

Towards Bayesian inference of friction parameters in ice sheets

Auteur : Remi, Arnaud

Promoteur(s) : Arnst, Maarten

Faculté : Faculté des Sciences appliquées

Diplôme : Master en ingénieur civil physicien, à finalité approfondie

Année académique : 2022-2023

URI/URL : <http://hdl.handle.net/2268.2/17210>

Avertissement à l'attention des usagers :

Tous les documents placés en accès ouvert sur le site le site MatheO sont protégés par le droit d'auteur. Conformément aux principes énoncés par la "Budapest Open Access Initiative"(BOAI, 2002), l'utilisateur du site peut lire, télécharger, copier, transmettre, imprimer, chercher ou faire un lien vers le texte intégral de ces documents, les disséquer pour les indexer, s'en servir de données pour un logiciel, ou s'en servir à toute autre fin légale (ou prévue par la réglementation relative au droit d'auteur). Toute utilisation du document à des fins commerciales est strictement interdite.

Par ailleurs, l'utilisateur s'engage à respecter les droits moraux de l'auteur, principalement le droit à l'intégrité de l'oeuvre et le droit de paternité et ce dans toute utilisation que l'utilisateur entreprend. Ainsi, à titre d'exemple, lorsqu'il reproduira un document par extrait ou dans son intégralité, l'utilisateur citera de manière complète les sources telles que mentionnées ci-dessus. Toute utilisation non explicitement autorisée ci-avant (telle que par exemple, la modification du document ou son résumé) nécessite l'autorisation préalable et expresse des auteurs ou de leurs ayants droit.



UNIVERSITÉ DE LIÈGE

*Towards Bayesian inference of friction
parameters in ice sheets*

Master's thesis submitted in partial fulfillment of the requirements
for the degree of Master of Science in Engineering Physics

Author
Arnaud REMI

Supervisor
Pr. Maarten ARNST

Contents

Introduction	4
Ice sheets in the climate system	4
Objectives of this thesis	6
I Forward problem	8
1 Continuum mechanics	8
1.1 Notations	8
1.2 Conservation laws	9
1.2.1 Reynolds transport theorem	9
1.2.2 Mass conservation	10
1.2.3 Linear momentum conservation	10
1.2.4 Angular momentum conservation	11
1.3 Constitutive equations	11
2 Ice-sheet governing equations	11
2.1 Full Stokes problem	11
2.1.1 Governing equations	12
2.1.2 Boundary conditions	14
2.1.3 Ice thickness equation	16
2.2 Shallow-shelf approximation	17
2.2.1 Implication on mass conservation	17
2.2.2 Implication on linear momentum conservation	17
3 Modelling the friction field	19
3.1 Notions of random fields	19
3.1.1 Gaussian random fields	19
3.1.2 Karhunen-Loève expansion	23
3.2 Friction field as a random field	23
3.2.1 Maximum entropy principle	23
3.2.2 Application to the friction field	24
3.2.3 Transformation from a Gaussian random field	24
3.2.4 Realization of the friction field	25
4 Surrogate model	26
4.1 Polynomial chaos surrogate model	26
4.2 Application to the SSA ice-flow model	26
II Inverse problem	30

5	Generalities on inverse problem	30
5.1	Context and notations	30
5.2	Deterministic framework	31
5.3	Bayesian frame work	32
5.4	Litterature review of inverse problem in glaciology	33
6	Sampling methods	36
6.1	Metropolis-Hastings algorithm	36
6.2	Adaptive methods	37
7	Neural posterior estimation	38
7.1	Variational inference	38
7.2	Normalizing flows	39
7.2.1	General principle	39
7.2.2	Autoregressive flows	40
8	Formulation of the inverse problem	42
8.1	Friction inversion	42
8.1.1	Prior and likelihood	42
8.1.2	Methods	43
8.2	Methodology	43
8.2.1	Validation phase	44
8.2.2	Testing phase	45
III	Results	46
9	Validation phase	46
9.1	Sampling methods	46
9.1.1	Metropolis-Hastings algorithm	46
9.1.2	Adaptive-Metropolis algorithm	48
9.2	Neural posterior estimation	51
9.2.1	Number of transformations	51
9.2.2	Size of the dataset	53
10	Testing phase	56
10.1	Adaptive Metropolis	56
10.2	Neural posterior estimation	58
10.3	AM and NPE results	60
	Conclusion	63
A	Parameters in the SSA equations	65
B	Complementary results	65

Acknowledgments

I would like to thank my supervisor Maarten Arnst, as well as Thomas Gregov and Romin Tomasetti for proposing a thesis topic that suited my interests, for their numerous advice, feedback, and their interest in my work, and for providing a stimulating working environment. In particular, I thank Thomas for the many discussions about ice sheets and the progress of this work.

I thank François Rozet for his assistance in getting started with the Lampe and Zuko libraries for normalizing flows.

I express my gratitude to Benjamin Dewals and Gilles Louppe for agreeing to be part of this jury.

Introduction

Ice sheets in the climate system

The Earth climate system is a complex and global system which includes five major components: the atmosphere, the lithosphere, the hydrosphere, the cryosphere, and the biosphere. All these components are interconnected such that a change in the state of one influences the state of the others. This field is of interest of many scientists as it contains tipping points. Tipping points are states which, when slightly perturbed, bifurcate towards another state. As the climate system is global, a bifurcation related to the state of one of its components will affect the others which could lead to a cascade of change of states in the global system. The identified tipping points in the climate system are represented in Fig. 1.

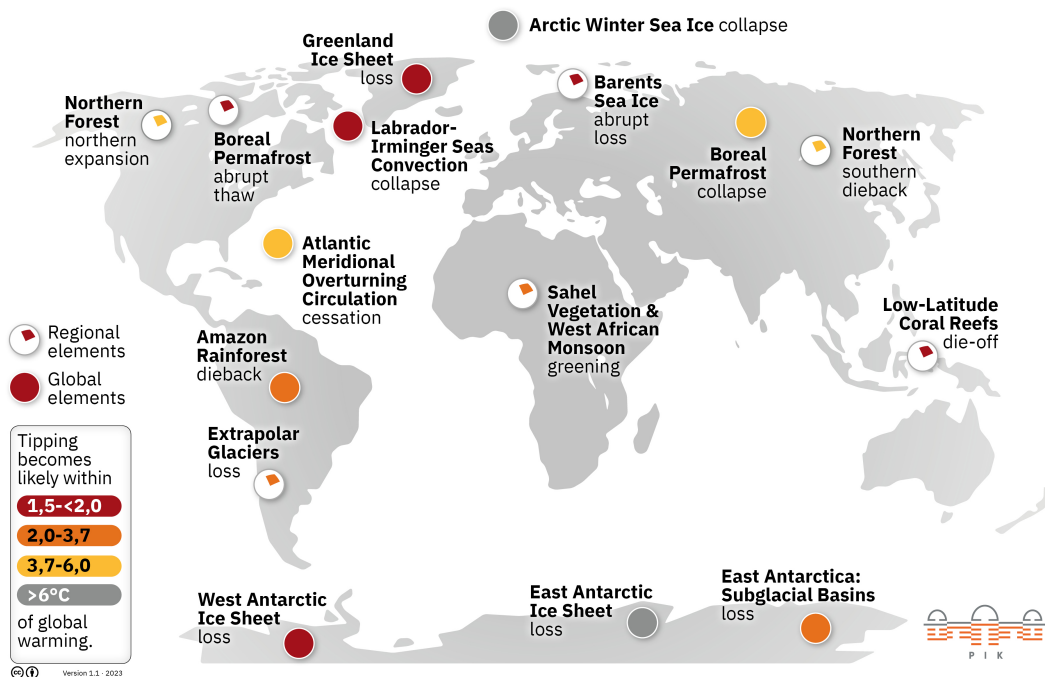


Figure 1: Tipping points in the climate system. Taken from [Potsdam \(2022\)](#), based on [Armstrong McKay et al. \(2022\)](#).

As it can be seen, some of the most critical tipping points concern the cryosphere, especially West Antarctica and Greenland ice sheets. Ice sheets are large-scale masses of ice that flow under the influence of their own weight. Two zones can be distinguished in an ice sheet: the grounded part where ice is in contact with the bedrock, and the floating part where ice is floating above water. These two zones are separated by the grounding line. This is illustrated in Fig. 2.




Figure 2: Ice-sheet schematic. Taken from [Huybrechts \(2009\)](#)

These ice sheets are characterized by several instabilities which lead to large ice loss. Among these instabilities figure

1. the marine ice-sheet instability (MISI), which concerns the potential irreversible retreat of the grounding line (*i.e.* the separation between the grounded part and the floating part of marine ice-sheets), which leads to ice losses ([Schoof \(2007\)](#));
2. the marine ice-cliff instability (MICI), which concerns the instability of the floating parts of the ice sheet that eventually fracture ([Pollard et al. \(2015\)](#)).

Currently, sea-level rise is dominated by water thermal expansion. However, due to tipping points, sea-level rise due to losses in marine ice sheet are expected to gain in relative importance, which makes of marine ice sheets key components in future sea-level rise predictions. Marine ice-sheet contribution to the sea-level rise is not only expected to be important, but is also highly uncertain (see Fig. 3). These uncertainties have several causes: (*i*) some parts of the physics are still poorly understood (*e.g.* MICI, nature of the friction between ice and the bedrock) ; (*ii*) some parameters cannot be measured and are uncertain (*e.g.* friction parameters). There is therefore an interest in assessing and reducing these uncertainties.

Melting Ice and Sea Level Rise

As Greenland and Antarctica lose ice and mountain glaciers melt, they add to sea level rise. These charts show how much each is projected to contribute depending how much greenhouse gas emissions rise.

ICE SHEET AND GLACIER LOSS

Mass loss as sea level equivalent, compared to 1986-2005 average

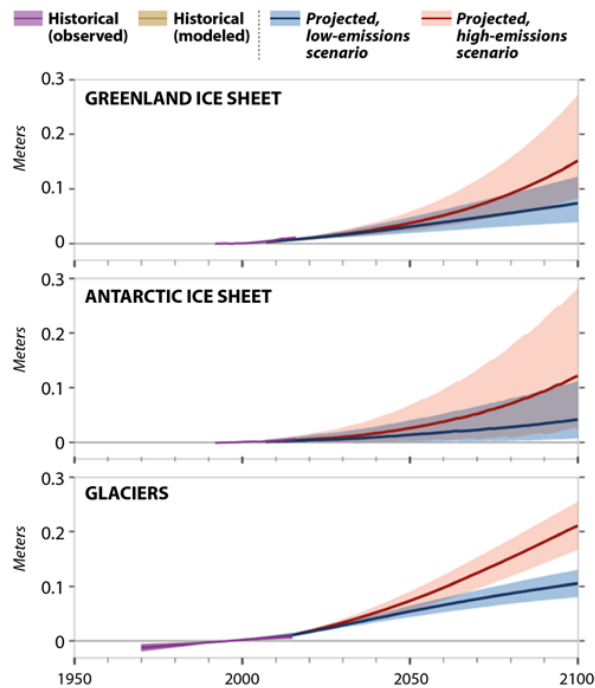


Figure 3: Contribution of Greenland ice sheet, Antarctic ice sheet, and glaciers to sea-level rise prediction. Taken from [IPCC \(2019\)](#).

Objectives of this thesis

This work is concerned in the identification of uncertain friction parameters from partial observations of the velocity. Such problems are referred to as *inverse problems*, as the equations that model the physics are travelled in the inverse direction, that is, from the solution to some of the parameters that led to the solution. Typically, such inverse problems are computationally expensive and mathematically ill-posed.

This field of research gains in interest, and is an active research area. In this thesis, this glaciological inverse problem will be solved in the probabilistic framework that allows to quantify the uncertainty on the identified parameters, which is crucial for sea-level rise predictions. Classical state-of-the-art methods will be implemented to solve this inverse problem. On the top of that, recent deep learning methods will be used to identify the potential contribution of deep learning to this research field, which is a strong novelty of this thesis.

This thesis will be divided in three parts

1. **Forward problem formulation.** Basic notions of continuum mechanics will be reviewed in Sec. 1. These notions will be particularized for the case of bidimensional marine ice sheets in Sec. 2, and a reduced model model will be described. The friction field will be modelled in Sec. 3, and a computationally inexpensive surrogate model will be built in Sec. 4;
2. **Inverse problem formulation.** The current state-of-the-art for the inverse problem in the

glaciology context will be described in Sec. 5, and methods for solving such inverse problem in a probabilistic framework will be reviewed in Sec. 6 and 7. The way these methods will be applied for solving the inverse problem is explained in Sec. 8;

3. **Results.** The expressiveness of the implemented methods will be assessed by solving a simple bidimensional problem in Sec. 9. Once the expressiveness of the methods are assessed, they are used to solve a more realistic inverse problem where the friction field has to be inferred from local observations of the ice velocity in Sec. 10.

Finally, conclusions will be drawn concerning the potential use of deep learning into the field, as a function of the results.

Part I

Forward problem

In this part, the forward problem will be formulated. The physical principles that are needed for the description of ice flow will be described starting from general continuum mechanics equations, towards specific governing equations of ice flows. Reduced-order models will then be introduced, and the shallow-shelf approximation (SSA) will be detailed. Uncertainty in the friction field that appears in these governing equations will be introduced by modelling the friction as a random field. A parametrization of this random field will be developed, which will make the friction field suited for the establishment of an input-output surrogate model.

1 Continuum mechanics

1.1 Notations

Let Ω be a material volume that deforms continuously in a time interval $T = [t_0, t_f]$. Let its configuration at time $t = t_0$ be a reference configuration Ω_0 and its configuration at a time $t \in T$ be the current configuration $\Omega(t)$. This situation is illustrated in Fig. 4.




Figure 4: Schematic of the deformation of a continuous material volume Ω in time.

At any current time $t \in T$, and with respect to the origin O , the position of a particle is noted $\mathbf{x}(t)$ and is linked with its reference position $\mathbf{x} = \mathbf{x}(t_0)$ through the one-to-one mapping ϕ in the following way:

$$\mathbf{x}(t) = \phi(\mathbf{X}, t) \quad \Leftrightarrow \quad \mathbf{X} = \phi^{-1}(\mathbf{x}(t)). \quad (1.1)$$

One can therefore distinguish the *Eulerian* quantities that are attached to a point of space $\mathbf{x}(t)$ and the *Lagrangian* quantities which are attached to a particle that was at \mathbf{X} at $t = t_0$. The velocity of a particle that is at position \mathbf{x} at a time t is noted \mathbf{u} , and can be obtained by differentiating the

current position with respect to time:

$$\mathbf{u}(\mathbf{x}, t) = \frac{d\mathbf{x}}{dt} = \frac{d\phi(\mathbf{X}, t)}{dt} = \frac{\partial\phi(\mathbf{X}, t)}{\partial t}. \quad (1.2)$$

The strain rate tensor \mathbf{D} is a symmetric second-order tensor defined as the symmetric part of the velocity gradient. It is written as

$$\mathbf{D} = \frac{1}{2} (\nabla_{\mathbf{x}}\mathbf{u} + (\nabla_{\mathbf{x}}\mathbf{u})^T). \quad (1.3)$$

The forces that are applied to the material volume Ω are separated into two categories. The volumic forces (or body forces) \mathbf{f} are applied to particles located inside the material volume ($\mathbf{x} \in \Omega$), and the surface forces \mathbf{t} are applied to particles that are located at the material volume boundaries ($\mathbf{x} \in \partial\Omega$). Let us consider a facet oriented towards a direction \mathbf{n} , inside the domain, on which is applied a traction force \mathbf{t} . The Cauchy stress tensor $\boldsymbol{\sigma}$ is a second-order tensor that is linked with \mathbf{t} through

$$\boldsymbol{\sigma} \cdot \mathbf{n} = \mathbf{t}. \quad (1.4)$$

1.2 Conservation laws

In this section, the conservation laws that will be necessary for the description of the ice flow will be established.

