice shelf within 45 km downstream of the grounding line. With this extended domain, we can confidently constrain the spatial gradients of the observation variables. Additionally, inclusion of the ice shelf also provides a means to validate the rheological parameters since basal drag is expected to be negligible on the shelf (seawater offers very little resistance to ice flow). Here, we use a characteristic temperature of -10° C to calculate an effective depth averaged value of A from tabulated values (Cuffey & Paterson, 2010) and a stress exponent $n = 3$ to compute ice viscosity (Appendix A). A more rigorous treatment of A would include the effect of flow rate enhancement due to shear weakening, crystallographic fabric, impurities, etc. (Cuffey & Paterson, 2010), which can affect the inferred distribution of basal drag (Zhao et al., 2018). Since the focus of this work is to demonstrate the recovery of time-dependent basal drag in the central trunk of RIS, where rheological variations are expected to be minimal, we proceed with a uniform value for A.

We train the network f_θ to predict the time-varying 2D horizontal velocity components and time-invariant ice thickness and surface elevation. At the fortnightly timescales, ice thickness and driving stress are assumed to be constant in time. By adding the surface elevation variable to the outputs of f_θ (as opposed to adding a known bed elevation to the thickness predictions as was done with the simulated data), we implicitly account for errors in the bed topography by treating $s = h + b$ as an additional noisy observation
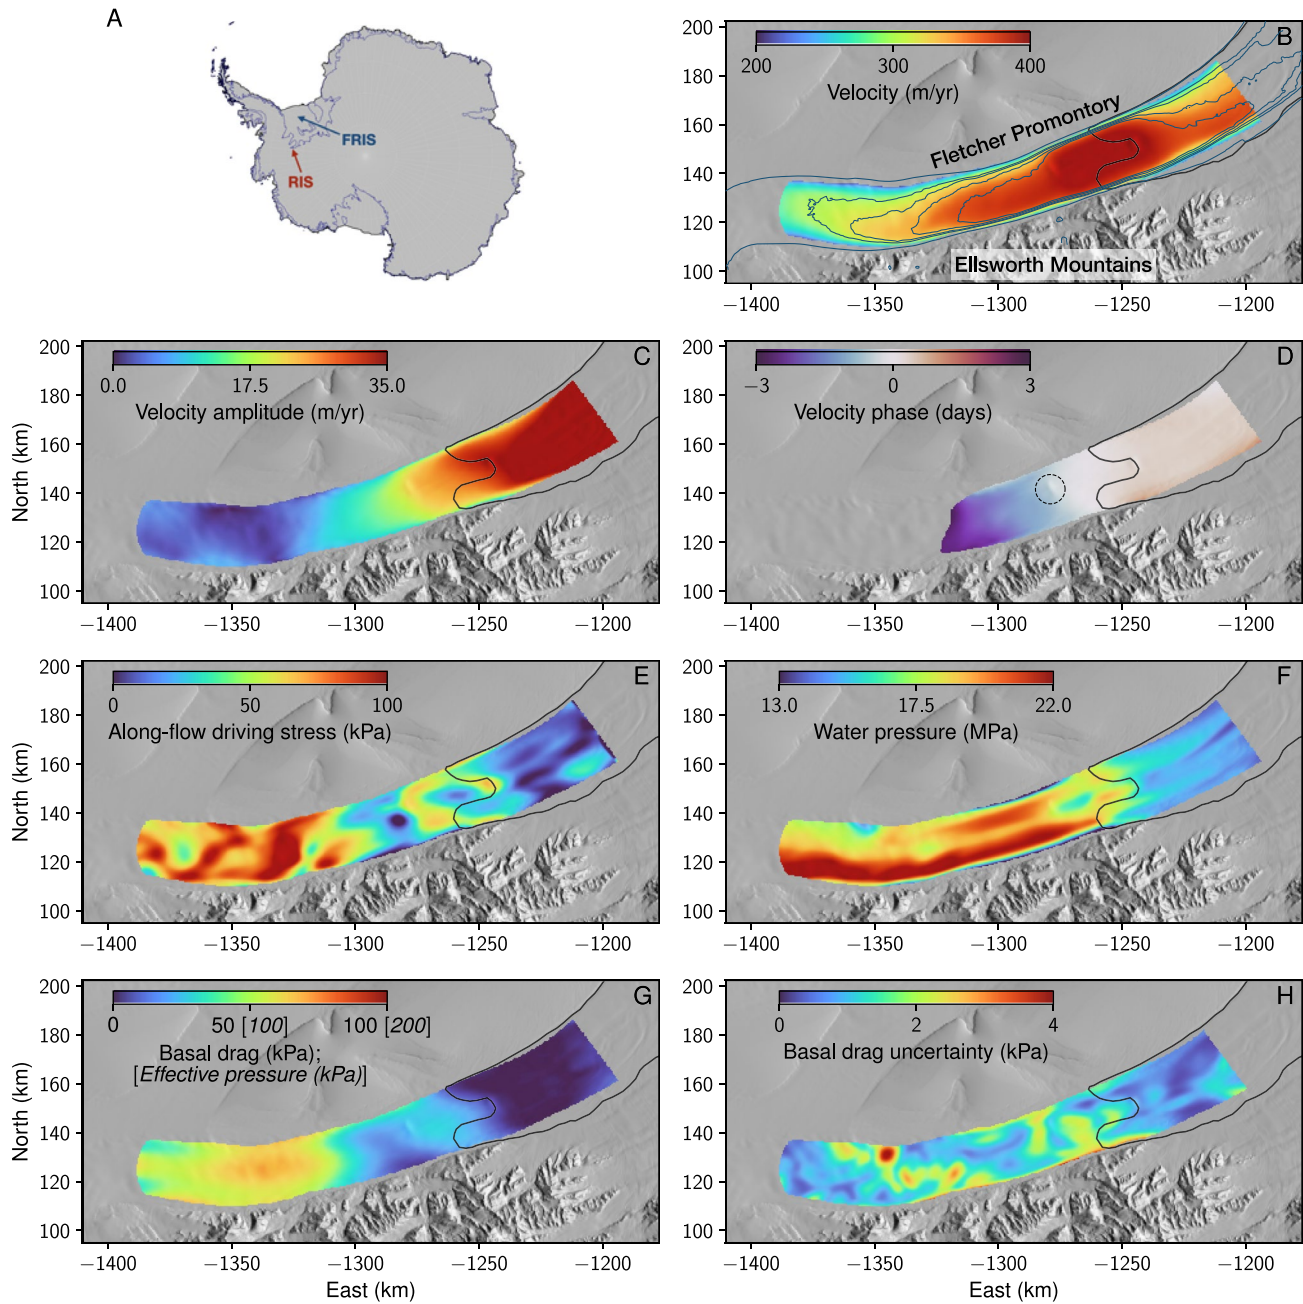
Figure 5. Rutford Ice Stream (RIS) study area and secular (time-averaged) surface velocity and basal drag. (a) RIS (red arrow) feeds into the Filchner-Ronne Ice Shelf, (blue arrow) in West Antarctica. (b) The neural-network-predicted secular velocity magnitude, which is in good agreement with the observed secular velocities from (Minchew et al., 2017) (blue contours at levels of 0, 100, 295, 320, 350, and 370 m/year). The predicted amplitude (c) and phase (d) of the time-dependent velocity variations, which are also in good agreement with (Minchew et al., 2017) upstream of the grounding line (areas of high phase uncertainty, due to low amplitude variations, masked out). Dashed black circle indicates prominent bump in bed topography. The driving stress (e) is mostly balanced by the neural-network-predicted basal drag (g). By assuming a plastic bed with yield stress determined by the Mohr-Coulomb yield criteria (Equation 13), effective pressure equals basal drag divided by the internal friction coefficient, $\mu = 0.5$ (pressure values shown in square brackets and italics in (G)). Secular water pressure (f) may then be derived from the effective pressure. Scalar uncertainties for predicted basal drag (h) are generally high in areas with relatively rapid changes in bed slope, such as the margins and near the grounding line.

subject to smoothing imposed by our physics-based loss functions. For the velocity components, rather than outputting the velocity values at any given input (x, t), we instead output the spatially varying coefficients of a periodic temporal model (independently for the u and v components), for example,:

$$u(x,t) = a(x)\cos(\omega_{sf}t) + b(x)\sin(\omega_{sf}t) + u_0(x), \tag{15}$$

where $\omega_{sf} = 2\pi/T_{sf}$ is the fortnightly angular frequency for the M_{sf} tidal constituent ($T_{sf} = 14.77$ days), and the coefficients $\left[a, b, u_0\right]$ vary in space only. This approach reduces the dimensionality of the neural network inputs to two spatial coordinates while providing physical constraints on the temporal form of the predictions. While the temporal model is not necessary since the neural network can learn the underlying time variability, we found that it substantially improved the convergence rate during training. Thus, for study areas where the time variability can be confidently constrained (e.g., for areas where daily GPS observations are available), inclusion of a suitable temporal model is likely beneficial.

For formulating the physics-based loss functions in $\mathcal{L}_{ph}$, we use the 2D SSA equations (Equations 1a and 1b) to predict the basal drag $\hat{\tau}_b$ in both spatial directions (east and north), which we then project to the along-flow direction using the predicted velocity vectors. This projection allows us to once again penalize the drag values with incorrect signs and to compute the Laplacian smoothness metric on a scalar field as opposed to a vector field. We assume an ice density $\rho_i = 917$ kg/m^3. Values for the hyperparameter controlling spatial smoothness of basal drag were chosen using a standard L-curve (Figure S5). After training f_θ, we Monte Carlo sample the time-dependent basal drag $\hat{\tau}_b$ in the along-flow direction (see Text S1 for validation on noise-free 2D ice flow simulations). As a post-processing step, we fit the predicted drag time series samples with the periodic model used for the velocity components (Equation 15) in order to reduce high-temporal-frequency drag variations. Additionally, uncertainties in the parameters of the final periodic model are formally computed by propagating uncertainties from the drag samples.