1.2.1 Reynolds transport theorem

The Reynolds transport theorem is used to differentiate an integrated quantity over a domain $\Omega(t)$ with boundary $\partial\Omega(t)$ that change with time. This is not a conservation law in itself, but rather a tool that is used to derive the conservation laws in their local form. In a one dimensional case, let $\Omega^{1D} =]x_1(t), x_2(t)[$ be the domain over which a quantity $F(x, t)$ is integrated. Its total derivative can be computed by using the Leibnitz integration rule

$$\frac{d}{dt} \int_{x_1(t)}^{x_2(t)} F(x, t) dx = \int_{x_1(t)}^{x_2(t)} \frac{\partial F(x, t)}{\partial t} dx + \left[F(x, t) \frac{dx}{dt} \right]_{x_1}^{x_2}. \quad (1.5)$$

This can be applied to a material volume Ω with boundary $\partial\Omega$ that deform continuously in a time interval $T = [t_0, t_f]$. The multidimensional generalization of Leibnitz integration rule can be written as

$$\frac{d}{dt} \int_{\Omega(t)} F(\mathbf{x}, t) dV = \int_{\Omega(t)} \frac{\partial F}{\partial t} dV + \int_{\partial\Omega(t)} F \mathbf{u} \cdot \mathbf{n} dS. \quad (1.6)$$

By using the divergence theorem, one obtains

$$\frac{d}{dt} \int_{\Omega(t)} F(\mathbf{x}, t) dV = \int_{\Omega(t)} \left(\frac{\partial F}{\partial t} + \nabla_{\mathbf{x}} \cdot (F \mathbf{u}) \right) dV, \quad (1.7)$$

which is known as the Reynolds transport theorem.

1.2.2 Mass conservation

Let Ω be a material volume of mass M and density ρ , that deforms continuously in a time interval $T = [t_0, t_f]$. At any time $t \in T$, mass conservation can be written as

$$\frac{d}{dt}M(t) = 0 \quad (1.8)$$

$$\Leftrightarrow \frac{d}{dt} \int_{\Omega(t)} \rho \, dV = 0 \quad (1.9)$$

$$\Leftrightarrow \int_{\Omega(t)} \left(\frac{\partial \rho}{\partial t} + \nabla_{\mathbf{x}} \cdot (\rho \mathbf{u}) \right) dV = 0 \quad (1.10)$$

$$\Leftrightarrow \int_{\Omega(t)} \left(\frac{\partial \rho}{\partial t} + \rho \nabla_{\mathbf{x}} \cdot \mathbf{u} + \mathbf{u} \cdot \nabla_{\mathbf{x}} \rho \right) dV = 0, \quad (1.11)$$

where the Reynolds transport theorem has been used. Eq. (1.11) is valid for any material volume Ω ; therefore one can write the mass conservation in a local form

$$\frac{\partial \rho}{\partial t} + \rho \nabla_{\mathbf{x}} \cdot \mathbf{u} + \mathbf{u} \cdot \nabla_{\mathbf{x}} \rho = 0, \quad \mathbf{x} \in \Omega, \quad t \in T. \quad (1.12)$$

If the material is incompressible, one has

$$\frac{\partial \rho}{\partial t} + \mathbf{u} \cdot \nabla_{\mathbf{x}} \rho = 0, \quad \mathbf{x} \in \Omega, \quad t \in T, \quad (1.13)$$

such that the mass conservation equation is reduced to the following divergence-free condition

$$\nabla_{\mathbf{x}} \cdot \mathbf{u} = 0, \quad \mathbf{x} \in \Omega, \quad t \in T. \quad (1.14)$$

1.2.3 Linear momentum conservation

The second law of Newton states that for any material point of mass m and velocity \mathbf{u} subjected to an external force \mathbf{f} ,

$$\mathbf{f} = \frac{d}{dt}(m\mathbf{u}). \quad (1.15)$$

This can be generalized for a material volume Ω subjected to volumic forces \mathbf{f} and surface forces \mathbf{t}

$$\frac{d}{dt} \int_{\Omega} \rho \mathbf{u} \, dV = \int_{\Omega} \rho \mathbf{f} \, dV + \int_{\partial\Omega} \mathbf{t} \, dS. \quad (1.16)$$

By definition of the Cauchy stress tensor $\boldsymbol{\sigma}$, one has $\boldsymbol{\sigma} \cdot \mathbf{n} = \mathbf{t}$ for all $\mathbf{x} \in \partial\Omega$, where \mathbf{n} is the outward pointing normal vector to $\partial\Omega$. Applying Reynolds transport theorem to the left hand side of Eq. (1.16) and the divergence theorem to its right hand side, one obtains:

$$\begin{aligned} \int_{\Omega} \left(\frac{\partial(\rho \mathbf{u})}{\partial t} + \mathbf{u} \cdot \nabla_{\mathbf{x}}(\rho \mathbf{u}) \right) dV &= \int_{\Omega} (\rho \mathbf{f} + \nabla_{\mathbf{x}} \cdot \boldsymbol{\sigma}) dV \\ \Leftrightarrow \int_{\Omega} \left(\mathbf{u} \frac{d\rho}{dt} + \rho \frac{d\mathbf{u}}{dt} \right) dV &= \int_{\Omega} (\rho \mathbf{f} + \nabla_{\mathbf{x}} \cdot \boldsymbol{\sigma}) dV. \end{aligned} \quad (1.17)$$

Since Eq. (1.17) is valid for any material volume, one can write it in a local form:

$$\mathbf{u} \frac{d\rho}{dt} + \rho \frac{d\mathbf{u}}{dt} = \rho \mathbf{f} + \nabla_{\mathbf{x}} \cdot \boldsymbol{\sigma}, \quad \mathbf{x} \in \Omega, \quad t \in T. \quad (1.18)$$

1.2.4 Angular momentum conservation

Let us define the angular momentum of a material volume Ω with respect to the origin O by

$$\mathbf{H}_O = \int_{\Omega} \mathbf{x} \times (\rho \mathbf{u}) \, dV. \quad (1.19)$$

In the case where there is no internal angular momentum, conservation of angular momentum states that the variation of angular momentum of the material volume Ω is due to applied (volumic and surfacic) angular momenta, that is,

$$\begin{aligned} \frac{d}{dt} \int_{\Omega} \mathbf{x} \times (\rho \mathbf{u}) \, dV &= \int_{\Omega} \mathbf{x} \times (\rho \mathbf{f}) \, dV + \int_{\partial\Omega} \mathbf{x} \times \mathbf{t} \, dS \\ &= \int_{\Omega} \mathbf{x} \times (\rho \mathbf{f}) \, dV + \int_{\Omega} \nabla_{\mathbf{x}} \cdot (\mathbf{x} \times \boldsymbol{\sigma}) \, dV. \end{aligned} \quad (1.20)$$

By using simple algebra, one could show from Eq. (1.20) that the Cauchy stress tensor should be symmetric, that is,

$$\boldsymbol{\sigma} = \boldsymbol{\sigma}^T, \quad \mathbf{x} \in \Omega, \quad t \in T. \quad (1.21)$$

1.3 Constitutive equations

The above conservation equations are not sufficient to describe the dynamics of the material volume Ω . The relationship between the Cauchy stress tensor $\boldsymbol{\sigma}$ and the strain rate tensor \mathbf{D} has to be modelled. Such relationships are referred to as constitutive equations.

In general, the Cauchy stress tensor is written as

$$\boldsymbol{\sigma} = \frac{1}{3} \text{tr}(\boldsymbol{\sigma}) \mathbf{I} + \boldsymbol{\tau}, \quad (1.22)$$

where $\text{tr}(\boldsymbol{\sigma}) \mathbf{I} / 3$ is a pressure term that is often written as $-p \mathbf{I}$, and $\boldsymbol{\tau}$ is the deviatoric part of the Cauchy stress tensor. In the case of viscous material, constitutive equations link the deviatoric part of the stress tensor to the strain-rate tensor. The choice of a constitutive equation must reflect the rheology of the material and is therefore problem-dependent.

2 Ice-sheet governing equations

In this section, the equations that have been expressed in a general way in Sec. 1 will be particularized to the case of marine ice sheet, based on [Greve and Blatter \(2009\)](#) and [Schoof and Hewitt \(2013\)](#). The full Stokes model will first be introduced, and reduced order models will then be described.

2.1 Full Stokes problem

Let us consider a bidimensional ice sheet in the x - z plane, x being the horizontal axis and z the vertical one. A schematic of a 2D ice-sheet is shown in Fig. 5.

As it is illustrated, the ice sheet is separated into two parts: (i) the part where ice is grounded (*i.e.* in contact with the bedrock) $x \in \Omega_g$; (ii) the part where ice is floating $x \in \Omega_f$. The whole spatial domain Ω is the union of the grounded part and the floating part: $\Omega = \Omega_g \cup \Omega_f$. The ice divide shows a zero velocity. There are three types of media which are in contact with the ice sheet. Let us denote the interface between the ice sheet and the bedrock by Γ_b , the interface between the ice sheet and water by Γ_w , and the interface between the ice sheet and the air (that is commonly referred to as the free interface) by Γ_f . The separations between the interfaces will be useful to establish boundary conditions. In the following, the upper part of the ice-sheet will be denoted by $s(x, t)$, the lower part by $l(x, t)$, and the ice thickness will be denoted by $h(x, t) = s(x, t) - b(x, t)$. The bedrock profile will be noted $b(x)$, and is assumed to not vary with time. Thus, $l(x, t) = b(x)$ for $x \in \Omega_g$. The velocity vector \mathbf{u} can be decomposed in horizontal and vertical directions as $\mathbf{u} = ue_x + we_z$.




Figure 5: Ice-sheet schematic in 1D. The arrows represent typical flow lines. The ice sheet is separated into two zones : the grounded part where the ice is in contact with the bed, and the floating part where the ice floats. This schematic is inspired from [Greve and Blatter \(2009\)](#).

2.1.1 Governing equations

Let us recall here the local form of the balance laws that have been established in the previous section:

$$\begin{cases} \frac{d\rho}{dt} + \rho \nabla_{\mathbf{x}} \cdot \mathbf{u} = 0, & \mathbf{x} \in \Omega, \quad t \in T, \\ \mathbf{u} \frac{d\rho}{dt} + \rho \frac{d\mathbf{u}}{dt} = \rho \mathbf{f} - \nabla_{\mathbf{x}} p + \nabla_{\mathbf{x}} \cdot \boldsymbol{\tau}, & \mathbf{x} \in \Omega, \quad t \in T. \end{cases} \quad (2.1)$$

Since the scale of the problem is large ($L^* = 1000$ km), fluctuations of the density are neglected and ice is considered incompressible such that $d\rho/dt = 0$. In the case of ice-sheets, the volumic force is given by the volumic weight (Schoof and Hewitt (2013)). That is

$$\mathbf{f} = \rho \mathbf{g}, \quad (2.2)$$

where \mathbf{g} is the gravity ($|\mathbf{g}| \approx 9.81$ m/s²). A constitutive law for $\boldsymbol{\tau}$ is still required to close the system. In glaciology, one typically consider viscous relationship between $\boldsymbol{\tau}$ and \mathbf{D}

$$\boldsymbol{\tau} = 2\eta(\mathbf{D})\mathbf{D}, \quad (2.3)$$

where η is a function of \mathbf{D} which highlights the fact that ice is considered as a non-Newtonian fluid. Most of the models use a power law for η :

$$\eta = \frac{1}{2} B D^{-1+1/n}, \quad (2.4)$$

where n is a constant, $D = \sqrt{\mathbf{D} : \mathbf{D}/2}$, and B captures the temperature dependency through an Arrhenius law ($B \sim \exp(1/T)$). In the rest of this work, the temperature effects will be disregarded and B will be a constant. This power law is known as the Glen's flow law. Experimental results suggest that $n = 3$ (Glen (1955)). The system of equations (2.1) can therefore be written

$$\begin{cases} \nabla_{\mathbf{x}} \cdot \mathbf{u} = 0, & \mathbf{x} \in \Omega, \quad t \in T, \\ \rho \frac{d\mathbf{u}}{dt} = \rho \mathbf{g} - \nabla_{\mathbf{x}} p + 2\nabla_{\mathbf{x}} \cdot \eta \mathbf{D}, & \mathbf{x} \in \Omega, \quad t \in T. \end{cases} \quad (2.5)$$

Now, let us show that the inertia term can be neglected with respect to the pressure gradient. For this purpose, one will introduce typical quantities which will be used to approximate the different terms of System (2.5) as in Greve and Blatter (2009)

$$\text{Typical horizontal length scale } L^* = 1000 \text{ km}, \quad (2.6)$$

$$\text{Typical vertical length scale } H^* = 1 \text{ km}, \quad (2.7)$$

$$\text{Typical horizontal velocity } U^* = 10^{-6} \text{ m/s}, \quad (2.8)$$

$$\text{Typical vertical velocity } V^* = 10^{-9} \text{ m/s}, \quad (2.9)$$

$$\text{Typical pressure } P^* = \rho g H^* \approx 10 \text{ MPa}, \quad (2.10)$$

$$\text{Typical time scale } t^* = L^*/U^* = 10^{12} \text{ s}. \quad (2.11)$$

The Froude number Fr is the ratio between the inertia and the pressure terms, that is

$$\text{Fr} = \frac{\rho U^*/t^*}{P^*/L^*} \approx 10^{-16}, \quad (2.12)$$

which confirms that the inertia term can be neglected. Therefore, System (2.5) becomes

$$\begin{cases} \nabla_{\mathbf{x}} \cdot \mathbf{u} = 0, & \mathbf{x} \in \Omega, \quad t \in T, \\ -\nabla_{\mathbf{x}} p + 2\nabla_{\mathbf{x}} \cdot \eta \mathbf{D} = -\rho \mathbf{g}, & \mathbf{x} \in \Omega, \quad t \in T, \end{cases} \quad (2.13)$$

which is known as the full Stokes model.

2.1.2 Boundary conditions

System (2.13) requires boundary conditions to be solved. These boundary conditions can be separated into the kinematic boundary conditions (that are applied to the velocity), and the dynamic boundary conditions (that are applied to the stress).

Kinematic boundary conditions:

In the following, the extreme left of the domain (which will correspond to the location $x = 0$) will be assumed to be the ice divide. Therefore, a homogeneous Dirichlet boundary condition is applied for the velocity

$$\mathbf{u} = \mathbf{0}, \quad \text{at } x = 0. \quad (2.14)$$

The upper surface is an interface that moves with velocity \mathbf{u}_s (see Fig. 6) that could be expressed as

$$F_s(x, z, t) \equiv z - s(x, t) = 0. \quad (2.15)$$

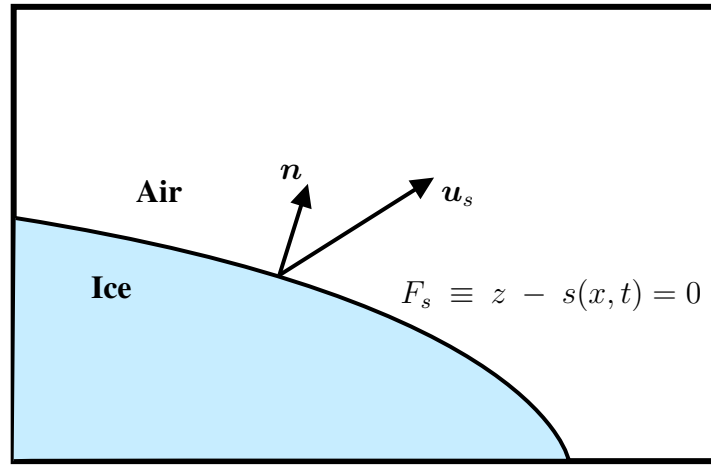


Figure 6: Free surface interface F_s that moves with velocity \mathbf{u}_s . Adapted from Greve and Blatter (2009).

Boundary condition for the ice flux at the interface can be expressed as

$$(\mathbf{u}_s - \mathbf{u}) \cdot \mathbf{n} = a_n, \quad (2.16)$$

where \mathbf{u}_s is the interface velocity, and a_n is the net accumulation rate of snow normal to the interface. Moreover, the definition of Eq. (2.15) implies that the total derivative of F_s should be zero, that is

$$\frac{dF_s}{dt} = \frac{\partial F_s}{\partial t} + \mathbf{u}_s \cdot \nabla_x F_s = 0. \quad (2.17)$$

By combining Eq. (2.16) and Eq. (2.17), one obtains the following relationship

$$\frac{\partial F_s}{\partial t} + \mathbf{u} \cdot \nabla_x F_s = -|\nabla_x F_s| a_n \quad (2.18)$$

$$\Leftrightarrow -\frac{\partial s}{\partial t} - u \frac{\partial s}{\partial x} + w = -\sqrt{1 + \left(\frac{\partial s}{\partial x}\right)^2} a_n, \quad \text{at } z = s(x, t). \quad (2.19)$$

The net accumulation rate a is assumed to be vertical, such that

$$a_n = \frac{a}{\sqrt{1 + \left(\frac{\partial s}{\partial x}\right)^2}}. \quad (2.20)$$

Therefore, one obtains

$$\frac{\partial s}{\partial t} + u \frac{\partial s}{\partial x} - w = a, \quad \text{at } F_s = 0. \quad (2.21)$$

The grounded surface is a fixed interface that is shown in Fig. 7

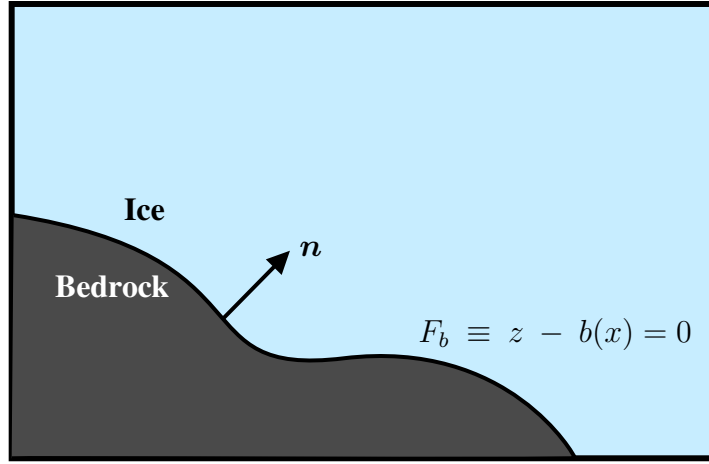


Figure 7: Fixed grounded surface interface F_b . Adapted from [Greve and Blatter \(2009\)](#).

The boundary condition associated with the ice–bedrock interface is a impermeability condition, that is

$$\mathbf{u} \cdot \mathbf{n} = 0. \quad (2.22)$$

The normal \mathbf{n} can be expressed as

$$\mathbf{n} = \frac{\nabla_x F_b}{|\nabla_x F_b|} = \frac{-\partial b / \partial x \mathbf{e}_x + \mathbf{e}_z}{\sqrt{1 + (\partial b / \partial x)^2}}. \quad (2.23)$$

By substitution of Eq. (2.23) in Eq. (2.22), one finds

$$-u \frac{\partial b}{\partial x} + w = 0, \quad \text{at } F_b = 0. \quad (2.24)$$

Dynamic boundary conditions:

At the free surface, *i.e.* at $z = s(x, t)$, or equivalently at Γ_f , a stress free condition is assumed such that

$$\boldsymbol{\sigma} \cdot \mathbf{n} = 0, \quad \text{at } \Gamma_f. \quad (2.25)$$

As represented in Fig. 8, the part of the ice-sheet that is in contact with water (the ice shelf) is submitted to water pressure such that

$$\boldsymbol{\sigma} \cdot \mathbf{n} = \rho_w g z \mathbf{n}, \quad \text{at } \Gamma_w. \quad (2.26)$$




Figure 8: Floating interface. The water pressure acts on the ice shelf perpendicularly.

At the base of the ice sheet, *i.e.* at $z = l(x, t)$, or equivalently at Γ_b , a tangential basal stress is assumed such that

$$\boldsymbol{\sigma} \cdot \mathbf{t}_b = \boldsymbol{\tau}_b, \quad \text{at } \Gamma_b, \quad (2.27)$$

where \mathbf{t}_b is a unitary vector tangential to the bedrock profile. Various friction laws exist for $\boldsymbol{\tau}_b$. In the following, this is the Weertman friction law that will be used. The Weertman friction law (Weertman (1957)) expresses the basal stress as a power law of the basal velocity, that is

$$\boldsymbol{\tau}_b = -\nu_b |\mathbf{u}_b|^{m-1} \mathbf{u}_b, \quad (2.28)$$

where ν_b is a friction coefficient, m is a constant, and \mathbf{u}_b is the basal velocity. Typically, $m = 1/3$ is chosen. A lot of other models for the friction law could have been used, as the Budd friction law (Budd et al. (1984)) or the Schoof friction law (Schoof (2005)) but they will not be detailed in this thesis.

2.1.3 Ice thickness equation

Now that the governing equations as well as the boundary conditions have been described, one can establish an equation for the ice thickness $h(x, t)$. Let us integrate the mass conservation equation along the z -axis:

$$\frac{\partial u}{\partial x} + \frac{\partial w}{\partial z} = 0 \quad (2.29)$$

$$\Leftrightarrow \int_{b(x)}^{s(x,t)} \frac{\partial u}{\partial x} dz + \int_{b(x)}^{s(x,t)} \frac{\partial w}{\partial z} dz = 0 \quad (2.30)$$

$$\Leftrightarrow \int_{b(x)}^{s(x,t)} \frac{\partial u}{\partial x} dz + w|_{z=s} - w|_{z=b} = 0. \quad (2.31)$$

By using the Leibniz rule Eq. (1.5), one obtains

$$\frac{\partial}{\partial x} \int_{b(x)}^{s(x,t)} u \, dz - \left(u|_{z=s} \frac{\partial s}{\partial x} - u|_{z=b} \frac{\partial b}{\partial x} \right) + w|_{z=s} - w|_{z=b} = 0. \quad (2.32)$$

By injecting the boundary conditions on the free surface and on the grounded surface, it comes

$$\frac{\partial s}{\partial t} + \frac{\partial}{\partial x} \int_{b(x)}^{s(x,t)} u \, dz = a. \quad (2.33)$$

By noticing that $\partial s/\partial t = \partial h/\partial t$, the ice thickness equation finally writes

$$\frac{\partial h}{\partial t} + \frac{\partial}{\partial x} \int_{b(x)}^{s(x,t)} u \, dz = a. \quad (2.34)$$

2.2 Shallow-shelf approximation

The Full Stokes model provides a mathematical description of the physics involved in ice-flow. However, it is computationally expensive to solve. In this subsection, the shallow-shelf approximation (SSA) will be presented.

Initially proposed by [Morland \(1987\)](#) and [MacAyeal \(1989\)](#), the SSA considers that ice moves in a sliding mode, that is,

$$\left| \frac{\partial u}{\partial x} \right| \gg \left| \frac{\partial u}{\partial z} \right|, \quad (2.35)$$

and that the vertical stress is hydrostatic. Formally, the SSA is valid in the floating part of the ice sheet, and in grounded regions close to the grounding line. It is a strong model assumption. The implications of the SSA on the mass conservation equation as well as on the momentum conservation will now be derived.

2.2.1 Implication on mass conservation

As the ice sheet is considered to move according to a sliding mode, the variation of horizontal the velocity u is assumed to be negligible in the vertical direction, that is

$$\frac{\partial u}{\partial z} \approx 0. \quad (2.36)$$

The integrated form of the mass conservation equation Eq. (2.34) becomes

$$\frac{\partial h}{\partial t} + \frac{\partial}{\partial x} (hu) = a. \quad (2.37)$$

2.2.2 Implication on linear momentum conservation

Let us project the linear momentum equation on x and z axis:

$$\begin{cases} -\frac{\partial p}{\partial x} + \frac{\partial \tau_{xx}}{\partial x} + \frac{\partial \tau_{xz}}{\partial z} = 0, \\ -\frac{\partial p}{\partial z} + \frac{\partial \tau_{xz}}{\partial x} + \frac{\partial \tau_{zz}}{\partial z} = \rho g. \end{cases} \quad (2.38)$$

In ice sheets, $|\tau_{xx}|/L^* \gg |\tau_{xz}|/H^*$ and $|\tau_{zz}|/H^* \gg |\tau_{xz}|/L^*$, which suggests to simplify System (2.38) into

$$\begin{cases} -\frac{\partial p}{\partial x} + \frac{\partial \tau_{xx}}{\partial x} = 0, \\ -\frac{\partial p}{\partial z} + \frac{\partial \tau_{zz}}{\partial z} = \rho g. \end{cases} \quad (2.39)$$

By integrating along the z -axis, one has

$$p - \tau_{zz} = \rho g(s - z). \quad (2.40)$$

Since $\tau_{zz} = -\tau_{xx}$ (τ is deviatoric), Eq. (2.40) can be substituted in the first equation of System (2.39) which leads to

$$-2 \frac{\partial \tau_{xx}}{\partial x} = \rho g \frac{\partial s}{\partial x}. \quad (2.41)$$

By integrating Eq. (2.41) along z , one has

$$2 \int_b^s \frac{\partial \tau_{xx}}{\partial x} dz = \int_b^s \rho g \frac{\partial s}{\partial x} dz \quad (2.42)$$

$$\Leftrightarrow 2 \frac{\partial}{\partial x} \int_b^s \tau_{xx} dz - [\tau_{xx}]_b^s = \rho g h \frac{\partial s}{\partial x} \quad (2.43)$$

$$\Leftrightarrow 2 \frac{\partial}{\partial x} \left\{ h A^{-n} \left| \frac{\partial u}{\partial x} \right|^{m-1} \frac{\partial u}{\partial x} \right\} + \mathbb{I}_g \tau_b = \rho g h \frac{\partial s}{\partial x}, \quad (2.44)$$

where $A = B^{-1/n}$, and

$$\mathbb{I}_g = \begin{cases} 1, & x \in \Omega_g, \\ 0, & \text{otherwise.} \end{cases} \quad (2.45)$$

Moreover, in the floating part (*i.e.* at $x \in \Omega_f$), Archimède's principle writes

$$\rho h = -\rho_w l \Leftrightarrow s = \left(1 - \frac{\rho_w}{\rho}\right) h, \quad (2.46)$$

which allows to write

$$2A^{-n} \left| \frac{\partial u}{\partial x} \right|^{m-1} \frac{\partial u}{\partial x} = \rho g \left(1 - \frac{\rho_w}{\rho}\right) h. \quad (2.47)$$

Finally, the SSA equations write

$$\boxed{\begin{cases} \frac{\partial h}{\partial t} + \frac{\partial}{\partial x}(hu) = a, & x \in \Omega, t \in T, \\ 2 \frac{\partial}{\partial x} \left\{ h A^{-n} \left| \frac{\partial u}{\partial x} \right|^{m-1} \frac{\partial u}{\partial x} \right\} + \mathbb{I}_g \tau_b = \rho g h \frac{\partial s}{\partial x}, & x \in \Omega_g, t \in T, \\ 2A^{-n} \left| \frac{\partial u}{\partial x} \right|^{m-1} \frac{\partial u}{\partial x} = \rho g \left(1 - \frac{\rho_w}{\rho}\right) h, & x \in \Omega_f, t \in T, \\ + \text{Boundary \& Initial conditions.} \end{cases}} \quad (2.48)$$

In what follows, a numerical implementation of the SSA equations with the Weertman friction law will be used. For more informations on this numerical implementation, please refer to [Bosten et al. \(2019\)](#); [Gregov et al. \(2022\)](#). Parameters of the SSA equations and the geometry of the bedrock are summarized in Appendix A.

3 Modelling the friction field

Spatially uncertain quantities are often described in a probabilistic framework as random fields. A spatial dependent random field can be seen as a collection of random variables that are indexed by spatial coordinates. In this section, notions of random fields will be introduced following [Pavliotis \(2014\)](#) and [Ghanem and Spanos \(2003\)](#). Afterwards, this framework will be applied to the friction field and the advantages of this description will be highlighted.

3.1 Notions of random fields

Let

$$\{\Xi(x) : x \in \Omega\}, \quad \Omega \subset \mathbb{R} \quad (3.1)$$

be a one-dimensional space dependent second order stochastic process. Ξ is therefore a collection of random variables indexed by the position $x \in \Omega$. Moreover, if Ξ is a Gaussian random field, it is fully described by its mean $\mathbb{E}(\Xi(x)) = \bar{\xi}(x)$ and its covariance function $C_{\Xi}(x, x')$.

3.1.1 Gaussian random fields

A second-order random field $\Xi : \Omega \rightarrow \mathbb{R}$ is Gaussian if for any $x \in \Omega$, the corresponding random variable $\Xi(x)$ is Gaussian. The covariance matrix of $\Xi(x)$ and $\Xi(x')$ should be symmetric and positive semi definite. There exists several models for the such covariance function. Most of them can be summarized by the Matérn covariance functions expressed as

$$C_{\nu}(x, x') = \sigma^2 \frac{2^{1-\nu}}{\Gamma(\nu)} \left(\frac{\sqrt{2\nu}|x - x'|}{l_c} \right)^{\nu} K_{\nu} \left(\frac{\sqrt{2\nu}|x - x'|}{l_c} \right), \quad (3.2)$$

where $\Gamma(\cdot)$ is the gamma function, $K_{\nu}(\cdot)$ is the modified Bessel function of the second kind, l_c is the correlation length and σ^2 is the variance of the random field. The parameter ν controls the smoothness of the resulting random field, as C_{ν} is $\lfloor \nu \rfloor$ times differentiable.

According to Mercer's theorem, such symmetric positive definite covariance kernels can be expressed using their spectral representation

$$C_{\Xi}(x, x') = \sum_{k=1}^{\infty} \lambda_k \phi_k(x) \phi_k(x'), \quad (3.3)$$

where $\{(\lambda_k, \phi_k)\}_{k=1, \dots, \infty}$ is the collection of eigen-values and eigen-functions of the covariance kernel, that are solutions of the Fredholm equation

$$\int_{\Omega} C_{\Xi}(x, x') \phi_k(x') \, dx' = \lambda_k \phi_k(x). \quad (3.4)$$

Particular values of ν for the Matérn covariance functions lead to well-known covariance kernels. For $\nu = 1/2$, the exponential covariance kernel is retrieved:

$$C_{1/2}(x, x') = \sigma^2 \exp \left(-\frac{|x - x'|}{l_c} \right). \quad (3.5)$$

The limit case $\nu \rightarrow \infty$ results in the infinitely differentiable Gaussian covariance function

$$C_\infty(x, x') = \sigma^2 \exp\left(\frac{-(x - x')^2}{2l_c^2}\right). \quad (3.6)$$

In Fig. 9 are shown both $C_{1/2}$ and C_∞ in $\Omega \times \Omega$ for $x_c = 1800$ km and $l_c = 200$ km.