4.2. Secular Velocity and Basal Drag Predictions

The predicted along-flow secular (steady-state) velocity magnitudes for RIS are in good agreement with the observed secular velocities with 1σ errors of 5 m/year in the central trunk (Figure 5b), while the velocity amplitude and phase are also in good agreement with prior studies (Minchew et al., 2017) with 1σ errors of 1.7 m/year and 0.7 days, respectively. The steady-state basal drag magnitudes show a region of very low basal drag from approximately 10–50 km upstream of the grounded line, transitioning to higher drag over short distances (Figure 5g). This transition from a weak to a stronger bed has been inferred in several prior studies (e.g., Joughin et al., 2006; Pralong & Gudmundsson, 2011) and has been associated with a transition from dilatant to stiff sediment (Smith et al., 2015).Since RIS is close to steady-state (Gudmundsson & Jenkins, 2009), drag variations are mostly in balance with the driving stress (Figure 5e). The drag magnitudes in our training area peak at around 100 km upstream of the grounding line, which is colocated with a local high in the basal topography (Figure S6). Previous numerical studies of RIS have shown that basal topography is the dominant control on surface undulations, which in turn implies that basal topography is the dominant control on secular drag variations at the spatial scale of tens of kilometers (De Rydt et al., 2013; Pralong & Gudmundsson, 2011). Under the functional form of Equation 13 where μ represents the internal friction coefficient for till, these results support the view that variations in the friction coefficient μ are at much longer wavelengths (> 20 ice thicknesses), with the exception of the low basal drag region. By further assuming a plastic bed with a uniform value of $\mu = 0.5$ (median of published values (Iverson, 2010)), the effective pressure is simply twice the basal drag (Figure 5e), and we can obtain an estimate of basal water pressure at RIS by subtracting the effective pressure from the overburden stress (Figure 5f). We explore the implications of bed plasticity on water pressure changes in a later discussion.

Uncertainties for the predicted drag are generally highest at the margins and areas with high bed slopes where data gradients are large and work against the Laplacian smoothing penalty on the basal drag (Figure 5h). Specifically, the high slope areas exist near the grounding line, as well as near a prominent bump in the bed topography about 30 km upstream of the grounding line (Figure 5d). By examining the uncertainties for the predicted velocity, ice thickness, and surface elevation, we can see that all three variables exhibit larger uncertainties in these areas and contribute to the total drag uncertainty (Figure S7). Mathematically, the uncertainties here have been inflated due to larger data misfits during training; the neural network

learns to increase the likelihood variance in these high misfit bias areas in order to increase the total log likelihood (see Section 5.4 for further discussion). Overall, quantification of drag and grounding line migration in these areas has proven challenging and will only improve once more high-quality bed topography data are acquired (Rosier & Gudmundsson, 2020).

4.3. Time-Dependent Velocity and Basal Drag Predictions

By quantifying the change in velocity and basal drag at different times within the M_{sf} tidal period, we observe significant basal drag variations propagating upstream with values spanning 4–6 kPa over the course of the tidal period (Figure 6). Perhaps the most interesting observation is that the upstream propagation of positive velocity variations is associated with a propagating decrease in basal drag, which suggests some form of a pressure wave driven by subglacial hydrology (analogous to Figure 4). During the initial speedup of the ice stream, the associated basal drag decreases only slightly, which may signify destructive interference of basal drag reduction and longitudinal stress perturbations originating from loss of buttressing stresses downstream (Robel et al., 2017; Rosier & Gudmundsson, 2020). We reiterate that the inferred basal drag near the grounding line is likely inaccurate since we do not incorporate elasticity of the ice into our stress calculations and bed slopes there are subject to larger uncertainties. However, later in the tidal cycle when velocity speedups have propagated to about 70 km upstream of the grounding line, the basal drag decrease has also propagated upstream while becoming more widespread within the ice stream (Figures 6e and 6f). We do observe a phase lag between the velocity and drag variations which can be confirmed by the elliptical relationship between τ_b versus $\|\boldsymbol{u}\|$ (Figure 7), a characteristic we previously observed for the 1D simulated pressure waves. The exceptions to this behavior are near the grounding line and in the weaker bed where drag variations are minimal compared to the velocity variations. At greater upstream distances, we can observe a gradual transition in the ellipse orientation, signifying a transition to a stress regime where longitudinal stresses become the primary driver of the velocity variations.

5. Discussion

The availability of time-dependent observations of surface velocity and elevation permit direct estimates of time-varying basal drag that satisfies global stress balance. Coupled with a machine learning model for reconstructing the spatiotemporal function for basal drag, we can retrieve important sliding parameters under certain stress and loading conditions. We discuss the robustness and implications of these results below.

5.1. Inference of Sliding Law Parameters

Under the condition that ice surface velocity variations are driven by processes other than changes in drag at the bed–such as longitudinal stress perturbations at the terminus or grounding line, as may be expected in some cases for seasonal calving cycles, ocean tide effects via hydrostatic stress differences, or changes in buttressing stresses from ice shelves–then, for a general power-law formulation of the sliding law, it is possible to recover both the prefactor and exponent from a linear fit of time-dependent sliding velocity and basal drag predictions in the log domain. Parameter estimation for recently proposed augmented sliding laws that combine the power-law sliding relationship at lower velocities and a rate-independent (plastic) relationship at higher velocities (Joughin et al., 2019; Minchew & Joughin, 2020; Zoet & Iverson, 2020) can be accomplished in a similar manner through a nonlinear optimization in the log domain. Verification and refinement of such a law from remote sensing data sets would make significant progress towards unification of laboratory, observational, and theoretical approaches towards understanding glacier sliding dynamics.

For all proposed sliding law functional forms, if the sliding law parameters are known to be time-invariant (but may be spatially varying), then simultaneous parameter recovery is possible. In cases where the parameters may vary in time, such as when changes in the prefactor are driven by subglacial hydrological processes, then simultaneous recovery is not possible, and one would need additional information about the physical properties of the bed, such as water pressure variations or bed plasticity (which is equivalent to knowing the value of the exponent in the power-law form of the sliding law, as discussed in the following section).
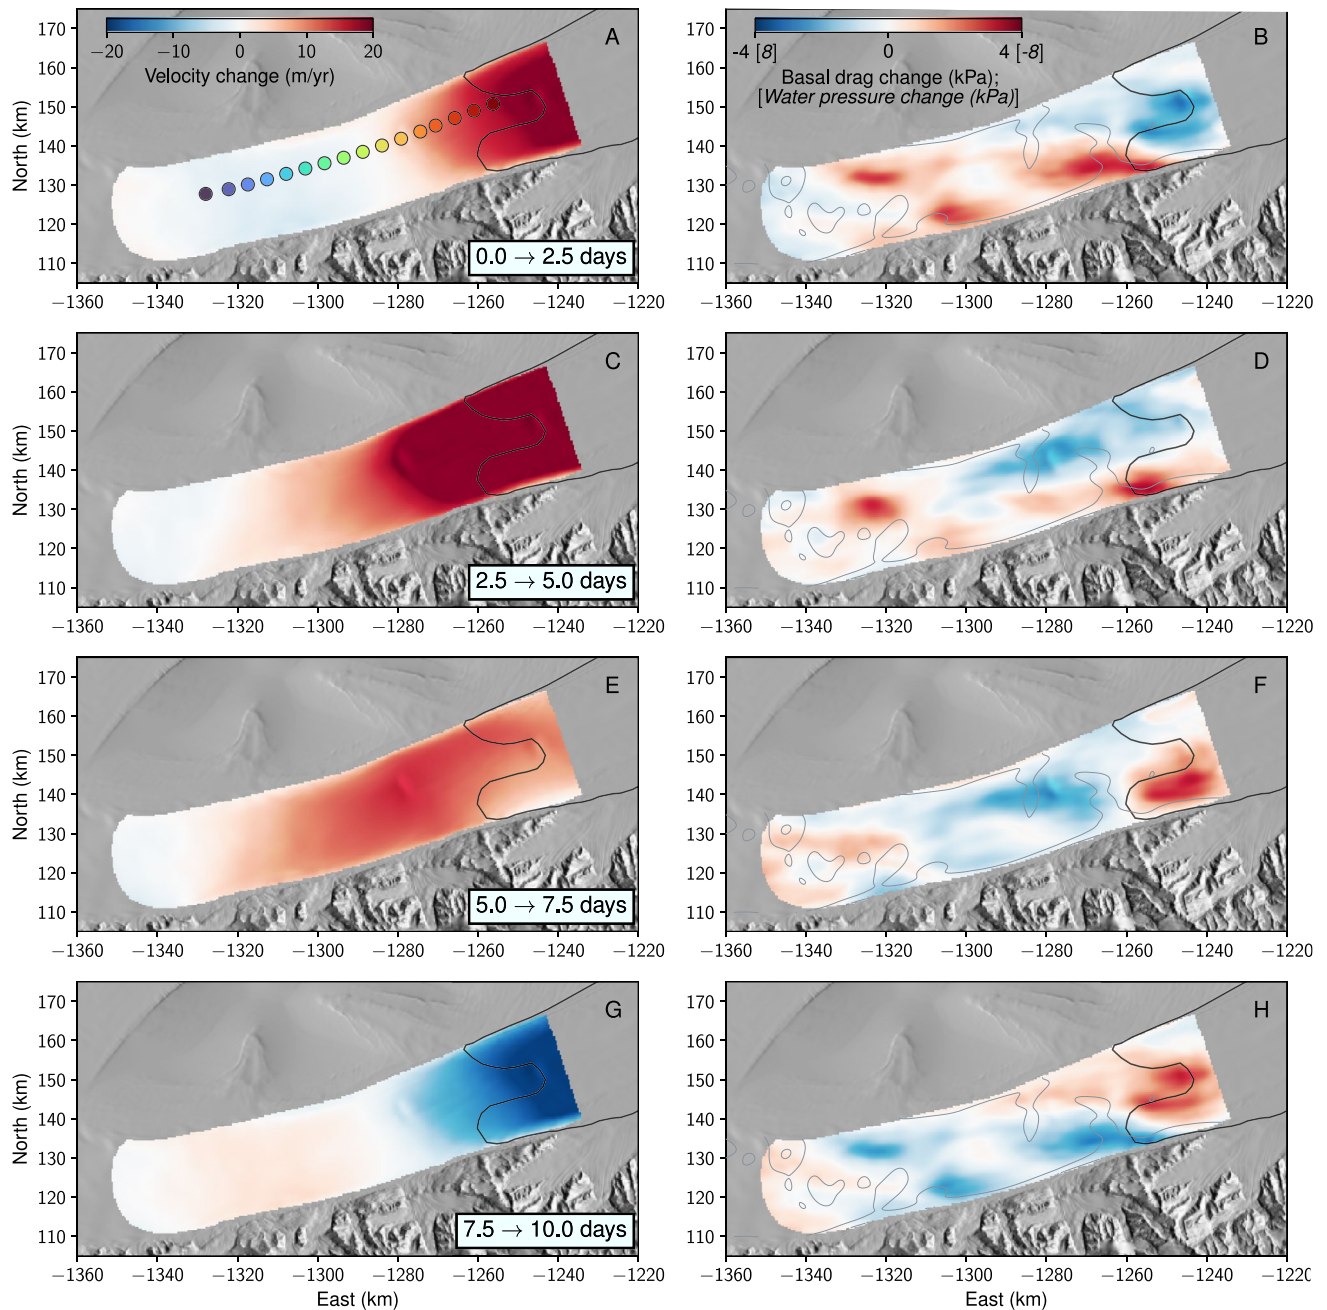
Figure 6. Time-dependent velocity and basal drag magnitude variations for Rutford Ice Stream, Antarctica. Velocity and basal drag variations are shown in the left and right columns, respectively. By assuming a plastic bed, basal water pressure variations can be inferred by scaling drag variations by the internal friction coefficient $\mu = 0.5$ (values indicated in square brackets and italics in (b)). Beginning at a reference time that approximately corresponds to the beginning of the M_{sf} cycle (minimum velocity near the grounding line), velocity and drag variations are measured in 2.5-day intervals: (a), (b) 0–2.5 days; (c), (d) 2.5–5 days; (e), (f) 5–7.5 days; and (g), (h) 7.5–10 days. Markers in (a) indicate points extracted for Figure 7. Gray contours for right-column plots correspond to the secular basal drag uncertainties in Figure 5h in intervals of 2.0 kPa. In general, an upstream-propagating increase in velocity is associated with an upstream-propagating decrease in basal drag, which suggests that a pressure wave in the subglacial till is responsible for the observed variations in surface velocity.

Under applicable conditions, successful recovery of sliding law parameters is largely dependent on the availability and temporal sampling of time-dependent velocity fields, as well as the noise wavelength and amplitude in the velocity and topography data. Static velocity snapshots allow only for the estimation of the magnitude of basal drag, which is equivalent to estimation of the joint probability distribution for the sliding law parameters in a Bayesian inference framework. Unique inference of one of the parameters would
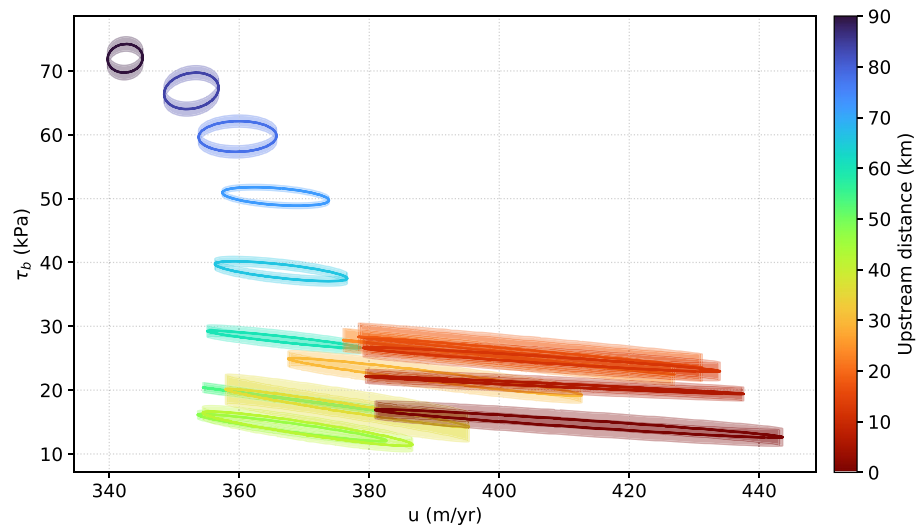
Figure 7. Time series of basal drag versus velocity magnitude for select points along a centerline. Points are colored by distance upstream of the grounding line with locations shown in Figure 6a. Shaded regions indicate 1σ uncertainties in predicted drag.

require some assumption on the value/distribution of the others, as well as an assumption on the form of the sliding law. Time-dependent velocity fields allow for quantification of time-dependent basal drag, permitting joint estimation of all sliding law parameters by quantifying the relationship between drag and sliding velocity at different points within the spatial domain. Furthermore, the larger the amplitude of velocity variability at any given location (e.g., amplitude of periodic variations due to ocean tides or seasonal effects), the better constrained the parameters (Figure 3). Of course, the level of noise or errors in the surface data will directly impact the uncertainty of the drag predictions, and as discussed in Section 3.1, errors with correlation lengthscales of several ice thicknesses will generally lead to larger artifacts in the predicted drag. For study areas where velocity and elevation measurements are more sparse or exhibit higher noise levels, the methods presented here would greatly benefit from a time series preprocessing stage that can fit some smoothly varying time function to the available data to inject stronger a priori knowledge about the underlying flow variations (Minchew et al., 2017; Riel et al., 2021), as was done for the RIS velocity data.

In this work, surface observations are used to compute the global stress balance directly via momentum balance equations under the assumptions that the surface velocities are approximately equal to basal sliding velocities and the rheology of the ice is reasonably well constrained. For the former, we note that the viscous nature of ice flow acts as a low-pass filter on basal stress variations such that variations with spatial scales < one ice thickness can result in similar surface velocities and elevations (Habermann et al., 2012). This non-uniqueness requires that any inversion technique (including ours) that assimilates noisy surface observations to estimate basal stress fields must use a regularization scheme to promote smoother basal stress or satisfy some other auxiliary constraint to obtain a unique solution (Habermann et al., 2012; Larour et al., 2012; Shapero et al., 2016). We do note that while it has been shown that the transfer function amplitude between variability in basal stress and surface velocities does decrease with decreasing spatial wavelengths, higher slip ratios (ratio of sliding to deformation velocity) can increase the transfer function amplitude (Gudmundsson, 2003). Thus, inversion for basal stress fields is generally more robust for fast-flowing glaciers which are dominated by basal sliding and where the spatial scale of basal stress variability is greater than the intrinsic resolution of the velocity fields (Stearns & Van der Veen, 2018).

When the ice rheology is subject to non-negligible uncertainties, any un-modeled variations in the rheology at the surface will lead to variations in the inferred basal drag via spatial gradients of the effective dynamic viscosity. When these gradients are small, as would be expected in the central trunk of the glacier, errors in the inferred drag should also be correspondingly small. In future work, we will explore joint estimation of rheological parameters and drag, perhaps by introducing a second auxiliary variable trained to predict a spatially varying flow rate parameter, $A(x, y)$. Since joint estimation of rheology and drag is an ill-posed

problem, we would likely need to incorporate additional conditioning/regularizing factors in the physics-based loss functions in order to constrain the solution space. As an example, we may encourage softer ice in high strain-rate areas (such as lateral shear margins) and anisotropic smoothing of the rheology and drag to enforce lower spatial gradients in the along-flow direction where strain-rates are orders of magnitude lower. As demonstrated by Ranganathan et al. (2020), such constraints can effectively reduce the inherent trade-offs between rheology and drag variations.

5.2. Rutford Ice Stream and Subglacial Hydrology

In general, speedups in ice flow respond to changes in driving, longitudinal, and basal stresses. A localized perturbation in longitudinal or driving stress (e.g., as a result of a calving event for a tidewater glacier) will result in a non-local redistribution of longitudinal stresses and velocity variations well away from the original perturbation. Similarly, a localized perturbation in basal drag will result in non-local redistribution of longitudinal stresses and velocity changes upstream (Joughin et al., 2019). On the other end of the spectrum, subglacial hydrological variations that result in traveling "pressure waves" are governed primarily by the properties of the hydraulic network, although indirect effects could arise from changes in surface slope (Hewitt & Fowler, 2008; Minchew & Meyer, 2020). In this case, surface velocity will respond to a combination of local reductions in basal drag (corresponding to the pressure wave front) and non-local variations in longitudinal stress. Consequently, quantifying velocity variation magnitudes without examining spatial gradients (i.e., strain rate variations) will not distinguish between these different forcing mechanisms, and an analysis of the stress states of the glacier is required (Rosier & Gudmundsson, 2016).