Figure 9: Particular cases of Matérn covariance functions for $l_c = 200$ km.

Mercer's theorem allows to approximate these covariance kernels by truncation of Eq. (3.3). Let us define \hat{C}_ν^K as the truncated spectral expansion of C_ν , that is

$$\hat{C}_\nu^K(x, x') = \sum_{k=1}^K \lambda_k \phi_k(x) \phi_k(x'). \quad (3.7)$$

One could now represent $\hat{C}_{1/2}^K$ and \hat{C}_∞^K graphically, as it is done in Fig. 10 for $K = 5$.




Figure 10: Particular cases of truncated Matérn covariance functions for $K = 5$ and $l_c = 200$ km.

It could be observed that the amount of information of $C_{1/2}$ that is covered by $\hat{C}_{1/2}^5$ is less important than the one of C_∞ that is covered by \hat{C}_∞^5 . This is because the eigen-value spectrum of C_∞ decays faster than the one of $C_{1/2}$, as it is shown in Fig. 11.




Figure 11: Eigen-value collapse of the Matérn covariance kernels $C_{1/2}$ and C_∞ with $l_c = 200$ km. For the case of $C_{1/2}$, the eigen-values decay asymptotically as n^{-1} , whereas for the case of C_∞ , the eigen-values decay exponentially.

The more rapidly the eigen-value spectrum decays, the less terms are needed in \hat{C}_ν to carry the principal informations of C_ν . As $C_{1/2}$ is not smooth at $x = x'$, the eigen-value spectrum cannot collapse that rapidly since the non-smoothness at $x = x'$ cannot be represented with a sum of smooth eigen-functions unless their respective frequencies tend to the infinity which is only the case of $\phi_\infty(x)$. Therefore, random fields represented with $C_{1/2}$ as the covariance kernel are non smooth and are difficult to generate. For this reason, $C_\Xi = C_\infty$ is often chosen for the representation of Gaussian random fields. The eigen-value collapse of Matérn covariance kernels depend not only on ν , but also on the correlation length l_c . In Fig. 12 the eigen-value collapse of C_∞ is represented for $l_c \in \{50, 100, 200\}$ km. It can be observed that the higher is the correlation length, the faster is the eigen-value decay. This result is intuitive since a small correlation length induces a higher fluctuations frequencies in the random field, such that higher frequency eigen-functions are needed to describe the field. First four eigen-functions of C_∞ with $l_c = 50$ km are shown in Fig. 13 in order to illustrate the frequency increase in the successive eigen-functions.




Figure 12: eigen-value spectrum of C_∞ for different correlation lengths and unit variance.




Figure 13: First four eigen-functions of C_∞ with $l_c = 50$ km.

3.1.2 Karhunen-Loève expansion

As a consequence to Eq. (3.3), the random field Ξ of mean $\bar{\xi}$ and covariance C_{Ξ} can be represented using a spectral expansion, known as the Karhunen-Loève (KL) expansion (Lévy (1965)):

$$\Xi(x) = \bar{\xi}(x) + \sum_{k=1}^{\infty} \theta_k \sqrt{\lambda_k} \phi_k(x), \quad (3.8)$$

where $\{\theta_k\}_{k=1, \dots, \infty}$ is a collection of uncorrelated random variables with zero mean and unit variance, and that are Gaussian in the case of Gaussian random fields. Therefore, given the mean field and the set of eigen-values and eigen-functions of the covariance function, realisations of such random fields can be generated by sampling from the joint probability density of $\{\theta_k\}_{k=1, \dots, \infty}$.

What is really advantageous in such representations is that the random field Ξ can be approximated by its K -th term truncation $\hat{\Xi}^K$

$$\hat{\Xi}^K(x) = \bar{\xi}(x) + \sum_{k=1}^K \theta_k \sqrt{\lambda_k} \phi_k(x). \quad (3.9)$$

Formally, $\hat{\Xi}^K$ is a second order random field of mean $\bar{\xi}$ and covariance function \hat{C}_{Ξ}^K . The truncation error between Ξ and $\hat{\Xi}^K$ is therefore related to the truncation error between C_{Ξ} and \hat{C}_{Ξ}^K . The rapid collapse of the eigen-values of C_{Ξ} ensures that this truncation is optimal, as it minimizes the mean square error term between $\hat{\Xi}^K$ and Ξ for K fixed (Ghanem and Spanos (2003)). This property allows to carry most of the information of Ξ in $\hat{\Xi}^K$ for small to moderate K .

3.2 Friction field as a random field

In the previous subsection, Gaussian random fields have been introduced. Here it will be explained why such representations are not suited for the friction field by application of the maximum entropy principle. This principle will then be used to modify the previous representation of such random fields.

3.2.1 Maximum entropy principle

Introduced by Jaynes (1957), the maximum entropy principle is a method that finds the most appropriate probability density function for a given random variable, given its structural informations. The given structural informations of a random variable constrain the set of possible probability density functions, and the maximum entropy principles states that among all the remaining densities, the one that is appropriate is the one that maximizes the entropy.

Formally, the entropy associated with a random variable $\mathbf{X} \in \mathcal{X}$ with unknown probability density function $\pi_{\mathbf{X}}(\mathbf{x})$ is defined as

$$S(\pi_{\mathbf{X}}) = - \int_{\mathcal{X}} \pi_{\mathbf{X}}(\mathbf{x}) \log(\pi_{\mathbf{X}}(\mathbf{x})) d\mathbf{x}. \quad (3.10)$$

The structural informations of the random variable are typically informations on its support, and on its moments. The maximum entropy principle can therefore be seen as an optimization problem with associated constraints:

$$\begin{aligned} \max \quad & S(\pi_{\mathbf{X}}) \\ \text{s.t.} \quad & \text{structural constraints on } \mathbf{X}, \end{aligned}$$

which is commonly solved with a Lagrange multiplier method.

3.2.2 Application to the friction field

Let

$$\{\nu_b(x) : x \in \Omega\}, \quad \Omega \subset \mathbb{R} \quad (3.11)$$

be a one-dimensional space dependent random field second-order an strictly positive stochastic process. That is, ν_b is a collection of strictly positive random variable of unknown probability density function π_{ν_b} , indexed by the position $x \in \Omega$ and provided with a mean field $\bar{\nu}(x)$ and a covariance function $C_\nu(x, x')$. The underlying structural informations of this definitions allow to write the maximum entropy principle for the friction field as

$$\max \quad S(\pi_{\nu_b}) \quad (3.12)$$

$$\text{s.t.} \quad \pi_{\nu_b}(\zeta) = 0 \quad \text{for} \quad \zeta \leq 0, \quad (3.13)$$

$$\int_{\mathbb{R}^+} \pi_{\nu_b}(\zeta) d\zeta = 1 \quad (3.14)$$

$$\int_{\mathbb{R}^+} \zeta \pi_{\nu_b}(\zeta) d\zeta = \bar{\nu} \quad (3.15)$$

$$\int_{\mathbb{R}^+} \log(\zeta) \pi_{\nu_b}(\zeta) d\zeta = c. \quad (3.16)$$

which finally leads to

$$\pi_{\nu_b}(\zeta) = \text{Gamma}(\zeta; \bar{\nu}, \sigma^2(\bar{\nu}, c)), \quad (3.17)$$

where the variance σ^2 can be deduced from the mean and the log-mean. The friction parameter ν_b should therefore not be described by a Gaussian random field, but by a gamma one.

3.2.3 Transformation from a Gaussian random field

A non Gaussian random field such as ν_b can be expressed as a non-linear transformation f of a Gaussian random field Ξ , such that

$$\nu_b(x) = f(\Xi(x)). \quad (3.18)$$

The transformation f can be seen as the composition of two transformations: $f = f_2 \circ f_1$. The first transformation f_1 can be taken as the cumulative distribution function of the normal distribution, that is,

$$f_1(x) = \frac{1}{2} \left(1 + \text{erf} \left(\frac{x - \bar{\xi}}{\sqrt{2}\sigma_\Xi} \right) \right), \quad (3.19)$$

where $\bar{\xi}$ and σ_Ξ are respectively the mean and the standard deviation of the Gaussian random field Ξ . Therefore, the random field resulting of the transformation $f_1(\Xi(x))$ is uniformly distributed

with support $[0, 1]$. This uniform random field can be now transformed by the inverse cumulative distribution function of the non Gaussian random field ν_b , which is a gamma random field here. The function f_2 can thus be expressed as

$$f_2(x) = \text{Inv-Gamma}(x; \bar{\nu}, \sigma_\nu^2). \quad (3.20)$$

One can finally express the gamma friction field as

$$\nu_b(x) = f_2 \circ f_1 \circ \Xi(x). \quad (3.21)$$

This sequence of transformations allows to generate a gamma random field by transformation of a Gaussian random field.

3.2.4 Realization of the friction field

It has been seen that realizations of the friction field could be generated based on a Gaussian random field of arbitrary mean and variance. However, the covariance function C_Ξ and its correlation length l_c have to be specified to generate realizations of the Gaussian random field, and the mean field $\bar{\nu}$ and the field variance σ_ν^2 have to be specified to apply the transformation f_2 .

The Gaussian covariance function is a common choice, and allows to represent smooth fields. Moreover, it has been shown that its eigen-value spectrum decays rapidly which allows to parametrize the corresponding random field with a small to moderate number of parameters. For these reasons, $C_\Xi(x, x') = C_\infty(x, x')$ will be kept. The corresponding correlation length will be chosen as $l_c = 100$ km, as it is not expected to be arbitrarily small, but smaller than the characteristic length of the ice-sheet which is here $x_c = 1000$ km. The mean field will be set to $\bar{\nu}_b = 7.624 \times 10^6$ [Pa m^{1/3} s^{-1/3}] by following the work of [Pattyn et al. \(2012\)](#). The standard deviation of the field will be set to 70 % of the mean field to allow large fluctuations. In Fig. 14 is shown some realization of the corresponding friction field.




Figure 14: Realizations of $\{\nu_b(x) : x \in \Omega\}$ for $C_\Xi(x, x') = C_\infty(x, x')$, $l_c = 100$ km, $\bar{\nu}_b = 7.624 \times 10^6$ [Pa m^{1/3} s^{-1/3}] and $\sigma_\nu = 0.7\bar{\nu}_b$.

4 Surrogate model

The finite-element model of the ice flow (under the Shallow Shelf Approximation) is time consuming. Typically, a simulation of the forward problem requires $O(10^3)$ [s] to converge which makes MCMC methods such as M-H intractable as they typically require $O(10^5)$ iterations to converge. An inexpensive surrogate model will be constructed based on polynomial chaos (PC) expansion.

4.1 Polynomial chaos surrogate model

A PC surrogate model is an approximation $g^p(\mathbf{m})$ of a function $g(\mathbf{m})$ that takes into account the prior uncertainty on its set of parameters \mathbf{m} . Formally, it can be expressed as

$$g^p(\mathbf{m}) = \sum_{i=0}^p c_i P_i(\mathbf{m}), \quad (4.1)$$

where p is the polynomial order of the surrogate, $\{c_i\}_{i=0}^p$ is a set of unknown coefficients, and $\{P_i(\mathbf{m})\}_{i=0}^p$ is a basis of polynomials that are orthogonal with respect to the prior probability density of \mathbf{m} over its support \mathcal{M} , that is,

$$\int_{\mathcal{M}} P_i(\mathbf{m}) P_j(\mathbf{m}) \pi_{\text{prior}}(\mathbf{m}) \, d\mathbf{m} = \delta_{ij}. \quad (4.2)$$

By multiplying Eq. (4.1) by $P_j(\mathbf{m})\pi(\mathbf{m})$ and by integrating over \mathcal{M} , we obtain

$$\int_{\mathcal{M}} g^p(\mathbf{m}) P_j(\mathbf{m}) \pi(\mathbf{m}) \, d\mathbf{m} = \int_{\mathcal{M}} \sum_{i=0}^p c_i P_i(\mathbf{m}) P_j(\mathbf{m}) \pi_{\text{prior}}(\mathbf{m}) \, d\mathbf{m} \quad (4.3)$$

$$= \sum_{i=0}^p c_i \int_{\mathcal{M}} P_i(\mathbf{m}) P_j(\mathbf{m}) \pi_{\text{prior}}(\mathbf{m}) \, d\mathbf{m} \quad (4.4)$$

$$= c_j \quad (4.5)$$

where the orthogonality property of the polynomial basis (Eq. (4.2)) has been used. Therefore, the set of unknown coefficients can be computed by approximating the remaining integral with an appropriate quadrature rule

$$c_j \approx \sum_{\alpha} g(\mathbf{m}_{\alpha}) P_j(\mathbf{m}_{\alpha}) w_{\alpha}, \quad (4.6)$$

where \mathbf{m}_{α} is the α^{th} node of integration, and w_{α} is its associated weight. The choice of polynomial basis depends of the nature of the uncertain parameter \mathbf{m} and on its probability density function.

4.2 Application to the SSA ice-flow model

In the framework of this thesis, an input-output model (where the inputs characterize the friction parameters, and the outputs characterize the observable velocity) of the ice-flow is needed. The friction field is parametrized by the unknown coefficients $\boldsymbol{\theta} = \{\theta_1, \dots, \theta_N\}$ of the KL expansion. Therefore, the surrogate model will be a function $g^p : \mathbb{R}^N \rightarrow (\mathbb{R}^+)^m$ that inputs a set of KL parameters $\boldsymbol{\theta} \in \mathbb{R}^N$ and outputs a set of velocities $\mathbf{u} \in (\mathbb{R}^+)^m$ that correspond to a set of observation points $\mathbf{x} \in \Omega^m$. According to Eq. 3.9, the prior probability density function over the KL parameter

space $\pi_{\text{prior}}(\boldsymbol{\theta})$ is a multivariate Gaussian with zero-mean and identity covariance matrix. Therefore, the appropriate polynomial basis to describe the model is the Hermite polynomial basis (Xiu and Karniadakis (2002)). The four first Hermite polynomials are shown in Fig. 15.




Figure 15: Four first Hermite polynomials over $x \in [-2, 2]$.

The surrogate model can therefore be expressed as

$$\mathbf{g}^p(\boldsymbol{\theta}) = \sum_{i=0}^p c_i H_i(\boldsymbol{\theta}), \quad (4.7)$$

where the set of coefficient is found by approximating integral of Eq. (4.5) with the Gauss-Hermite quadrature rule

$$c_j \approx \sum_{\alpha} \mathbf{g}(\boldsymbol{\theta}_{\alpha}) H_j(\boldsymbol{\theta}_{\alpha}) w_{\alpha}. \quad (4.8)$$

The integration grid is the N -dimensional extension of the roots of the r^{th} order Hermite polynomial H_r (where r is referred to as the level). That is, for a level $r > p$, the integration grid will be composed of r^N points with associated grid weights that are obtained by multiplying the corresponding one-dimensional roots. A schematic of a bi-dimensional integration grid chosen with respect to the 5th order surrogate model is shown in Fig.16.

As it can be observed, the number of integration nodes scales exponentially with the dimensionality of the problem which makes the construction of a surrogate model challenging for large dimensional input space. To tackle this problem, sparse grid quadrature methods have been developed which makes the integration cheaper (Le Maître and Knio (2010)) but will not be described in detail here.