While a comprehensive comparison of the stress response to the different forcing mechanisms is reserved for future work, a simplified analysis of the evolution of longitudinal normal stresses can be used to infer the sign of the corresponding change in basal drag (Text S2). For RIS, gradients of longitudinal normal stresses decrease (become more negative) in response to increases in surface velocity. Therefore, assuming that lateral shear stress also becomes more resistive for increases in velocity, it follows that a decrease in basal drag is driving the velocity increases upstream. This simplified analysis, which is not subject to any modeling assumptions other than bulk ice rheology, further supports the inferred pressure wave.

Several recent modeling studies have proposed sub-glacial hydrology as the primary driver of long-period along-flow velocity variations near the grounding zone for RIS (e.g., Rosier & Gudmundsson, 2020; Warburton et al., 2020), as well as the high velocity variation amplitudes further upstream (Rosier et al., 2015). From the perspective of the work presented here, we implicitly assume that the basal drag is varying only at the M_{sf} frequency when we enforce the periodic time representation. This assumption is likely to be valid for the upstream portions of RIS (greater than a few ice thicknesses from the grounding zone) where elastic stress variations have decayed, thus limiting the response of ice flow to viscous effects (Rosier et al., 2015; Thompson et al., 2014). Simultaneous tracking of the velocity and basal drag variations suggests that the possible pressure wave lags behind the traveling wave of surface velocity (Figures 6 and 7), which is consistent with a pressure wave speed below the speed of longitudinal stress transmission. This constraint is almost certainly valid for real-world glaciers since pressure-driven subglacial water flow will be resisted by drag in the hydraulic network, so even for highly connected distributed systems, longitudinal stresses will propagate faster than basal drag variations (Warburton et al., 2020).

As previously discussed, estimation of the underlying sliding law parameters from surface observations is not possible without additional information if the parameters vary in time. However, we may still consider different endmembers for the sliding law exponents to explore the implications of the basal water pressure variations. Let us again consider the case where the bed is perfectly plastic such that the secular basal drag is equal to the yield stress, τ_y, of the bed. In this case, drag variations are entirely determined by the Mohr-Coulomb yield criteria in Equation 13. By assuming that ice thickness is approximately constant over (fortnightly) tidal timescales, then it follows that variations in drag take the form $\Delta \tau_b = -\mu \Delta p_w$, so that changes in drag are proportional to changes in basal water pressure. For an estimated basal drag variation amplitude of 4–5 kPa roughly 20 km upstream of the grounding line, the corresponding water pressure variation would then be 8–10 kPa for internal friction coefficient $\mu = 0.5$ (Figure 6). Following the subglacial hydrological model of Rosier et al. (2015), which assumes subglacial hydrology at RIS can be described as a homogenous porous medium, changes in water pressure can be related to changes in hydrologic head.

At the grounding line where the hydrological system is in direct contact with the ocean, hydrologic head is equal to the ocean elevation, and basal water pressure variations can be computed as:

$$\Delta p_w = \rho_w g \Delta S, \tag{16}$$

where ρ_w is the density of seawater and S is the height of the ocean surface. From tidal models and GPS records, tidal amplitudes are approximately 3 meters at RIS (Minchew et al., 2017; Padman et al., 2018; Rosier et al., 2015), which would lead to water pressure amplitudes of approximately 30 kPa at the grounding line. Thus, our estimate of 10 kPa basal water pressure change at a distance of 20 km upstream indicates an e-folding distance of approximately 20 km. Note that doubling the basal water pressure change to 20 kPa is equivalent to an e-folding distance of approximately 50 km, which is the same value for the velocity amplitudes (Minchew et al., 2017). The amplitudes of basal drag variations estimated for RIS are likely on the lower end of plausible values due to our higher choice for the smoothing hyperparameter (Figure S5), which was necessary to handle assumed errors in the surface and bed topographies. Thus, it is reasonable to expect that estimates of basal water pressure variations are as high as 15–20 kPa.

The upstream diffusion of hydrological head variations is a function of the conductivity of the hydrological system and the temporal frequency of the tidal forcing. While estimation of head variations over the grounded ice is beyond the scope of the work, the relative consistency between the estimated basal water pressure variations assuming a plastic bed and those predicted from a simple subglacial hydrological model provides some support for the possibility that the bed of RIS deforms plastically. If the sliding exponent was instead closer to $m = 3$, the negative feedback defined by the increased basal drag resistance caused by the velocity speedup would necessitate a nearly factor of two larger water pressure variation (e.g., Figure 4), which would be on the higher end of plausible values. Therefore, independent measurements of time-dependent basal water pressures would likely provide substantial information for constraining the sliding law exponent for RIS.

5.3. Ice Dynamics and Physics-Informed Neural Networks

The use of the SSA momentum balance to compute basal drag as a target metric for enforcing physical consistency follows the overall strategy of physics-informed neural networks (PINNs), wherein known physical relationships and constraints are used as auxiliary "data" for neural network training (Raissi et al., 2019). PINNs themselves are similar in nature to PDE-constrained optimization problems (see Morlighem et al. (2017) and Brinkerhoff and Johnson (2015) for cryosphere applications). The two primary advantages of PINNs for our purposes are: (a) the ability to evaluate the physics constraints at arbitrary space-time coordinates within the training domain, which prevents overfitting of the surface observations; and (b) negating the need for running (and then backpropagating gradients through) a sophisticated ice flow model. As discussed in Section 2.2.1, we are able to use this approach due to our narrower focus on inference of the time-evolution of basal drag for a given glacier without having to consider the physics governing key boundary conditions, particularly the dynamics of the grounding line and the surface mass balance. We consider these boundary conditions as implicit in the surface observations, allowing us to focus directly on the stress distribution in the momentum balance. One important implication of this approach is a lack of generalizability to other glaciers since neural networks simply function as physics-aware interpolators of the surface data. Of course, these methods are not intrinsically limited to a single glacier and could be scaled to a much wider region if surface data and bed topography are available. In this case, a neural network architecture with more and possibly wider layers would likely be necessary to capture the wider spatiotemporal variability of the data. One promising method for improving the computational efficiency of the PINN framework for our case is to treat each time slice of data as a high-dimensional training example for the neural network as opposed to scattered point examples as used in this work. While this method would require spatially continuous observations at all time epochs, it would utilize the efficiency of convolutional neural network architectures for computing spatial gradients and would likely decrease training time from 1–2 h to a few minutes due to training examples consisting of full images rather than individual points (Zhu et al., 2019). For very large modeling domains, one could improve scalability by reconstructing image patches rather than full images, although care would need to be taken for patch boundaries.

An alternative learning approach to the PINNs discussed here is to learn the full ice dynamics for a given glacier. Essentially, a neural network could be trained to predict the time evolution of ice velocity and

thickness completely from velocity and thickness time series without utilizing physical information from the momentum balance equations (Raissi, 2018). In this way, the representation of the glacier's dynamics would be purely generic and could be learned with minimal supervision, that is, "end-to-end" learning. However, the main challenge for this approach is also generalizability. In order for a pure neural network model to robustly predict the time evolution of a glacier or ice stream not seen during training, one would have to train the network with many different simulations spanning the expected parameter sets of all glaciers and ice streams over the globe. In other words, as the distribution of desired testing examples becomes wider, the distribution of training examples would also have to become wider to ensure that predictions are done in an interpolatory manner rather than an extrapolatory one. Considering the wide variety of bed topographies, sliding conditions, ice shelf conditions, climatic environment, and ice geometries, the training data would need to be prohibitively large in order to ensure generalizability without using any prior physics information. One potentially promising area of research utilizes flexible relational inductive biases encoded in graphs for improving generalizability of neural networks (Battaglia et al., 2018). This type of learning would relax the usage of a specific set of momentum balance equations while still utilizing additional information known from physical interactions between velocity and thickness.

5.4. Uncertainty Quantification

By prescribing the neural network f_θ to predict the distribution of the surface variables (via means and standard deviations of independent Gaussian distributions), we are able to obtain uncertainties on the predictions of those variables and on the derived basal drag. This uncertainty is governed by the misfit between the surface observations and the mean hypersurface learned by the neural network, that is, the neural network will inflate the prediction standard deviations in areas where the predictions deviate the most from the observations (Figure S7). This deviation is itself driven by a combination of data noise and the physics constraints on the basal drag. Thus, we would expect larger surface variable uncertainty when data noise is large or when the basal drag field is enforced to be smoother. Consequently, we expect uncertainties to vary with the choice of the smoothing penalty λ in Equation 10. We note that this uncertainty behavior arises from the likelihood loss $\mathcal{L}_{nll}$ in Equation 11 and specification of the data distribution as a trainable distribution such that heteroscedastic data variances are estimated along with the data means. For a fixed error model where data covariances are known a priori, one would likely need to use probabilistic neural networks (e.g., Bayesian neural networks (MacKay, 1995)) in order to recover an uncertainty measure that incorporates both data noise and the strength of the physics-based loss functions. For all of these approaches, we emphasize that the main derived utility is in comparison of relative uncertainties between different regions within the study domain, which is useful for outlining areas where observation noise is high or where velocity/topographic variations lead to SSA-based drag predictions that are not well-aligned to the prescribed level of spatial smoothing. Obtaining a reliable estimate of absolute uncertainties requires some form of uncertainty calibration with validation data (e.g., comparison of uncertainties with misfits between neural network predictions and data unseen during training). We leave uncertainty calibration and comparison of uncertainties with more rigorous MCMC-based methods (e.g., Petra et al., 2014) for future work.

From a broader perspective, we believe that the probabilistic learning framework takes a significant step towards general quantification of both data and modeling uncertainties within a geophysical context while lowering the burden to run computationally expensive MCMC methods. This uncertainty quantification for hybrid physical and machine learning models has proven to be useful in related fields such as atmospheric dynamics (Stuart & Teckentrup, 2018). Other probabilistic machine learning models, such as Gaussian processes (Rasmussen, 2003), may also be a suitable surrogate for the basal drag, and recent advances in variational Gaussian processes that allow for training on large datasets make them a compelling machine learning model for future work (Hensman et al., 2015).

6. Conclusion

We have presented a hybrid machine learning framework for learning the time-evolution of basal mechanics for glaciers and ice streams. This approach integrates into the learning procedure well-known ice flow momentum balance equations at various levels of complexity. The a priori physical knowledge allows for the transformation of ice velocity, thickness, and surface elevation measurements into a domain where a

neural network can directly predict basal drag. Furthermore, we demonstrated the utility of probabilistic loss functions for quantifying uncertainties for the basal drag predictions, which will prove to be invaluable for subsequent interpretation of the drag, inference of sliding law parameters, and development of future data acquisition plans. As a real-world example, application of these techniques to time-dependent velocity data over Rutford Ice Stream, Antarctica, resulted in observational evidence of subglacial hydrological effects during the tidal cycle.

From a broader perspective, this work demonstrates a new and rapidly advancing approach for combining the physical knowledge gained from decades of theoretical and experimental work with modern data-driven techniques in order to address an outstanding problem in glacier dynamics, mainly determination of the sliding mode via the form of the inferred sliding law. Under certain forcing environments, we demonstrated that estimation of the value and uncertainty of the exponent in the power-law form of the sliding law is possible with these methods. The exponentially increasing data volume over the fastest flowing areas in the cryosphere demands techniques that combine data efficiency, modeling flexibility, and robustness in the presence of noise, data gaps, and modeling uncertainties. The methods presented here take an important step towards those requirements and present a path forward for future data assimilation tasks for a multitude of disparate data sources.

Appendix A: Ice Flow Model

A1. Governing Equations

In its most general form, glacier flow can be described as an incompressible Stokes flow:

$$\nabla \cdot \boldsymbol{\sigma} + \rho_i \mathbf{g} = \mathbf{0}, \tag{A1}$$

$$\mathrm{Tr}\left(\dot{\epsilon}\right) = 0, \tag{A2}$$

where $\nabla \cdot \boldsymbol{\sigma}$ is the divergence of the Cauchy stress tensor $\boldsymbol{\sigma}$, ρ_i is the density of ice, $\mathbf{g}$ is the gravitational acceleration, $\dot{\epsilon}$ is the strain rate tensor, and Tr is the trace operator (here, bold font indicates tensor and vector quantities while regular font represents scalars). Equation A1 describes the stress balance (also referred to as the momentum balance) while Equation A2 represents the incompressibility of ice. The strain rate tensor describes the rate of deformation of ice and is calculated as the symmetric component of the velocity gradient:

$$\dot{\epsilon} = \frac{1}{2}\left(\nabla \boldsymbol{u} + \left(\nabla \boldsymbol{u}\right)^T\right), \tag{A3}$$

where $\boldsymbol{u} = \left[u, v, w\right]$ is the velocity vector. To relate the stress tensor components in Equation A1 to velocity components, the constitutive law for incompressible viscous fluids is used:

$$\boldsymbol{\tau} = 2\eta\dot{\epsilon}, \tag{A4}$$

where $\boldsymbol{\tau} = \boldsymbol{\sigma} + p\mathbf{I}$ is the deviatoric stress tensor, $p = -\mathrm{Tr}(\boldsymbol{\sigma})/3$ is the isotropic pressure, and $\mathbf{I}$ is the identity matrix. The non-Newtonian effective ice dynamic viscosity, η, is given as:

$$\eta = \frac{1}{2}A^{-\frac{1}{n}}\dot{\epsilon}_e^{\frac{1-n}{n}}, \tag{A5}$$

where n is the stress exponent in Glen's flow law, $\dot{\epsilon}_e$ is the effective strain rate (square root of the second invariant of the strain rate tensor), and A is the flow rate factor which depends on properties of the ice (e.g., temperature, interstitial liquid water content, crystal size/orientation, and impurity content). In practice, many studies have found that various approximations to the computationally expensive full Stokes equations (Equations A1 and A2) are able to reconstruct observed velocity fields fairly well for certain glacier geometries and result in similar implied glacier mechanics. For the types of glaciers and ice streams we examine in this study, the Shallow Ice Shelf/Stream Approximation (SSA) (MacAyeal, 1989) is most commonly used and assumes: (a) ice thickness is much smaller than the horizontal span; (b) most of the forward motion of fast-flowing glaciers is due to sliding at the bed (i.e., vertical shearing is negligible); and (c) total vertical normal stress is equal to the ice overburden pressure. Under these assumptions, the 3D momentum balance can be depth averaged along the z-dimension, and by using the constitutive law in Equation A4, the approximate 2D momentum balance is expressed as (identical to the main text):

$$\frac{\partial}{\partial x}\left(2\eta h\left(2\frac{\partial u}{\partial x}+\frac{\partial v}{\partial y}\right)\right)+\frac{\partial}{\partial y}\left(\eta h\left(\frac{\partial u}{\partial y}+\frac{\partial v}{\partial x}\right)\right)-\tau_{bx}=\rho_i g h\frac{\partial s}{\partial x}, \tag{A6a}$$

$$\frac{\partial}{\partial y}\left(2\eta h\left(2\frac{\partial v}{\partial y}+\frac{\partial u}{\partial x}\right)\right)+\frac{\partial}{\partial x}\left(\eta h\left(\frac{\partial u}{\partial y}+\frac{\partial v}{\partial x}\right)\right)-\tau_{by}=\rho_i g h\frac{\partial s}{\partial y}, \tag{A6b}$$

where h is the ice thickness, τ_{bx} and τ_{by} represent the x- and y-components of basal shear stress, and s is the ice surface elevation. The vertical velocity component w can be recovered using the incompressibility condition. The above momentum balance states that the gravitational horizontal driving stresses of ice flow (terms on the right-hand side) are balanced by a combination of horizontal gradients of deviatoric stresses and drag at the base of the glacier, τ_b.

A2. 1D Shallow Ice Stream Model Boundary Conditions

Following previous work on 1D flowline models, we enforce two Neumann boundary conditions at the edges of the spatial domain (Nick et al., 2009; Vieli & Payne, 2005). At the ice divide boundary condition ($x = 0$), a symmetric ice sheet is assumed such that $\partial u/\partial x = 0$. At the grounding line (assuming no ice shelf), the boundary condition is derived from the difference between the hydrostatic pressure of water and ice:

$$\frac{\partial u}{\partial x}=f_s A\left[\frac{1}{4}\rho_i g h\left(1-\frac{\rho_i}{\rho_w}\right)\right]^n, \tag{A7}$$

where ρ_w is the density of ocean water and f_s is a scalar factor used to apply a time-varying force on the calving face (Nick et al., 2009). While Equation A7 is not strictly applicable to marine-terminating ice streams with no ice shelf, it provides a convenient way to apply longitudinal stress perturbations originating at the terminus. Thus, we can generate time-dependent ice velocity and thickness fields representative of those observed at tidewater glaciers that respond to changes in regional oceanic and climate conditions. For all 1D simulations, we use a flow rate factor $A = 1.2 \times 10^{-24}\,\mathrm{s}^{-1}\mathrm{Pa}^{-3}$, which corresponds to a temperature of -5°C and an exponent $n = 3$ (Cuffey & Paterson, 2010). We solve these equations in a staggered fashion by solving for u in Equation 3 under the stated boundary conditions for a given thickness profile, h, using Newton's method. Mass continuity gives the time evolution of ice thickness, h:

$$\frac{\partial h}{\partial t}=a-\frac{\partial q}{\partial x}, \tag{A8}$$

where a is the surface mass balance (difference between snow accumulation and ablation) and $q = hu$ is the width-averaged ice flux. Thus, we implement a forward Euler step for Equation A8 to update the thickness profile.

Appendix B: Network Architecture

We use feedforward networks for all neural networks in this work. The hidden layers have the form $\mathbf{Wx} + \mathbf{b}$ followed by an activation with a hyperbolic tangent (tanh) function. During our experiments, we found that activation functions that were at least twice differentiable (e.g., tanh or exponential rectified linear units (ELU)) resulted in smoother spatial gradients of output variables than the rectified linear unit (ReLU). We found very little difference in training convergence speed between tanh and ELU activations, so we use tanh throughout this work. The outputs of all networks are linear (i.e., no activation is applied).

Neural networks tasked with reconstructing velocity and thickness observations were prescribed 4–6 hidden layers where each hidden layer consisted of 50 or 100 hidden units. The exact architectures varied for each problem and were qualitatively chosen based on a balance between reconstruction accuracy, spatial smoothness of the reconstruction, and computational efficiency. Regardless, the tradeoffs between the metrics were minor, and the data reconstructions for all network architectures were largely consistent. The neural networks for basal drag predictions were prescribed 4 hidden layers where each hidden layer consisted of 50 units. In this way, we effectively applied more regularization for these networks as compared with the data networks since our prior assumption for the spatial distribution of basal shear stress is one that is smooth.