Surrogate models of order p with level r can therefore be constructed. For the purpose of illustration, let us show the construction of a surrogate model with level $r = 8$ with varying the order p . The surrogate will output a single velocity observation ($m = 1$) u^{obs} at location $x^{\text{obs}} = 500\text{km}$, and input the two first KL coefficients θ_1 and θ_2 . This surrogate model is shown in Fig. 17 for $p \in \{1, 3, 5, 7\}$.

It could be observed that the surrogate model accuracy increases with p at high density regions of $\pi_{\text{prior}}(\boldsymbol{\theta})$, but simple observations cannot conclude about the global accuracy of the surrogate




Figure 16: Integration grid based on $H_5(x)$ shown with the contour plot of the bivariate normal distribution of zero mean and identity covariance. The corresponding weights are the largest for nodes close to the origin.

model with respect to p . Let us define the weighted error R_w between the surrogate and the true model with respect to $\pi_{\text{prior}}(\boldsymbol{\theta})$ as

$$R_w = \int_{\mathbb{R}^N} |g^p(\boldsymbol{\theta}) - g(\boldsymbol{\theta})|^2 \pi_{\text{prior}}(\boldsymbol{\theta}) \, d\boldsymbol{\theta} \quad (4.9)$$

$$\approx \sum_{\alpha} w_{\alpha} (g^p(\boldsymbol{\theta}_{\alpha}) - g(\boldsymbol{\theta}_{\alpha}))^2, \quad (4.10)$$

which can be evaluated as a function of the polynomial order p to check that it indeed decreases. As shown in Fig. 18, R_w does indeed decrease with p . In order to avoid overfitting, the order is chosen smaller than the level ($p < r$), so that $p = r - 1$ will be kept.

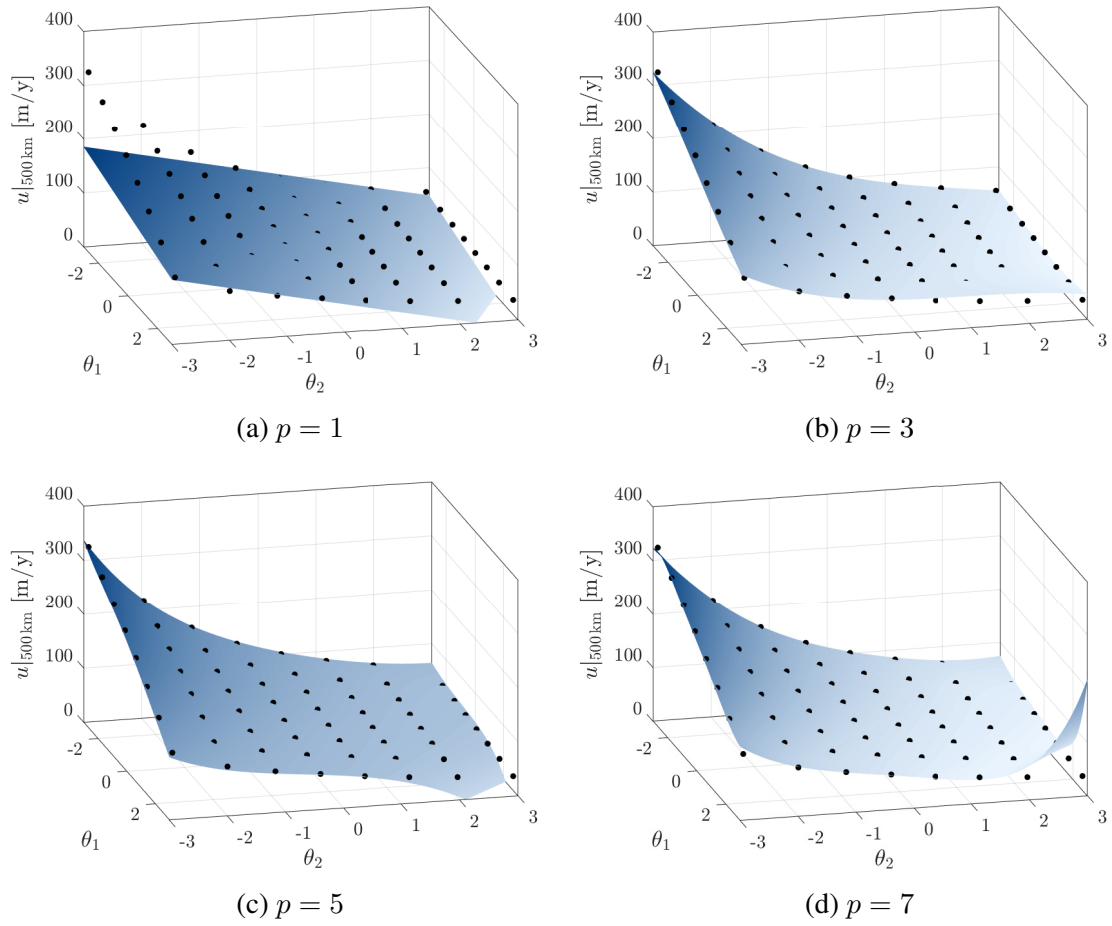


Figure 17: Surface plot of the surrogate model $g^p(\boldsymbol{\theta}) \approx g(\boldsymbol{\theta}) = u|_{500\text{km}}$ for $r = 8$ and varied order p . Discrete points denote the true model evaluations $g(\boldsymbol{\theta})$.




Figure 18: Weighted error R_w as a function of the surrogate model order p , with corresponding scaling. The input dimensionality and the level are respectively $N = 2$ and $r = 8$, as in Fig. 17.

Part II

Inverse problem

5 Generalities on inverse problem

In this section, the inverse problem will be formulated using different approaches. At first, the context and notations that are necessary to understand how the forward and inverse problems are linked are briefly described. Afterwards, the inverse problem will be formulated in a general way based on, among others, [Tarantola \(2005\)](#) and [Stuart \(2010\)](#). The deterministic approach will be first introduced. The inverse problem will be subsequently formulated in the Bayesian framework, and the similarities between these two approaches will be highlighted. Finally, the general formulations will be particularized to the glaciology problem of the inversion of the friction field given observations of the surface velocity.

5.1 Context and notations

Let us consider a model $f : \mathcal{X} \times \mathcal{M} \rightarrow \mathbb{R}^p$ which, given a set of parameters $m \in \mathcal{M}$, has as a solution a state $x \in \mathcal{X}$ that verifies

$$f(x; m) = 0. \quad (5.1)$$

The model f is commonly a system of partial differential equations with associated boundary conditions. The solution state x gives rise to an observable $d \in \mathcal{D}$ defined generally as

$$d = h(x), \quad (5.2)$$

where $h : \mathcal{X} \rightarrow \mathcal{D}$ is a state to observable mapping. In the framework of inverse problems, one is interested in the parameter to observable mapping $g : \mathcal{M} \rightarrow \mathcal{D}$ that given any set of parameters $m \in \mathcal{M}$ finds the corresponding observable $d \in \mathcal{D}$ such that

$$d = g(m). \quad (5.3)$$

The evaluation of the parameter-to-observable mapping therefore requires the evaluation of the model and of the state-to-observable mapping, as it is illustrated in Fig. 19.




Figure 19: Input output schematic of the parameter-to-observable map g , and its equivalence in terms state-to-observable map h and of model f .

In general, an inverse problem consists in, from a set of (possibly noisy) observations $d^{\text{obs}} \in \mathcal{D}$, retrieving the informations about the set of parameter m that led to these observations. In practice,

such problems are ill-posed (*i.e.* their solution is not guaranteed to exist nor to be unique, and is not stable with respect to perturbations in the given observations). These problems are therefore solved in frameworks that are able to ensure the well-posedness of the problems.

5.2 Deterministic framework

The first framework in which an inverse problem can be solved is the deterministic one. In the deterministic framework, the inverse problem is formulated as an optimization problem which has as a solution the set a set of parameter $\mathbf{m}^* \in \mathcal{M}$.

In general, the observations \mathbf{d}^{obs} are noisy (*i.e.* contain error). Thus, they can be expressed as

$$\begin{aligned}\mathbf{d}^{\text{obs}} &= \mathbf{d} + \boldsymbol{\epsilon} \\ &= \mathbf{g}(\mathbf{m}) + \boldsymbol{\epsilon},\end{aligned}\tag{5.4}$$

where $\boldsymbol{\epsilon}$ includes (i) model errors, which are due to the incapacity of the model to represent the true process that links the observable quantity to the set of parameter; (ii) measurement errors which are due to the imperfection of the measurement instrumentation; (iii) numerical errors which are introduced by the numerical methods used for the model to be solved. These errors imply that it is likely that no set of parameter $\mathbf{m} \in \mathcal{M}$ allow the observable \mathbf{d} to fit exactly the observation \mathbf{d}^{obs} . For this reason, the problem has to be formulated as an optimization problem, that is

$$\mathbf{m}^* = \arg \min_{\mathbf{m} \in \mathcal{M}} \mathcal{J}(\mathbf{m}; \mathbf{d}^{\text{obs}}),\tag{5.5}$$

where J is a cost function. In practice, the cost function is often chosen as a least-squares error between observable and observation

$$\mathcal{J}(\mathbf{m}; \mathbf{d}^{\text{obs}}) = \frac{1}{2} \|\mathbf{g}(\mathbf{m}) - \mathbf{d}^{\text{obs}}\|_{\mathcal{D}}^2,\tag{5.6}$$

where $\|\cdot\|_{\mathcal{D}}$ denotes the l^2 -norm defined in the \mathcal{D} space and the factor of one half is added for convenience. Such least-squares formulations are often ill-posed, as their solution is not guaranteed to be unique nor stable. These formulations have therefore to be regularized. The regularization of such ill-posed problems can be done by adding a penalization term to Eq. (5.6) (Tikhonov (1963)). A common choice for the regularization term is proportional to the norm of the differences between the parameters to be inferred and a reference value for these parameters, which leads to

$$\mathcal{J}_R(\mathbf{m}; \mathbf{m}_0, \mathbf{d}^{\text{obs}}) = \frac{1}{2} \|\mathbf{g}(\mathbf{m}) - \mathbf{d}^{\text{obs}}\|_{\mathcal{D}}^2 + \frac{\alpha_R}{2} \|\mathbf{m} - \mathbf{m}_0\|_{\mathcal{M}}^2,\tag{5.7}$$

where α_R controls the magnitude of the regularization, prior knowledge can be incorporated through $\mathbf{m}_0 \in \mathcal{M}$ such that value of the set of parameter that are far from \mathbf{m}_0 are penalized, and $\|\cdot\|_{\mathcal{M}}$ denotes the l^2 -norm defined in the \mathcal{M} space. The solution of Eq. (5.5) with the cost function defined in Eq. (5.7) therefore exists for sufficiently large α_R , and is a point estimate $\mathbf{m}^* \in \mathcal{M}$, which achieves a balance between being sufficiently close to \mathbf{m}_0 and minimizing the error between the observations and the model predictions.

5.3 Bayesian frame work

In the Bayesian framework, the solution to the inverse problem is not a set of parameter \mathbf{m}^* but a probability density function (pdf) for the parameters \mathbf{m} given the observations \mathbf{d}^{obs} , that is the conditional pdf $\pi_{\text{post}}(\mathbf{m}|\mathbf{d}^{\text{obs}})$. The Bayes' rule reads as

$$\pi_{\text{post}}(\mathbf{m}|\mathbf{d}^{\text{obs}}) = \frac{\pi_{\text{like}}(\mathbf{d}^{\text{obs}}|\mathbf{m}) \times \pi_{\text{prior}}(\mathbf{m})}{\int_{\mathcal{M}} \pi_{\text{like}}(\mathbf{d}^{\text{obs}}|\mathbf{m})\pi_{\text{prior}}(\mathbf{m}) d\mathbf{m}}, \quad (5.8)$$

where $\pi_{\text{like}}(\mathbf{d}^{\text{obs}}|\mathbf{m})$ is the likelihood, $\pi_{\text{prior}}(\mathbf{m})$ is the prior pdf, and the denominator is the normalization constant. Equation (5.8) should be interpreted as an update of the prior knowledge on \mathbf{m} given the observations \mathbf{d}^{obs} , that leads to a posterior pdf $\rho(\mathbf{m}|\mathbf{d}^{\text{obs}})$, as it is represented in Fig. 20.




Figure 20: Illustration of the Bayes rule. In (c), the posterior pdf (that is proportional to the product of the prior and the likelihood) can be seen as an update of the prior, given the observations taken into account by the likelihood.

The inverse problem in the Bayesian framework can be solved with sampling methods which, given the likelihood and the prior, will explore the posterior. However, exploring the posterior pdf is challenging for large scale inverse problems, as the space to be explored increases exponentially with the number of variables. Alternatively, from this problem can be defined an optimization problem that is to maximize the posterior (or equivalently, to maximize the product between the prior and the likelihood without having to compute the normalization constant in Eq. (5.8)). The set of parameter that solves this optimization problem is referred to as the maximum a posteriori (MAP) $\mathbf{m}^{\text{MAP}} \in \mathcal{M}$ and follows

$$\begin{aligned} \mathbf{m}^{\text{MAP}} &= \arg \max_{\mathbf{m} \in \mathcal{M}} \pi_{\text{post}}(\mathbf{m}|\mathbf{d}^{\text{obs}}) \\ &= \arg \max_{\mathbf{m} \in \mathcal{M}} \{ \pi_{\text{like}}(\mathbf{d}^{\text{obs}}|\mathbf{m}) \times \pi_{\text{prior}}(\mathbf{m}) \}. \end{aligned} \quad (5.9)$$

A common (but questionable) assumption is that the noise in Eq. (5.4) follows a centered normal distribution, $\epsilon \sim \mathcal{N}(\mathbf{0}, \Sigma_{\text{noise}})$, where Σ_{noise} is the noise covariance matrix that is diagonal if the noises are independent, and directly proportional to the identity matrix if the noises are iid

($\Sigma_{\text{noise}}|_{\text{iid}} = \sigma^2 \mathbf{I}$). Under the assumption of Gaussian noise, one has

$$\begin{aligned}\pi_{\text{like}}(\mathbf{d}^{\text{obs}}|\mathbf{m}) &\propto \exp\left(-\frac{1}{2}(\mathbf{g}(\mathbf{m}) - \mathbf{d}^{\text{obs}})^T \Sigma_{\text{noise}}^{-1}(\mathbf{g}(\mathbf{m}) - \mathbf{d}^{\text{obs}})\right) \\ &= \exp\left(-\frac{1}{2}\|\mathbf{g}(\mathbf{m}) - \mathbf{d}^{\text{obs}}\|_{\Sigma_{\text{noise}}^{-1}}^2\right).\end{aligned}\quad (5.10)$$

Then, if the prior is uniform with infinite support (*i.e.* non-informative), the optimization problem of Eq. (5.9) is reduced to

$$\begin{aligned}\mathbf{m}^{\text{MAP}} &= \arg \max_{\mathbf{m} \in \mathcal{M}} \left\{ \exp\left(-\frac{1}{2}\|\mathbf{g}(\mathbf{m}) - \mathbf{d}^{\text{obs}}\|_{\Sigma_{\text{noise}}^{-1}}^2\right) \right\} \\ &= \arg \max_{\mathbf{m} \in \mathcal{M}} \left\{ -\frac{1}{2}\|\mathbf{g}(\mathbf{m}) - \mathbf{d}^{\text{obs}}\|_{\Sigma_{\text{noise}}^{-1}}^2 \right\} \\ &= \arg \min_{\mathbf{m} \in \mathcal{M}} \left\{ \frac{1}{2}\|\mathbf{g}(\mathbf{m}) - \mathbf{d}^{\text{obs}}\|_{\Sigma_{\text{noise}}^{-1}}^2 \right\},\end{aligned}\quad (5.11)$$

which shows how maximizing the likelihood in the Bayesian framework is equivalent to the inverse problem of Eq. (5.5) from the deterministic framework, with the cost function defined in Eq. (5.6). If the prior is chosen as Gaussian centered at \mathbf{m}_0 with a given covariance matrix Σ_{prior} such that

$$\begin{aligned}\pi_{\text{prior}}(\mathbf{m}) &\propto \exp\left(-\frac{1}{2}(\mathbf{m} - \mathbf{m}_0)^T \Sigma_{\text{prior}}^{-1}(\mathbf{m} - \mathbf{m}_0)\right) \\ &= \exp\left(-\frac{1}{2}\|\mathbf{m} - \mathbf{m}_0\|_{\Sigma_{\text{prior}}^{-1}}^2\right),\end{aligned}\quad (5.12)$$

so that Eq. (5.11) becomes

$$\mathbf{m}^{\text{MAP}} = \arg \min_{\mathbf{m} \in \mathcal{M}} \left\{ \frac{1}{2}\|\mathbf{g}(\mathbf{m}) - \mathbf{d}^{\text{obs}}\|_{\Sigma_{\text{noise}}^{-1}}^2 + \frac{1}{2}\|\mathbf{m} - \mathbf{m}_0\|_{\Sigma_{\text{prior}}^{-1}}^2 \right\}, \quad (5.13)$$

which shows how maximizing the posterior with informative prior is linked with the regularized deterministic inverse problem.