B1. Training

Weight matrices for all networks are initialized from a normal distribution with variances specified by $s = 1/\sqrt{a}$ where a is the number of input hidden units for each layer. Inputs to all networks are normalized to be zero-mean with unit variance. We use the Adam optimizer (Kingma & Ba, 2014) with a learning rate of 0.0002 and train for 500–1,000 epochs (each epoch is defined as a complete pass through the training data). We use the Python API for TensorFlow (Abadi et al., 2015) and TensorFlow Probability (Dillon et al., 2017) for neural network construction and training.

We use a train/test split where 85% of the data is used for training and the remaining 15% is used for validation where the split is performed in pixel space, that is, random points throughout the space-time volume. We also experimented with a train/test split across time slices where entire spatial fields are held out in the test set. In the cases examined here, both splitting schemes resulted in similar test set losses, which is likely due to the density of surface data available (Figure B1). We do observe a slightly larger basal drag smoothness cost for the time slice splitting scheme (Figure B1f). Overall, when data are more sparse, we expect the time slice splitting scheme to give a more challenging test set, which may be mitigated by larger drag smoothing constraints or regularization of the neural network weights.
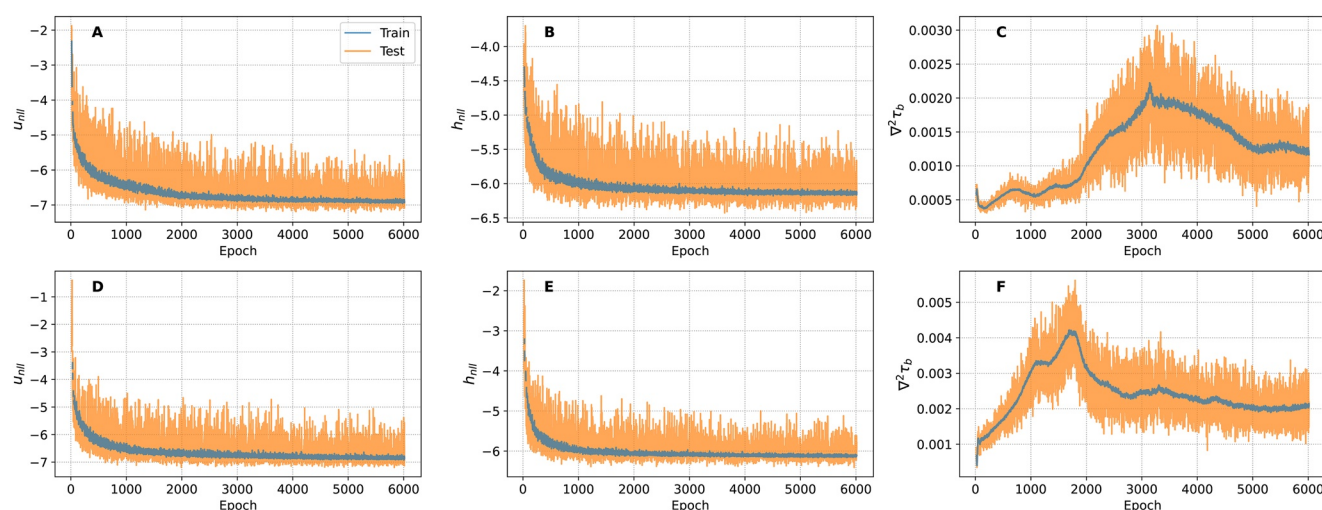

Figure B1. Training performance for example 1D simulation in Section 2.3. The top plots show the evaluated cost functions for the velocity likelihood, thickness likelihood, and basal drag smoothing (A, B, and C, respectively) using a train/test split in pixel space. Blue lines represent the epoch-averaged training losses and orange lines represent the test losses. The bottom plots show the same cost functions but for a train/test split across time slices.

Conflict of Interest

The authors declare that they have no competing interests.

Data Availability Statement

All simulation data and deep learning training code can be downloaded at https://doi.org/10.5281/zenodo.4755695. Derived velocity components for Rutford Ice Stream, Antarctica can be downloaded at https://doi.org/10.5281/zenodo.5508397. Original COSMO-SkyMed products (copyright ASI Agenzia Spaziale Italiana, 2013–2016) are archived at Jet Propulsion Laboratory and were processed under license from ASI as part of a collaborative project between CIDOT and JPL/Caltech. Bed topography for Antarctica available at https://nsidc.org/data/NSIDC-0756/versions/1.

References

Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., et al. (2015). TensorFlow: Large-scale machine learning on heterogeneous systems. *Proceedings of the 12th USENIX Symposium on Operating Systems Design and Implementation*. Retrieved from https://www.tensorflow.org/ (Software available from tensorflow.org).

Aschwanden, A., Fahnestock, M. A., Truffer, M., Brinkerhoff, D. J., Hock, R., Khroulev, C., & Khan, S. A. (2019). Contribution of the greenland ice sheet to sea level over the next millennium. *Science Advances*, *5*(6), eaav9396. https://doi.org/10.1126/sciadv.aav9396

Battaglia, P. W., Hamrick, J. B., Bapst, V., Sanchez-Gonzalez, A., Zambaldi, V., Malinowski, M., & others (2018). *Relational inductive biases, deep learning, and graph networks*. arXiv preprint arXiv:1806.01261.

Baydin, A. G., Pearlmutter, B. A., Radul, A. A., & Siskind, J. M. (2017). Automatic differentiation in machine learning: A survey. *Journal of Machine Learning Research*, *18*(1), 5595–5637.

Bengio, Y., Courville, A., & Vincent, P. (2013). Representation learning: A review and new perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, *35*(8), 1798–1828. https://doi.org/10.1109/tpami.2013.50

Bölcskei, H., Grohs, P., Kutyniok, G., & Petersen, P. (2019). Optimal approximation with sparsely connected deep neural networks. *SIAM Journal on Mathematics of Data Science*, *1*(1), 8–45. https://doi.org/10.1137/18m118709x

Bondzio, J. H., Morlighem, M., Seroussi, H., Kleiner, T., Rückamp, M., Mouginot, J., & Humbert, A. (2017). The mechanisms behind jakobshavn isbræ's acceleration and mass loss: A 3-D thermomechanical model study. *Geophysical Research Letters*, *44*(12), 6252–6260. https://doi.org/10.1002/2017GL073309

Brinkerhoff, D., & Johnson, J. (2015). A stabilized finite element method for calculating balance velocities in ice sheets. *Geoscientific Model Development*, *8*(5), 1275–1283. https://doi.org/10.5194/gmd-8-1275-2015

Church, J. A., Clark, P. U., Cazenave, A., Gregory, J. M., Jevrejeva, S., Levermann, A., et al. (2013). *Sea level change (Tech. Rep.)*: PM Cambridge University Press. http://drs.nio.org/drs/handle/2264/4605

Cover, T. M., & Thomas, J. A. (1999). *Elements of information theory*. John Wiley & Sons.

Cuffey, K. M., & Paterson, W. S. B. (2010). *The physics of glaciers*. Academic Press.

Cybenko, G. (1989). Approximation by superpositions of a sigmoidal function. *Mathematics of control, signals and systems*, *2*(4), 303–314. https://doi.org/10.1007/bf02551274

De Rydt, J., Gudmundsson, G. H., Corr, H., & Christoffersen, P. (2013). Surface undulations of Antarctic ice streams tightly controlled by bedrock topography. *The Cryosphere*, *7*, 407–417. https://doi.org/10.5194/tc-7-407-2013

Delalleau, O., & Bengio, Y. (2011). Shallow vs. deep sum-product networks. *Advances in Neural Information Processing Systems*, *24*, 666–674.

Dillon, J. V., Langmore, I., Tran, D., Brevedo, E., Vasudevan, S., Moore, D., & Saurous, R. A. (2017). *Tensorflow distributions*. arXiv preprint arXiv:1711.10604.

Duputel, Z., Agram, P. S., Simons, M., Minson, S. E., & Beck, J. L. (2014). Accounting for prediction uncertainty when inferring subsurface fault slip. *Geophysical Journal International*, *197*(1), 464–482. https://doi.org/10.1093/gji/ggt517

Enderlin, E. M., O'Neel, S., Bartholomaus, T. C., & Joughin, I. (2018). Evolving environmental and geometric controls on columbia glacier's continued retreat. *Journal of Geophysical Research: Earth Surface*, *123*(7), 1528–1545. https://doi.org/10.1029/2017jf004541

Flowers, G. E. (2015). Modelling water flow under glaciers and ice sheets. *Proceedings of the Royal Society A: Mathematical, Physical & Engineering Sciences*, *471*(2176), 20140907. https://doi.org/10.1098/rspa.2014.0907

Gillet-Chaulet, F. (2020). Assimilation of surface observations in a transient marine ice sheet model using an ensemble kalman filter. *The Cryosphere*, *14*(3), 811–832. https://doi.org/10.5194/tc-14-811-2020

Gillet-Chaulet, F., Durand, G., Gagliardini, O., Mosbeux, C., Mouginot, J., Rémy, F., & Ritz, C. (2016). Assimilation of surface velocities acquired between 1996 and 2010 to constrain the form of the basal friction law under Pine Island Glacier. *Geophysical Research Letters*, *43*(19), 10–311. https://doi.org/10.1002/2016gl069937

Goldberg, D., Heimbach, P., Joughin, I., & Smith, B. (2015). Committed retreat of Smith, Pope, and Kohler Glaciers over the next 30 years inferred by transient model calibration. *The Cryosphere*, *9*(6), 2429–2446. https://doi.org/10.5194/tc-9-2429-2015