5.4 Litterature review of inverse problem in glaciology

The governing equations of ice-sheet flows contain both phenomenologically uncertain parameters (*e.g.* ice rheology, friction law) and parameters that are non-available due to the impossibility of measuring them directly, such as the friction parameter field. If physical reasonings allow to constrain the phenomenologically uncertain parameters, the non measurable quantities have to be inferred from observation data through an inversion process. For this reason, the inverse problem which consists in inferring the friction field from observations of the surface velocity has garnered considerable interest among numerous scientists. In this subsection, the work of some of these scientists will be reviewed. The deterministic approach will be addressed first, and will be followed by the Bayesian approach.

Deterministic approach: As the parameter to be inferred is a bounded function from Ω to \mathbb{R}^+ , the underlying inverse problem is an infinite-dimensional inverse problem where to each local point

$x \in \Omega$, a value $\nu_b(x) \in \mathbb{R}^+$ has to be inferred. In a deterministic framework, this leads to the following least-squares minimization problem

$$\nu_b^* = \arg \min_{\nu_b(\mathbf{m})} \left\{ \|\mathbf{u}^{\text{obs}} - \mathbf{g}(\mathbf{m})\|_{\Sigma_{\text{noise}}^{-1}}^2 + R(\mathbf{m}) \right\}, \quad (5.14)$$

where R is a regularization term, often taken to penalize large gradients in the friction field which smoothens it (Petra et al. (2012); Morlighem et al. (2010)), or to penalize friction fields that are far to an expected prior field. This inverse problem is formulated as an optimization problem that can be solved with gradient descent methods. Such methods require the computation of the gradient of a cost function with respect to the set of parameter at each iteration. The most basic method to compute these gradients relies on forward finite differentiation (Larour et al. (2012)) which is computationally expensive as it requires $N + 1$ forward model evaluation for N -dimensional problems. There exist methods allowing to compute these cost function gradients in a more efficient way. A common approach is to analytically derive an adjoint problem where only one single additional forward model evaluation is needed for the computation of the gradient (MacAyeal (1993); Vieli and Payne (2003); Morlighem et al. (2013)). Automatic differentiation (AD) can be used as an alternative to compute gradients (Larour et al. (2014)).

All of these methods require the set of parameters (here the friction field) to be finite dimensional. The friction field will therefore be discretized, and expressed in a general form:

$$\nu_b^N(x) = \sum_{k=1}^N m_k \phi_k(x) = \mathbf{m}^T \boldsymbol{\phi}(x), \quad (5.15)$$

where $\{\phi_k(x)\}_{k=1}^N$ is a finite-dimensional set of pre-determined shape functions, and $\{m_k\}_{k=1}^N$ is a finite-dimensional set of parameter that can be inferred from observations. Finite-element discretizations of ν_b are common (e.g. MacAyeal et al. (1995); Goldberg and Sergienko (2011)), but lead to large scale inverse problems as the number of elements is usually large in ice-sheet numerical models. It has been shown that Newton's methods should be used instead of basic gradient descent methods for large scale inverse problems, as its convergence rate is insensitive to the problem dimensions (Petra et al. (2012)). Such methods use the information of the hessian to accelerate the convergence. As the computation of the Hessian is intractable for computationally expensive forward model evaluations, a low rank approximation of the Hessian can be used instead as its eigen values often collapse rapidly to zero (Villa et al. (2021)).

Other types of discretizations can be used in order to decrease the problem dimensionality. For example, a finite element discretization with a reduced number of degrees of freedom is often used. Fourier modes have been used in MacAyeal (1993).

Bayesian approach: This inverse problem can be solved in the Bayesian framework in order to quantify the uncertainty associated with the inferred parameters. Mathematically, this leads to the proportionality relationship

$$\pi_{\text{post}}(\mathbf{m} | \mathbf{u}^{\text{obs}}) \propto \pi_{\text{prior}}(\mathbf{m}) \pi_{\text{like}}(\mathbf{u}^{\text{obs}} | \mathbf{m}), \quad (5.16)$$

where the finite set of parameter to be inferred \mathbf{m} is derived from the general discretization of the friction field Eq. (5.15). Under Gaussian observation noise and Gaussian prior assumptions, one has

$$\pi_{\text{post}}(\mathbf{c} | \mathbf{u}^{\text{obs}}) \propto \exp \left(-\frac{1}{2} \|\mathbf{g}(\mathbf{m}) - \mathbf{u}^{\text{obs}}\|_{\Sigma_{\text{noise}}^{-1}}^2 - \frac{1}{2} \|\mathbf{m} - \mathbf{m}_{\text{prior}}\|_{\Sigma_{\text{prior}}^{-1}}^2 \right), \quad (5.17)$$

with $\mathbf{m}_{\text{prior}}$ and Σ_{prior} are the mean and the covariance matrix of the Gaussian prior respectively, and Σ_{noise} is the noise covariance matrix. Note that, as the mapping $\mathbf{g}(\mathbf{m})$ is non linear, the Gaussianity of π_{prior} does not imply π_{post} to be Gaussian. The posterior is therefore likely to be a complex distribution that cannot be fully described by its two first moments (that are its mean and its covariance matrix). Sampling methods such as Monte Carlo Markov Chain methods (MCMC) are typically used to explore the posterior pdf in the parameter space, as they only require the prior information and the likelihood to be pointwise evaluable. These methods generate a Markov chain of states in the parameter space by, at each step, sampling a proposal around the current state and accepting it with a probability that requires the evaluation of a likelihood to be computed. These methods are popular but are facing two major issues. The first is that the more parameters there are in the space, the harder it is to explore it, which makes these methods challenging for large-scale inverse problems. The second is that, as these methods require the evaluation of the likelihood (and therefore, of the forward model) at each step and as the number of steps is typically of the order of the million for these methods to converge, such methods are intractable for computationally expensive forward models such as ice sheet numerical models. Such methods can however be accelerated. [Isaac et al. \(2015\)](#) have proposed a Gaussian approximation of the posterior, centered at the maximum a posteriori with a covariance matrix derived from a low rank approximation of the hessian of the misfit expressed in Eq. (5.6) at the MAP point. The creation of a surrogate model, constructed through a fixed number of forward model evaluations, can render the likelihood evaluation tractable which strongly accelerates the MCMC-based methods ([Marzouk et al. \(2007\)](#)), which have been applied in the case of computationally expensive ice-sheet models in [Bulthuis et al. \(2019\)](#).

6 Sampling methods

As it has been mentioned in Sec. 5.3, the computation of the posterior requires, according to Eq. (5.8), the computation of a non tractable integral. Therefore, the analytical direct computation of the posterior is generally not an option. Alternatively, Markov chain Monte Carlo (MCMC) sampling methods can be used to estimate the posterior. The idea of MCMC methods is to build a Markov chain which explores the parameters space in accordance to a target posterior. In this section, the Metropolis-Hastings (Metropolis et al. (1953); Hastings (1970)) algorithm will be introduced and a procedure to enhance the convergence of this method will be described.

6.1 Metropolis-Hastings algorithm

The Metropolis-Hastings algorithm generates a Markov chain (*i.e.* a chain whose state depend exclusively on the previous state). At each step of this method, a candidate state \mathbf{m}^* is sampled from a pre-defined proposal distribution q . Since the target posterior π_{post} can be pointwise evaluated, one can evaluate how likely is the candidate state compared to the previous one. The new candidate has a probability α to be accepted. This procedure is defined properly in Algorithm 1.

Algorithm 1 Metropolis-Hastings Algorithm

- 1: Define proposal density $q(\mathbf{m}_1|\mathbf{m}_2)$
 - 2: Initialize state $\mathbf{m}^{(0)}$
 - 3: **for** $k = 1$ to N_{max} **do**
 - 4: Sample $\mathbf{m}^* \sim q(\mathbf{m}^*|\mathbf{m}^{(k-1)})$
 - 5: Compute $\alpha = \pi_{\text{post}}(\mathbf{m}^*)q(\mathbf{m}^*|\mathbf{m}^{(k-1)})/\pi_{\text{post}}(\mathbf{m}^{(k-1)})q(\mathbf{m}^{(k-1)}|\mathbf{m}^*)$
 - 6: Sample $u \sim U_{[0,1]}$
 - 7: **if** $u < \alpha$ **then**
 - 8: Set $\mathbf{m}^{(k)} = \mathbf{m}^*$
 - 9: **else**
 - 10: Set $\mathbf{m}^{(k)} = \mathbf{m}^{(k-1)}$
 - 11: **end if**
 - 12: **end for**=0
-

A key point in the performance of M-H MCMC is the choice of proposal density q . A common choice of proposal is the Gaussian pdf centered in the current state \mathbf{m} , that is

$$q(\mathbf{m}^*, \mathbf{m}) = \frac{1}{\sqrt{(2\pi)^d \det[\Sigma]}} \exp\left(-\frac{1}{2}(\mathbf{m} - \mathbf{m}^*)^T [\Sigma]^{-1} (\mathbf{m} - \mathbf{m}^*)\right), \quad (6.1)$$

where d is the dimensionality of \mathbf{m} . Such symmetric proposals (*i.e.* proposals for which $q(\mathbf{m}^*, \mathbf{m}) = q(\mathbf{m}, \mathbf{m}^*)$) lead to the particular case of the random walk Metropolis-Hastings (RWMH) algorithm. The choice of covariance matrix $[\Sigma]$ will strongly affect the exploration of the target posterior as it will dictate how the proposal is confident (locally concentrated) or conservative (diffused). If the proposal density is too confident, the posterior might be badly explored and the Markov chain might be stuck in local modes. If the proposal density is too conservative, the acceptance rate might be too low which leads to bad performance of the algorithm. Several methods can be proposed to choose appropriate proposals. The first is to rely on prior knowledge, which can

come from a previous burn in phase. For instance, the RWMH algorithm can be performed with an arbitrarily tuned covariance matrix, and statistics can be assessed from the generated Markov chain such as its covariance matrix and its MAP point. RWMH can therefore be performed starting from the previously found MAP point, with the so found covariance matrix. This method has been shown to work, and does indeed lead to a Markovian process as the burn in phase is separated from the final phase. The finally generated Markov chain is shown to verify some ergodicity properties, which guarantee that it will converge towards a stationary distribution that is the posterior.

6.2 Adaptive methods

More systematic methods have been proposed for tuning the proposal properties. Adaptive proposal (AP) have been proposed in [Haario et al. \(1999\)](#). They proposed Gaussian proposals with a covariance matrix that is computed from the U last states where U is a memory length. This method therefore produces a non-Markovian chain as the proposal depends on the history and not only on the current state. However, this method has been shown to not verify the ergodicity property that is required for the chain to converge towards a stationary distribution. It is two years later that the authors have proposed the Adaptive Metropolis (AM) algorithm which verifies the ergodicity property for bounded posterior with a bounded support ([Haario et al. \(2001\)](#)). Its idea is to update the proposal by taking into account the whole history. Formally, the proposal covariance matrix at the k -th point of the chain is given by

$$[\Sigma]^{(k)} = \begin{cases} [\Sigma]^{(0)}, & k \leq k_0 \\ s_d \text{cov}(\mathbf{m}^{(0)}, \dots, \mathbf{m}^{(k-1)}) + s_d \epsilon I_d, & k > k_0, \end{cases} \quad (6.2)$$

where s_d depends exclusively on the dimensionality, I_d is the d -dimensional identity matrix, k_0 is the length of the initial period, and $\epsilon > 0$ is a parameter that ensures the positive definiteness of the covariance matrix. It is typically chosen very small compared to the size of the support. A common choice for s_d is

$$s_d = \frac{2.4^2}{d}, \quad (6.3)$$

from [Gelman et al. \(1996\)](#). Here again, one can rely on prior informations or on a previous burn-in phase for the choice of $[\Sigma]^{(0)}$.

7 Neural posterior estimation

Instead of using MCMC methods to sample from the pointwise evaluable posterior $p(\boldsymbol{\theta}|\mathbf{x})$, an estimator $q \in \mathcal{Q}$ of the posterior from which it is possible to draw samples can be learned with neural posterior estimation (NPE). Inference process with such posterior estimators is referred to as *likelihood-free inference*, as once the estimator is constructed, there are no need for the simulator to be evaluated. Moreover, NPEs are powerful as they are amortized estimators, *i.e.* estimators that approximate a general posterior that is not pre-conditioned on a given \mathbf{x} . In this section, the construction of such posterior estimator will be turned into an optimization problem through the concept of variational inference, and normalizing flows will be introduced to parametrize the family \mathcal{Q} according to [Papamakarios et al. \(2021\)](#); [Kobyzev et al. \(2020\)](#).

7.1 Variational inference

Variational inference methods turn the search for an estimator $q \in \mathcal{Q}$ of a posterior p into an optimization problem. This optimization problem will be to minimize the discrepancy between $q(\boldsymbol{\theta}|\mathbf{x})$ and $p(\boldsymbol{\theta}|\mathbf{x})$, with respect to the parameters that parametrize q . Therefore, this discrepancy has to be quantified. The most usual way of quantifying the discrepancy between two distributions is to compute their Kullback-Leibler divergence, written D_{KL} , that is

$$D_{KL}(p(\boldsymbol{\theta}|\mathbf{x})||q(\boldsymbol{\theta}|\mathbf{x})) = \mathbb{E}_{\boldsymbol{\theta} \sim p(\boldsymbol{\theta}|\mathbf{x})} \left\{ \log \left(\frac{p(\boldsymbol{\theta}|\mathbf{x})}{q(\boldsymbol{\theta}|\mathbf{x})} \right) \right\}. \quad (7.1)$$

However, in the present case, samples from the posterior are not available. Thus, it is the expected Kullback-Leibler divergence \tilde{D}_{KL} that is used instead, that is

$$\tilde{D}_{KL}(p(\boldsymbol{\theta}|\mathbf{x})||q(\boldsymbol{\theta}|\mathbf{x})) = \mathbb{E}_{\mathbf{x} \sim p(\mathbf{x})} \mathbb{E}_{\boldsymbol{\theta} \sim p(\boldsymbol{\theta}|\mathbf{x})} \left\{ \log \left(\frac{p(\boldsymbol{\theta}|\mathbf{x})}{q(\boldsymbol{\theta}|\mathbf{x})} \right) \right\}. \quad (7.2)$$

The chain rule of probability writes

$$p(\boldsymbol{\theta}, \mathbf{x}) = p(\mathbf{x})p(\boldsymbol{\theta}|\mathbf{x}) = p(\boldsymbol{\theta})p(\mathbf{x}|\boldsymbol{\theta}), \quad (7.3)$$

such that Eq. (7.2) becomes

$$\tilde{D}_{KL}(p(\boldsymbol{\theta}|\mathbf{x})||q(\boldsymbol{\theta}|\mathbf{x})) = \mathbb{E}_{(\mathbf{x}, \boldsymbol{\theta}) \sim p(\boldsymbol{\theta}, \mathbf{x})} \left\{ \log \left(\frac{p(\boldsymbol{\theta}|\mathbf{x})}{q(\boldsymbol{\theta}|\mathbf{x})} \right) \right\}. \quad (7.4)$$

The expected Kullback-Leibler divergence \tilde{D}_{KL} can therefore be evaluated if samples of the joint distribution $p(\boldsymbol{\theta}, \mathbf{x})$ are available. According to Eq. (7.3), samples $(\boldsymbol{\theta}^*, \mathbf{x}^*) \sim p(\boldsymbol{\theta}, \mathbf{x})$ can be sequentially drawn by first sampling $\boldsymbol{\theta}^* \sim p(\boldsymbol{\theta})$, and then sampling $\mathbf{x}^* \sim p(\mathbf{x}|\boldsymbol{\theta}^*)$. In the framework of variational inference, optimal estimator q^* can be found by minimizing \tilde{D}_{KL} which leads to the following optimization problem

$$q^*(\boldsymbol{\theta}|\mathbf{x}) = \arg \min_{q \in \mathcal{Q}} \tilde{D}_{KL}(p(\boldsymbol{\theta}|\mathbf{x})||q(\boldsymbol{\theta}|\mathbf{x})) \quad (7.5)$$

$$= \arg \min_{q \in \mathcal{Q}} \mathbb{E}_{(\mathbf{x}, \boldsymbol{\theta}) \sim p(\boldsymbol{\theta}, \mathbf{x})} \left[\log \left(\frac{p(\boldsymbol{\theta}|\mathbf{x})}{q(\boldsymbol{\theta}|\mathbf{x})} \right) \right] \quad (7.6)$$

$$= \arg \min_{q \in \mathcal{Q}} \mathbb{E}_{(\mathbf{x}, \boldsymbol{\theta}) \sim p(\boldsymbol{\theta}, \mathbf{x})} (-\log q(\boldsymbol{\theta}|\mathbf{x})), \quad (7.7)$$

which shows how minimizing the Kullback-Leibler divergence is equivalent to maximizing the data likelihood in the context of variational inference. In the following, the negative data likelihood will be referred to as the loss.

7.2 Normalizing flows

The principle of variational inference allows to turn the search for an estimator $q(\boldsymbol{\theta}|\boldsymbol{x})$ of a posterior $p(\boldsymbol{\theta}|\boldsymbol{x})$ into an optimization problem. However, the choice of estimator family \mathcal{Q} still needs to be specified. The choice of such families will be seen as parametrizations of q , such that the optimization problem can be solved with gradient descent methods in the space of these parameters. Normalizing flows (Tabak and Vanden-Eijnden (2010); Tabak and Turner (2013)) can be used to constrain the family \mathcal{Q} in the context of variational inference (Rezende and Mohamed (2015)). In this subsection, the general principle of normalizing flows will be described, and a particular architecture will be reviewed.

7.2.1 General principle

Let \boldsymbol{z} be a d -dimensional vector in \mathbb{R}^d from which we would like to know the joint probability density function $p_Z(\boldsymbol{z})$. The vector \boldsymbol{z} can be expressed as a function of another vector $\boldsymbol{u} \in \mathbb{R}^d$ sampled from a known distribution $p_U(\boldsymbol{u})$ by using an invertible transformation T such that

$$\boldsymbol{z} = T(\boldsymbol{u}) \Leftrightarrow \boldsymbol{u} = T^{-1}(\boldsymbol{z}). \quad (7.8)$$

The joint probability density of \boldsymbol{z} , noted $p_Z(\boldsymbol{z})$ can be expressed with respect to $p_U(\boldsymbol{u})$ since $\boldsymbol{z} = T(\boldsymbol{u})$, which leads to

$$p_Z(\boldsymbol{z}) = p_U(\boldsymbol{u}) |\det J_T(\boldsymbol{u})|^{-1} \quad (7.9)$$

$$= p_U(T^{-1}(\boldsymbol{z})) |\det J_{T^{-1}}(T^{-1}(\boldsymbol{z}))| \quad (7.10)$$

where J_T is the jacobian of the forward transformation ($J_T = \partial T(\boldsymbol{u})/\partial \boldsymbol{u}$), and $J_{T^{-1}}$ is the jacobian of the inverse transformation ($J_{T^{-1}} = \partial T^{-1}(\boldsymbol{z})/\partial \boldsymbol{z}$). The jacobian determinant is here a normalization constant, which will account for the change of volume due to the transformation. The transformation of the probability densities therefore require T and its inverse T^{-1} to be differentiable. Such differentiable invertible transforms are known as diffeomorphism. An important property of diffeomorphism is that they can be defined as a composition such that

$$T = T_N \circ \dots \circ T_1, \quad (7.11)$$

where T_1, \dots, T_N are diffeomorphisms as well. The determinant of J_T can therefore be expressed as

$$\det J_T = \det J_{T_1}(\boldsymbol{u}) \cdot \det J_{T_2}(T_1(\boldsymbol{u})) \cdot \dots \cdot J_{T_N}(T_{N-1}(\dots(T_1(\boldsymbol{u}))). \quad (7.12)$$

Formally, a normalizing flow is the transformation of a simple base distribution (such as $p_U(\boldsymbol{u})$) into an arbitrarily complex distribution $p_Z(\boldsymbol{z})$ through a sequence of invertible and differentiable transformations $\{T_k\}_{k=1, \dots, N}$. This process is illustrated in Fig. 21.

The final density $q_N(\boldsymbol{z})$ can be sampled from by sampling from the base distribution and by applying the successive transformations to the samples, and can be evaluated through Eq. (7.10). To




Figure 21: Schematic of the normalizing flow mechanism. A simple base distribution $q_0(\mathbf{z}_0)$ (equivalent to $p_U(\mathbf{u})$) is successively transformed by a sequence of well defined transformations into a complex distribution $q_N(\mathbf{z}_N)$ (equivalent to $p_Z(\mathbf{z})$). Samples $\mathbf{z}_N \sim q_N(\mathbf{z})$ can be drawn by sampling $\mathbf{z}_0 \sim q_0(\mathbf{z}_0)$ and by applying successive transformations such that

$$\mathbf{z}_N \sim T_N \circ \dots \circ T_1(\mathbf{z}_0).$$

summarize, in practice, the transformations inside a normalizing flow should be (i) differentiable since the determinant of the jacobian has to be computed for sampling, and determinant of the inverse jacobian has to be computed for evaluating the density; (ii) invertible since T is needed for sampling and T^{-1} is required for evaluating the density; (iii) computationally efficient in terms of forward and reverse transformations, and in terms of jacobian determinant computation; (iv) sufficiently expressive for the final density $q_N(\mathbf{z})$ to be a good estimator of the target posterior. In what follows, autoregressive transformations will be described as transformations which satisfy the above conditions.

7.2.2 Autoregressive flows

Autoregressive flows are normalizing flows composed of autoregressive transformations. Autoregressive transformations are transformations of the form

$$z_i = T(u_i, \phi_i(\mathbf{u}_1, \dots, \mathbf{u}_{i-1})) = T(u_i, \phi_i(\mathbf{u}_{1:i-1})), \quad (7.13)$$

where T is called the transformer, and ϕ_i is the conditioner of the i -th transform. The corresponding inverse transformation is given by

$$u_i = T^{-1}(z_i, \phi_i(\mathbf{u}_{1:i-1})). \quad (7.14)$$

Conditioners expressed as $\phi_i(\mathbf{u}_{1:i-1})$ highlight the fact that z_i depends exclusively on components of \mathbf{u} that have lower indices, which is the idea of an autoregressive transformation. Such transformations have a triangular jacobian as

$$[J_T(\mathbf{u})]_{ij} = \frac{\partial z_i}{\partial u_j} = \frac{\partial T(u_i, \phi_i(\mathbf{u}_{1:i-1}))}{\partial u_j} = 0, \quad \forall i < j. \quad (7.15)$$

The determinant of such triangular matrices is the product of its diagonal elements

$$\det J_T(\mathbf{u}) = \prod_{k=1}^d \frac{\partial T(u_k, \phi_k)}{\partial u_k}. \quad (7.16)$$

The autoregressive property can be seen in the inverse direction to Eq. (7.13). Indeed, one could express

$$z_i = T(u_i, \phi_i(\mathbf{z}_{1:i-1})), \quad (7.17)$$

so that the i -th conditioner depends on $\mathbf{z}_{1:i-1}$ instead of $\mathbf{u}_{1:i-1}$. Normalizing flows with autoregressive transformations formulated as Eq. (7.13) and Eq. (7.17) have been respectively proposed in Papamakarios et al. (2017) where they introduce masked autoregressive flows (MAF), and Kingma et al. (2016) where inverse autoregressive flows (IAF) are introduced. Both MAF and IAF use affine autoregressive transforms

$$\mathbf{z} = \boldsymbol{\alpha} \odot \mathbf{u} + \boldsymbol{\beta}, \quad (7.18)$$

where $\alpha_i = \alpha_i(u_i, \phi_i(\mathbf{u}_{1:i-1}))$ and $\beta_i = \beta_i(u_i, \phi_i(\mathbf{u}_{1:i-1}))$ in the case of MAF, and $\alpha_i = \alpha_i(u_i, \phi_i(\mathbf{z}_{1:i-1}))$ and $\beta_i = \beta_i(u_i, \phi_i(\mathbf{z}_{1:i-1}))$ in the case of IAF. The conditioners ϕ_i are typically multilayer perceptrons (MLPs). Such MLPs can be made autoregressive by applying masks (Germain et al. (2015)). The idea of masked MLPs is simple: a degree between 1 and d is assigned to each unit of each layer of the MLP, and a mask is applied such that for a given layer, no unit feeds into units of lower or equal degree (Papamakarios et al. (2021)). Such autoregressive flows are differentiable, invertible if $\boldsymbol{\alpha} \neq \mathbf{0}$, and their jacobian determinant computation is simply

$$\det J_T = \prod_{i=1}^d \alpha_i. \quad (7.19)$$

A great expressivity can be achieved with MAF by stacking several transformations of type Eq. (7.18). However, it is an open question if MAF are universal approximators Papamakarios et al. (2017). In the present work, the architecture of normalizing flows that will be used is the MAF, which are implemented using Zuko library (Rozet (2022)).

8 Formulation of the inverse problem

In this section, all the material that has been reviewed so far will be used to formulate properly the inverse problem which will be tackled in this thesis, as well as the underlying methodology. Here, the inverse problem consists in inferring a friction field from a collection of m noisy velocity observations $\mathbf{u}^{\text{obs}} \in \mathbb{R}^m$, taken at predefined observation points $\mathbf{x}^{\text{obs}} \in \Omega^m$. In this work, the observations will be synthetic. A realization of $\{\nu_b : x \in \Omega\}$ will be drawn (see Sec. 3), and the corresponding observed velocity \mathbf{u}^{obs} will be synthesized using the numerical implementation of the SSA equations. A Gaussian noise $\boldsymbol{\eta} \in \mathbb{R}^m$ will be added, where $\boldsymbol{\eta} \sim \mathcal{N}(\mathbf{0}, \sigma_\eta \mathbf{I})$, $\sigma_\eta = 10^{-7}$ m/s.

8.1 Friction inversion

As the complete inversion of ν_b leads to an infinite dimensional problem, a truncated field $\{\nu_b : x \in \Omega\}$ will be inferred instead by truncation of Eq. (3.9), so that $\boldsymbol{\theta} = \{\theta_k\}_{k=1, \dots, K}$ define the truncated gaussian field $\hat{\Xi}^K$, and $\hat{\nu}_b = f_2 \circ f_1(\hat{\Xi}^K)$. The Bayesian inference problem is then be formulated as follows: Given noisy observations \mathbf{u}^{obs} generated by the friction field ν , find the conditional distribution (*i.e.* the posterior) $\pi_{\text{post}}(\boldsymbol{\theta} | \mathbf{u}^{\text{obs}})$, where $\boldsymbol{\theta}$ represents the parameters defining $\hat{\nu}_b$.

8.1.1 Prior and likelihood

The inversion methods require the prior and likelihood to be specified.

1. The prior pdf $\pi_{\text{prior}}(\boldsymbol{\theta})$ represents the a priori knowledge in the realization of $\boldsymbol{\theta}$. Friction field realizations are drawn by sampling an infinite sequence $\{\theta_k\}_{k=1, \dots, \infty}$ where $\theta_k \sim \mathcal{N}(0, 1)$. Thus, the prior for the finite set of parameters $\boldsymbol{\theta}$ is chosen as

$$\pi_{\text{prior}}(\boldsymbol{\theta}) = \mathcal{N}(\mathbf{0}, \mathbf{I}). \quad (8.1)$$

2. The likelihood $\pi_{\text{like}}(\mathbf{u} | \boldsymbol{\theta})$ is implicitly defined by the simulator. However, the bayesian inversion process requires a large number of forward model evaluation. Since the numerical implementation of the SSA equations is time consuming (this numerical implementation requires $O(100)$ [s] to be evaluated), it is intractable for bayesian inference. A computationally-inexpensive surrogate model is built instead, as presented in Sec. 4. Such surrogate models \mathbf{g} take finite set of parameters $\boldsymbol{\theta}$ and return the velocity \mathbf{u} at the observation points \mathbf{x}^{obs}

$$\mathbf{g}(\boldsymbol{\theta}) = \mathbf{u}. \quad (8.2)$$

Given the observation velocity \mathbf{u}^{obs} , the likelihood is defined with respect to the Gaussian noise model

$$\pi_{\text{like}}(\mathbf{u} | \boldsymbol{\theta}) = \exp\left(-\frac{1}{2\sigma_\eta^2} \|\mathbf{u} - \mathbf{u}^{\text{obs}}\|^2\right). \quad (8.3)$$

No normalization constant is needed in π_{like} as it is not constrained to be a probability density function in itself.

8.1.2 Methods

The Bayesian inference problem will be solved with both sampling methods and variational methods involving deep learning.

Adaptive Metropolis: For the AM method to be used, the initial point $\boldsymbol{\theta}^{(0)}$ has to be defined, as well as the initial covariance matrix $[\Sigma]^{(0)}$ and the initial period t_0 before the initial covariance matrix is updated. Here,

$$\begin{cases} \boldsymbol{\theta}^{(0)} &= \mathbf{0}, \\ [\Sigma]^{(0)} &= [I], \\ t_0 &= 20. \end{cases} \quad (8.4)$$

The expressiveness of this method will be assessed on a toy problem where the posterior is known and complex, and the choice of proposal covariance matrix $[\Sigma]^{(0)}$ will be discussed. Afterwards, the methods will be applied to a multidimensional problem for the inversion of a friction field.