Gudmundsson, G. H. (2003). Transmission of basal variability to a glacier surface. *Journal of Geophysical Research: Solid Earth*, *108*(B5), 2253. https://doi.org/10.1029/2002jb002107

Gudmundsson, G. H. (2006). Fortnightly variations in the flow velocity of rutford ice stream, west Antarctica. *Nature*, *444*(7122), 1063–1064. https://doi.org/10.1038/nature05430

Gudmundsson, G. H. (2007). Tides and the flow of Rutford Ice Stream, West Antarctica. *Journal of Geophysical Research: Earth Surface*, *112*(F4). https://doi.org/10.1029/2006JF000731

Gudmundsson, G. H., & Jenkins, A. (2009). Ice-flow velocities on Rutford Ice Stream, West Antarctica, are stable over decadal timescales. *Journal of Glaciology*, *55*(190), 339–344. https://doi.org/10.3189/002214309788608697

Gudmundsson, G. H. (2011). Ice-stream response to ocean tides and the form of the basal sliding law. *The Cryosphere*, *5*, 259–270. https://doi.org/10.5194/tc-5-259-2011

Habermann, M., Maxwell, D., & Truffer, M. (2012). Reconstruction of basal properties in ice sheets using iterative inverse methods. *Journal of Glaciology*, *58*(210), 795–808. https://doi.org/10.3189/2012jog11j168

Habermann, M., Truffer, M., & Maxwell, D. (2013). Changing basal conditions during the speed-up of Jakobshavn Isbræ, Greenland. *The Cryosphere*, *7*(6). https://doi.org/10.5194/tc-7-1679-2013

Hensman, J., Matthews, A., & Ghahramani, Z. (2015). Scalable variational Gaussian process classification. *Proceedings of the Eighteenth International Conference on Artificial Intelligence and Statistics*, *38*, 351–360.

Hewitt, I. J. (2013). Seasonal changes in ice sheet motion due to melt water lubrication. *Earth and Planetary Science Letters*, *371*, 16–25. https://doi.org/10.1016/j.epsl.2013.04.022

Hewitt, I. J., & Fowler, A. C. (2008). Seasonal waves on glaciers. *Hydrological Processes*, *22*(19), 3919–3930. https://doi.org/10.1002/hyp.7029

Howat, I. M., Porter, C., Smith, B. E., Noh, M.-J., & Morin, P. (2019). The reference elevation model of Antarctica. *The Cryosphere*, *13*(2), 665–674. https://doi.org/10.5194/tc-13-665-2019

Iken, A., & Bindschadler, R. A. (1986). Combined measurements of subglacial water pressure and surface velocity of Findelengletscher, Switzerland: Conclusions about drainage system and sliding mechanism. *Journal of Glaciology*, *32*(110), 101–119. https://doi.org/10.3189/s0022143000006936

Iverson, N. R. (2010). Shear resistance and continuity of subglacial till: Hydrology rules. *Journal of Glaciology*, *56*(200), 1104–1114. https://doi.org/10.3189/002214311796406220

Joughin, I., Bamber, J. L., Scambos, T., Tulaczyk, S., Fahnestock, M., & MacAyeal, D. R. (2006). Integrating satellite observations with modelling: Basal shear stress of the Filcher-Ronne ice streams, Antarctica. *Philosophical Transactions of the Royal Society A: Mathematical, Physical & Engineering Sciences*, *364*(1844), 1795–1814. https://doi.org/10.1098/rsta.2006.1799

Joughin, I., Smith, B., Howat, I., & Scambos, T. (2011). *MEaSUREs Greenland ice velocity: Selected glacier site velocity maps from InSAR*. NASA National Snow and Ice Data Center Distributed Active Archive Center. https://doi.org/10.5067/MEASURES/CRYOSPHERE/nsidc-0481.001

Joughin, I., Smith, B. E., Howat, I. M., Floricioiu, D., Alley, R. B., Truffer, M., & Fahnestock, M. (2012). Seasonal to decadal scale variations in the surface velocity of Jakobshavn Isbrae, Greenland: Observation and model-based analysis. *Journal of Geophysical Research: Earth Surface*, *117*(F2), F02030. https://doi.org/10.1029/2011JF002110

Joughin, I., Smith, B. E., & Schoof, C. G. (2019). Regularized Coulomb friction laws for ice sheet sliding: Application to Pine Island Glacier, Antarctica. *Geophysical Research Letters*, *46*(9), 4764–4771. https://doi.org/10.1029/2019GL082526

Kamb, B., & Echelmeyer, K. A. (1986). Stress-gradient coupling in glacier flow: I. longitudinal averaging of the influence of ice thickness and surface slope. *Journal of Glaciology*, *32*(111), 267–284. https://doi.org/10.1017/s0022143000015604

Karpatne, A., Atluri, G., Faghmous, J. H., Steinbach, M., Banerjee, A., Ganguly, A., & Kumar, V. (2017). Theory-guided data science: A new paradigm for scientific discovery from data. *IEEE Transactions on Knowledge and Data Engineering*, *29*(10), 2318–2331. https://doi.org/10.1109/tkde.2017.2720168

Kingma, D. P., & Ba, J. (2014). *Adam: A method for stochastic optimization*. arXiv preprint arXiv:1412.6980.

Larour, E., Seroussi, H., Morlighem, M., & Rignot, E. (2012). Continental scale, high order, high spatial resolution, ice sheet modeling using the Ice Sheet System Model (ISSM). *Journal of Geophysical Research: Earth Surface*, *117*(F1), F01022. https://doi.org/10.1029/2011jf002140

Larour, E., Utke, J., Csatho, B., Schenk, A., Seroussi, H., Morlighem, M., & Khazendar, A. (2014). Inferred basal friction and surface mass balance of the Northeast Greenland Ice Stream using data assimilation of ICESat (Ice Cloud and land Elevation Satellite) surface altimetry and ISSM (Ice Sheet System Model). *The Cryosphere*, *8*(6), 2335–2351. https://doi.org/10.5194/tc-8-2335-2014

LeCun, Y., Bengio, Y., & Hinton, G. (2015). Deep learning. *Nature*, *521*(7553), 436–444. https://doi.org/10.1038/nature14539

Lu, Z., Pu, H., Wang, F., Hu, Z., & Wang, L. (2017). The expressive power of neural networks: A view from the width. *Proceedings of the 31st International Conference on Neural Information Processing Systems*. arXiv preprint arXiv:1709.02540.

MacAyeal, D. R. (1989). Large-scale ice flow over a viscous basal sediment: Theory and application to ice stream B, Antarctica. *Journal of Geophysical Research: Solid Earth*, *94*(B4), 4071–4087. https://doi.org/10.1029/jb094ib04p04071

MacAyeal, D. R. (1993). A tutorial on the use of control methods in ice-sheet modeling. *Journal of Glaciology*, *39*(131), 91–98. https://doi.org/10.3189/s0022143000015744

MacKay, D. J. (1995). Bayesian neural networks and density networks. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, *354*(1), 73–80. https://doi.org/10.1016/0168-9002(94)00931-7

McCormack, F. S., Roberts, J. L., Jong, L. M., Young, D. A., & Beem, L. H. (2019). A note on digital elevation model smoothing and driving stresses. *Polar Research*, *38*. https://doi.org/10.33265/polar.v38.3498

Minchew, B., & Joughin, I. (2020). Toward a universal glacier slip law. *Science*, *368*(6486), 29–30. https://doi.org/10.1126/science.abb3566

Minchew, B., & Meyer, C. R. (2020). Dilation of subglacial sediment governs incipient surge motion in glaciers with deformable beds. *Proceedings of the Royal Society A*, *476*(2238), 20200033. https://doi.org/10.1098/rspa.2020.0033

Minchew, B., Simons, M., Björnsson, H., Pálsson, F., Morlighem, M., Seroussi, H., et al. (2016). Plastic bed beneath Hofsjökull Ice Cap, central Iceland, and the sensitivity of ice flow to surface meltwater flux. *Journal of Glaciology*, *62*(231), 147–158. https://doi.org/10.1017/jog.2016.26

Minchew, B., Simons, M., Riel, B., & Milillo, P. (2017). Tidally induced variations in vertical and horizontal motion on Rutford Ice Stream, West Antarctica, inferred from remotely sensed observations. *Journal of Geophysical Research: Earth Surface*, *122*(1), 167–190. https://doi.org/10.1002/2016JF003971

Morlighem, M., Rignot, E., Binder, T., Blankenship, D., Drews, R., Eagles, G., et al. (2020). Deep glacial troughs and stabilizing ridges unveiled beneath the margins of the Antarctic ice sheet. *Nature Geoscience*, *13*(2), 132–137. https://doi.org/10.1038/s41561-019-0510-8

Morlighem, M., Rignot, E., Seroussi, H., Larour, E., Ben Dhia, H., & Aubry, D. (2010). Spatial patterns of basal drag inferred using control methods from a full-stokes and simpler models for pine island glacier, west Antarctica. *Geophysical Research Letters*, *37*(14). https://doi.org/10.1029/2010gl043853

Morlighem, M., Rignot, E., Seroussi, H., Larour, E., Ben Dhia, H., & Aubry, D. (2011). A mass conservation approach for mapping glacier ice thickness. *Geophysical Research Letters*, *38*(19), L19503. https://doi.org/10.1029/2011GL048659

Morlighem, M., Williams, C. N., Rignot, E., An, L., Arndt, J. E., Bamber, J. L., & Zinglersen, K. B. (2017). Bedmachine v3: Complete bed topography and ocean bathymetry mapping of Greenland from multibeam echo sounding combined with mass conservation. *Geophysical Research Letters*, *44*(21), 11051–11061. https://doi.org/10.1002/2017GL074954

Mouginot, J., Rignot, E., & Scheuchl, B. (2019). Continent-wide, interferometric SAR phase, mapping of Antarctic ice velocity. *Geophysical Research Letters*, *46*(16), 9710–9718. https://doi.org/10.1029/2019GL083826

Murray, T., Smith, A. M., King, M. A., & Weedon, G. P. (2007). Ice flow modulated by tides at up to annual periods at Rutford Ice Stream, West Antarctica. *Geophysical Research Letters*, *34*(18), L18503. https://doi.org/10.1029/2007gl031207

Nick, F. M., Vieli, A., Howat, I. M., & Joughin, I. (2009). Large-scale changes in Greenland outlet glacier dynamics triggered at the terminus. *Nature Geoscience*, *2*(2), 110. https://doi.org/10.1038/ngeo394

Padman, L., Siegfried, M. R., & Fricker, H. A. (2018). Ocean tide influences on the Antarctic and Greenland ice sheets. *Reviews of Geophysics*, *56*(1), 142–184. https://doi.org/10.1002/2016RG000546

Petra, N., Martin, J., Stadler, G., & Ghattas, O. (2014). A computational framework for infinite-dimensional bayesian inverse problems, part ii: Stochastic Newton MCMC with application to ice sheet flow inverse problems. *SIAM Journal on Scientific Computing*, *36*(4), A1525–A1555. https://doi.org/10.1137/130934805

Porter, C., Morin, P., Howat, I., Noh, M.-J., Bates, B., Peterman, K., & Bojesen, M. (2018). *ArcticDEM*. Harvard Dataverse. https://doi.org/10.7910/DVN/OHHUKH

Pralong, M. R., & Gudmundsson, G. H. (2011). Bayesian estimation of basal conditions on Rutford Ice Stream, West Antarctica, from surface data. *Journal of Glaciology*, *57*(202), 315–324. https://doi.org/10.3189/002214311796406004

Raissi, M. (2018). Deep hidden physics models: Deep learning of nonlinear partial differential equations. *Journal of Machine Learning Research*, *19*(1), 932–955.

Raissi, M., Perdikaris, P., & Karniadakis, G. E. (2019). Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, *378*, 686–707. https://doi.org/10.1016/j.jcp.2018.10.045

Ranganathan, M., Minchew, B., Meyer, C. R., & Gudmundsson, G. H. (2020). A new approach to inferring basal drag and ice rheology in ice streams, with applications to West Antarctic Ice Streams. *Journal of Glaciology*, *67*, 1–14. https://doi.org/10.1017/jog.2020.95

Rasmussen, C. E. (2003). Gaussian processes in machine learning. In *Advanced Lectures on Machine Learning: ML Summer Schools 2003, Canberra, Australia* (pp. 63–71). Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-540-28650-9_4

Reichstein, M., Camps-Valls, G., Stevens, B., Jung, M., Denzler, J., Carvalhais, N., et al. (2019). Deep learning and process understanding for data-driven Earth system science. *Nature*, *566*(7743), 195–204. https://doi.org/10.1038/s41586-019-0912-1

Riel, B., Minchew, B., & Joughin, I. (2021). Observing traveling waves in glaciers with remote sensing: New flexible time series methods and application to Sermeq Kujalleq (Jakobshavn Isbræ), Greenland. *The Cryosphere*, *15*(1), 407–429. https://doi.org/10.5194/tc-15-407-2021

Rignot, E., Mouginot, J., & Scheuchl, B. (2011). Ice Flow of the Antarctic Ice Sheet. *Science*, *333*(6048), 1427–1430. https://doi.org/10.1126/science.1208336

Ritz, C., Edwards, T. L., Durand, G., Payne, A. J., Peyaud, V., & Hindmarsh, R. C. (2015). Potential sea-level rise from Antarctic ice-sheet instability constrained by observations. *Nature*, *528*(7580), 115–118. https://doi.org/10.1038/nature16147

Robel, A. A., Seroussi, H., & Roe, G. H. (2019). Marine ice sheet instability amplifies and skews uncertainty in projections of future sea-level rise. *Proceedings of the National Academy of Sciences*, *116*(30), 14887–14892. https://doi.org/10.1073/pnas.1904822116

Robel, A. A., Tsai, V. C., Minchew, B., & Simons, M. (2017). Tidal modulation of ice shelf buttressing stresses. *Annals of Glaciology*, *58*(74), 12–20. https://doi.org/10.1017/aog.2017.22

Rosier, S. H. R., & Gudmundsson, G. H. (2016). Tidal controls on the flow of ice streams. *Geophysical Research Letters*, *43*(9), 4433–4440. https://doi.org/10.1002/2016GL068220

Rosier, S. H. R., & Gudmundsson, G. H. (2020). Exploring mechanisms responsible for tidal modulation in flow of the filchner–ronne ice shelf. *The Cryosphere*, *14*(1), 17–37. https://doi.org/10.5194/tc-14-17-2020

Rosier, S. H. R., Gudmundsson, G. H., & Green, J. A. M. (2015). Temporal variations in the flow of a large Antarctic ice-stream controlled by tidally induced changes in the subglacial water system. *The Cryosphere*, *9*(4), 1649–1661. https://doi.org/10.5194/tc-9-1649-2015

Schoof, C. (2005). The effect of cavitation on glacier sliding. *Proceedings of the Royal Society of London - Series A: Mathematical and Physical Sciences*, *461*, 609–627. https://doi.org/10.1098/rspa.2004.1350

Schoof, C. (2007). Marine ice-sheet dynamics. Part 1. The case of rapid sliding. *Journal of Fluid Mechanics*, *573*, 27. https://doi.org/10.1017/s0022112006003570

Schoof, C. (2010). Ice-sheet acceleration driven by melt supply variability. *Nature*, *468*(7325), 803–806. https://doi.org/10.1038/nature09618

Shapero, D. R., Joughin, I. R., Poinar, K., Morlighem, M., & Gillet-Chaulet, F. (2016). Basal resistance for three of the largest Greenland outlet glaciers. *Journal of Geophysical Research: Earth Surface*, *121*(1), 168–180. https://doi.org/10.1002/2015jf003643

Smith, E., Smith, A., White, R., Brisbourne, A., & Pritchard, H. (2015). Mapping the ice-bed interface characteristics of Rutford Ice Stream, West Antarctica, using microseismicity. *Journal of Geophysical Research: Earth Surface*, *120*(9), 1881–1894. https://doi.org/10.1002/2015jf003587

Stearns, L., & Van der Veen, C. (2018). Friction at the bed does not control fast glacier flow. *Science*, *361*(6399), 273–277. https://doi.org/10.1126/science.aat2217

Stevens, L. A., Hewitt, I. J., Das, S. B., & Behn, M. D. (2018). Relationship between greenland ice sheet surface speed and modeled effective pressure. *Journal of Geophysical Research: Earth Surface*, *123*(9), 2258–2278. https://doi.org/10.1029/2017jf004581

Stuart, A., & Teckentrup, A. (2018). Posterior consistency for Gaussian process approximations of Bayesian posterior distributions. *Mathematics of Computation*, *87*(310), 721–753. https://doi.org/10.1090/mcom/3244

Thompson, J., Simons, M., & Tsai, V. C. (2014). Modeling the elastic transmission of tidal stresses to great distances inland in channelized ice streams. *The Cryosphere*, *8*(6), 2007–2029. https://doi.org/10.5194/tc-8-2007-2014

Tulaczyk, S., Kamb, W. B., & Engelhardt, H. F. (2000). Basal mechanics of Ice Stream B, west Antarctica: 1. Till mechanics. *Journal of Geophysical Research: Solid Earth*, *105*(B1), 463–481. https://doi.org/10.1029/1999JB900329

Van der Veen, C. J. (2013). *Fundamentals of glacier dynamics*. CRC Press. https://doi.org/10.1201/b14059

Vieli, A., & Payne, A. J. (2005). Assessing the ability of numerical ice sheet models to simulate grounding line migration. *Journal of Geophysical Research: Earth Surface*, *110*(F1), F01003. https://doi.org/10.1029/2004jf000202

Warburton, K., Hewitt, D., & Neufeld, J. (2020). Tidal grounding-line migration modulated by subglacial hydrology. *Geophysical Research Letters*, *47*(17), e2020GL089088. https://doi.org/10.1029/2020GL089088

Weertman, J. (1957). On the sliding of glaciers. *Journal of Glaciology*, *3*(21), 33–38. https://doi.org/10.3189/S0022143000024709

Zhao, C., Gladstone, R. M., Warner, R. C., King, M. A., Zwinger, T., & Morlighem, M. (2018). Basal friction of Fleming Glacier, Antarctica–part 1: Sensitivity of inversion to temperature and bedrock uncertainty. *The Cryosphere*, *12*(8), 2637–2652. https://doi.org/10.5194/tc-12-2637-2018

Zhu, Y., Zabaras, N., Koutsourelakis, P.-S., & Perdikaris, P. (2019). Physics-constrained deep learning for high-dimensional surrogate modeling and uncertainty quantification without labeled data. *Journal of Computational Physics*, *394*, 56–81. https://doi.org/10.1016/j.jcp.2019.05.024

Zoet, L. K., & Iverson, N. R. (2020). A slip law for glaciers on deformable beds. *Science*, *368*(6486), 76–78. https://doi.org/10.1126/science.aaz1183