Neural posterior estimation: The neural posterior estimation method that will be used estimates the posterior by a masked autoregressive flow (MAF), whose transformation parameters are tuned by maximizing the loss which is here the data negative log-likelihood. The optimization is performed with a stochastic gradient descent algorithm. Here, the Adam optimizer with weight decay (AdamW) will be used, with a learning rate $\gamma = 0.001$. The parameters of the dataset are sampled with respect to the prior: $\boldsymbol{\theta}^* \sim \pi_{\text{prior}}(\boldsymbol{\theta})$. The corresponding velocities are generated with the surrogate model: $u^* = \mathbf{g}(\boldsymbol{\theta}^*)$. The dataset will be separated into three parts: (i) the training dataset is the dataset with respect to which the loss will be minimized through a stochastic gradient method; (ii) the validation dataset is the dataset on which the loss will be evaluated in order to obtain an unbiased estimate. The best estimator will be considered as the one that minimizes the loss over the validation dataset. This dataset allows to select an estimator during the stochastic gradient descent process; (iii) the testing set will be over which the quality of the chosen estimator will be assessed. To quantify the quality of the posterior estimator, the expected coverage probability (ECP) will be computed (see [Hermans et al. \(2021\)](#)). The parameters of the MAF such as the number of transformations, the number of layers in each masked MLP, and the number of units in each hidden layer have to be specified. The number of transformations in the MAF as well as the size of the dataset used to train the MAF parameters are expected to have an influence which will be studied, but the number of layers and units in the masked MLPs will be fixed during the whole study.

8.2 Methodology

The problem that has been defined will be solved by using both sampling methods and variational inference with normalizing flows. Two different inverse problems will be solved: (i) a two-dimensional *toy problem* in which the target posterior is known and rendered complex enough so that the expressivity of the methods can be assessed. This will consist in a validation phase; (ii) a five-dimensional problem which will involve an unknown posterior, but where the quality of the results will be assessed, among others, in terms of inverted friction field compared to the one that gave rise to the observed data. This will consist in a testing phase.

8.2.1 Validation phase

For the validation phase, a two dimensional surrogate model is generated to implicitly represent the likelihood. For this, a 8×8 integration grid suited for the Gauss-Hermite quadrature is generated (see Sec. 4). The construction of this surrogate model requires 64 evaluations of the numerical implementation of the SSA equations. The order of the surrogate model is $p = 7$. In this phase, a single point observation at $x^{\text{obs}} = 500$ km is considered. The surrogate model takes as input the bidimensional set of parameter θ and outputs the corresponding velocity u at $x = 500$ km, that is $g(\theta) = u|_{x=500 \text{ km}}$. In fig. 22 is shown the corresponding surrogate model.




Figure 22: 2D Surrogate model for the validation phase, $r = 8$, $p = 7$. This surrogate model takes in input two parameters θ_1 and θ_2 , and outputs the velocity at $x^{\text{obs}} = 500$ km.

As it can be observed, the fit between the surrogate model and the discrete point (*i.e.* the evaluations of the numerical implementation of the SSA equations) is great close to the origin but gets worse far from the origin, as the Gauss-Hermite quadrature rule is built to be precise for $\theta \sim \mathcal{N}(0, \mathbf{I})$. However, one could observe that close to the origin, the surrogate model is linear. Since the noise and the prior are both Gaussian, this will lead to a Gaussian posterior, which is not a complex enough problem to conclude on the expressiveness of the methods that will be used. For this reason, and exclusively in the context of this validation phase, the prior will be taken as

$$\pi_{\text{prior}}(\theta) = U_{[-5, 5] \times [-5, 5]}, \quad \theta = \{\theta_1, \theta_2\}. \quad (8.5)$$

The conservativeness of this prior ensures that non-linear regions of the surrogate model are not penalized. From a physical reasoning, this choice of prior is bad as the surrogate model is built for normally distributed parameters, but this inverse problem will be used for evaluating the expressiveness of the methods exclusively. By evaluating the surrogate model in a grid in $\theta \in [-5, 5] \times [-5, 5]$, the posterior can be drawn. Fig. 23 shows the target posterior for an observation $u^{\text{obs}} = 96.7$ m/y.




Figure 23: Target bidimensional posterior for an observation $u^{\text{obs}} = 96.7 \text{ m/y}$.

As it can be observed, this distribution is complex and multimodal. This is due to the fact that the surrogate model is not injective.

8.2.2 Testing phase

For the testing phase, a five dimensional surrogate model is generated on a $4 \times 4 \times 4 \times 4 \times 4$ integration grid. The construction of this surrogate model therefore involves $4^5 = 1024$ evaluations of the numerical implementation of the SSA equations. The order of this surrogate model is $p = 3$. Ten equally spaced observation points are considered between 50 km and 600 km ($\mathbf{x}^{\text{obs}} = \{50, 100, \dots, 600\} \text{ km}$). The surrogate model \mathbf{g} is therefore a vectorial function that outputs the 12 corresponding velocities given a five dimensional set $\boldsymbol{\theta}$. For the observations to be synthesized with the SSA equations implementation, the target friction field is generated and shown in Fig. 24a. This friction field leads to the observation shown in Fig. 24b




Figure 24: (a) Target friction field ν_b^* sampled from Eq. (3.21); (b) Synthetized velocity \mathbf{u} at \mathbf{x}^{obs} , a noise $\boldsymbol{\eta}$ is applied so that $\mathbf{u}^{\text{obs}} = \mathbf{u}|_{\mathbf{x}^{\text{obs}}} + \boldsymbol{\eta}$.

Part III

Results

In this part, the Bayesian inverse problems that have been established in the previous part will be solved. At first, the two dimensional inverse problem described in Sec. 8.2.1 will be solved to validate the methods. Afterwards, the five-dimensional inverse problem described in Sec. 8.2.2 will be solved to infer the friction field. Both MCMC methods and normalizing flows will be used, and the advantages of both methods will be discussed.

9 Validation phase

9.1 Sampling methods

The two sampling methods that have been described in Sec. 6 are used here. The one that performs the best will be kept for solving the five-dimensional inverse problem.

9.1.1 Metropolis-Hastings algorithm

Let us begin with the MH algorithm with a gaussian proposal density. Let the first state of the Markov chain be at the origin ($\theta^{(0)} = \mathbf{0}$), and let us study the impact of the proposal covariance matrix $[\Sigma]$. For this purpose, two different proposal covariance matrices will be tested: (i) $[\Sigma] = 0.1 \times [I]$; (ii) $[\Sigma] = [I]$. The algorithm will be performed until the chain contains 200000 states, and its convergence will be assessed by showing the statistics of the chain. In Fig. 25 are shown the evolution of the states during of the MH algorithm, for the first and last ten thousand iterations (Complete evolution of the states are shown in Fig. 44, Appendix B). One observes that the number of iterations is higher when $[\Sigma] = 0.1 \times [I]$ to reach the number of 200000 states. In terms of acceptance rate, one observes 6% of accepted states for $[\Sigma] = 0.1 \times [I]$ versus 10% for $[\Sigma] = [I]$. Moreover, one could notice that regions close to $\theta_1 \approx 5$ are more frequently explored with $[\Sigma] = [I]$, which suggests that some regions are poorly explored when $[\Sigma] = 0.1 \times [I]$.

In order to evaluate if the chain has converged towards a stationary distribution, it is common to represent graphically some of the statistics of the chain such as the mean states of the chain with respect to the number of accepted samples. In Fig. 26 is shown the evolution of the mean states of the chain. As it can be observed, the convergence behavior of the mean states is better for $[\Sigma] = [I]$ than for $[\Sigma] = 0.1 \times [I]$. Moreover, one could notice that the mean states do not converge towards the same values for $[\Sigma] = 0.1 \times [I]$ that for $[\Sigma] = [I]$, which suggests that the generated chains do not converge towards the same distribution.

The part of the chain that corresponds to the converged states can be shown in the parameter space to conclude to which distribution the MH algorithm has converged. In Fig. 27 is shown the posterior estimated from converged MH samples, together with a target posterior level curve. As expected, the choice of proposal covariance matrix influences the stationary distribution to which the generated chain converges. For $[\Sigma] = 0.1 \times [I]$, MH does not explore the two modes as the steps generated by the proposal is small compared to the gap between the two modes. The approximate posterior is satisfying for $[\Sigma] = [I]$.

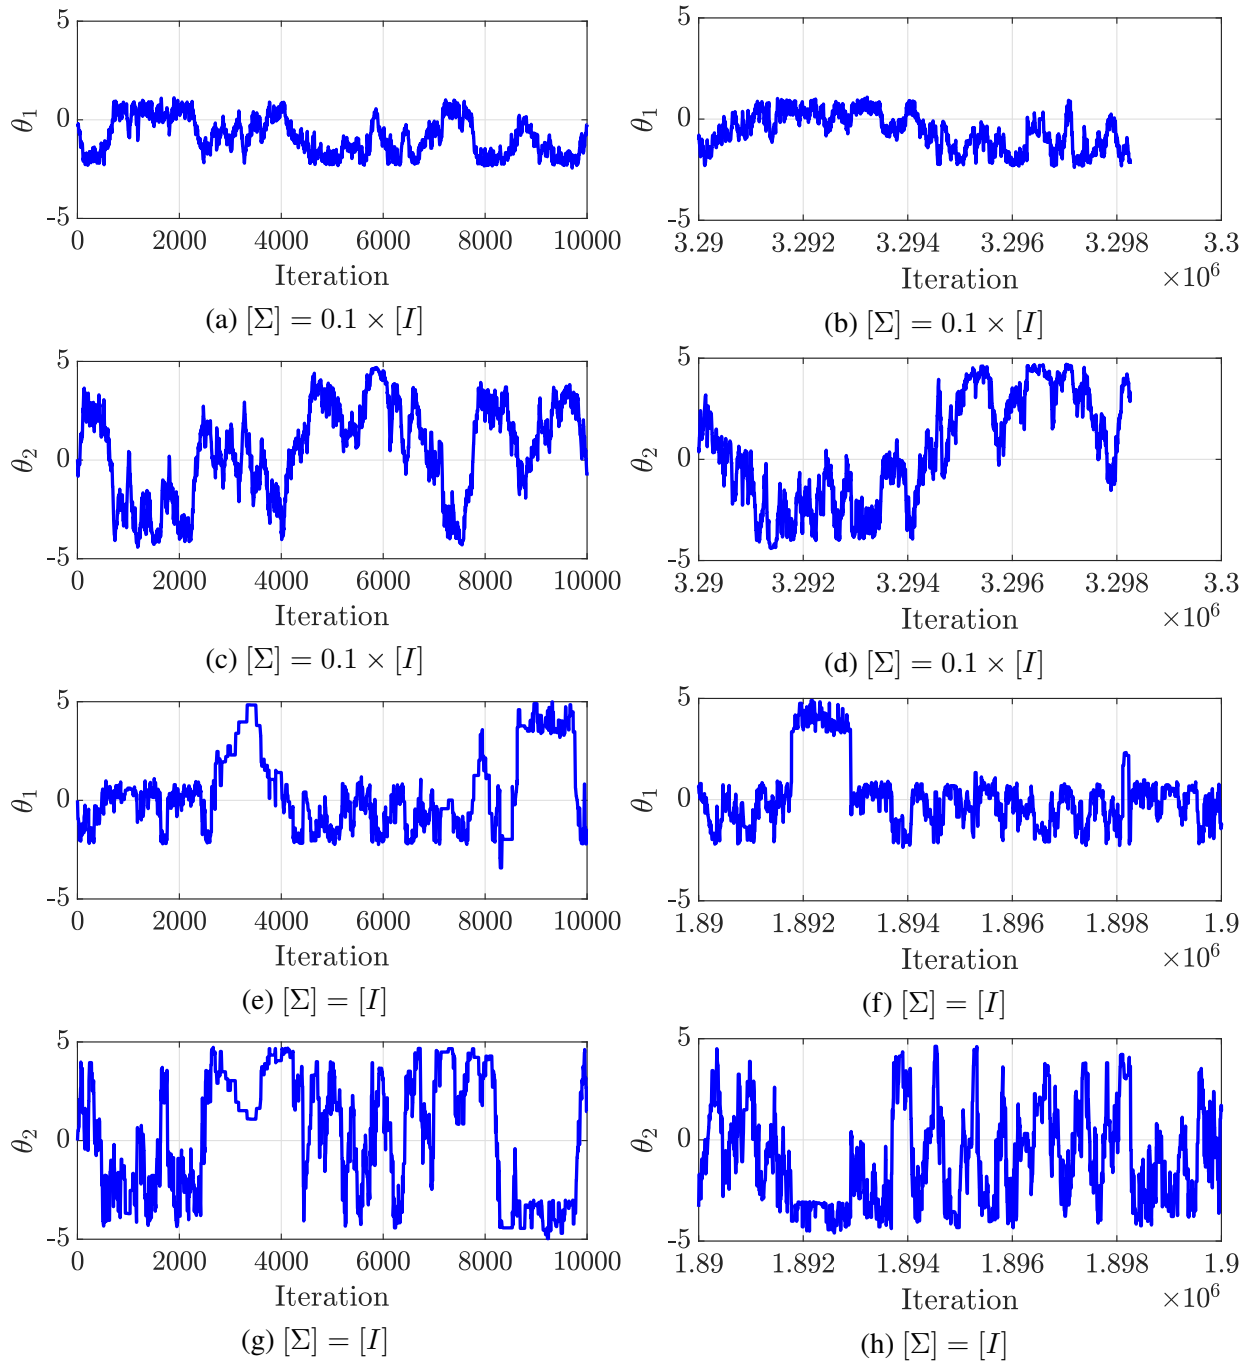


Figure 25: Chain of states for the Metropolis-Hastings (MH) algorithm. On the left column, the first ten thousand iterations of the MH algorithm are shown whereas on the right column it is the last ten thousand.




Figure 26: Mean statistics of the states in the markov chain generated by the MH algorithm with $[\Sigma] = 0.1 \times [I]$ (in blue) and $[\Sigma] = [I]$ (in red).




Figure 27: Posterior estimated from converged samples of the MH algorithm. Light blue correspond to density regions where $\pi_{\text{post}}^{\text{MH}}(\theta|u^{\text{obs}}) < 0.01$. Dark blue correspond to density regions where $\pi_{\text{post}}^{\text{MH}}(\theta|u^{\text{obs}}) \geq 0.01$. Dashed-dot black curve correspond to the target posterior level curve $\pi_{\text{post}}(\theta|u^{\text{obs}}) = 0.01$.

9.1.2 Adaptive-Metropolis algorithm

The same two-dimensional validation problem can be solved with the AM algorithm, starting the chain at the origin. Two starting proposal covariance matrices $[\Sigma]^{(0)}$ will be tested as for the case of the MH algorithm. Since this method is adaptive, the choice of starting covariance matrix is not expected to change the stationary distribution, except if the choice of starting covariance matrix influences the converged covariance matrix. In Fig. 28 is shown the evolution of the states over the first and last ten thousand iteration of the AM algorithm (Evolution over all the iterations is shown

in Fig. 45, Appendix B). One observes that once the covariance matrices have adapted (after a large number of iterations), regions where θ_1 is close to 5 are greatly explored. Moreover, the acceptance rate seems to decrease once the proposal covariance matrices are adapted.

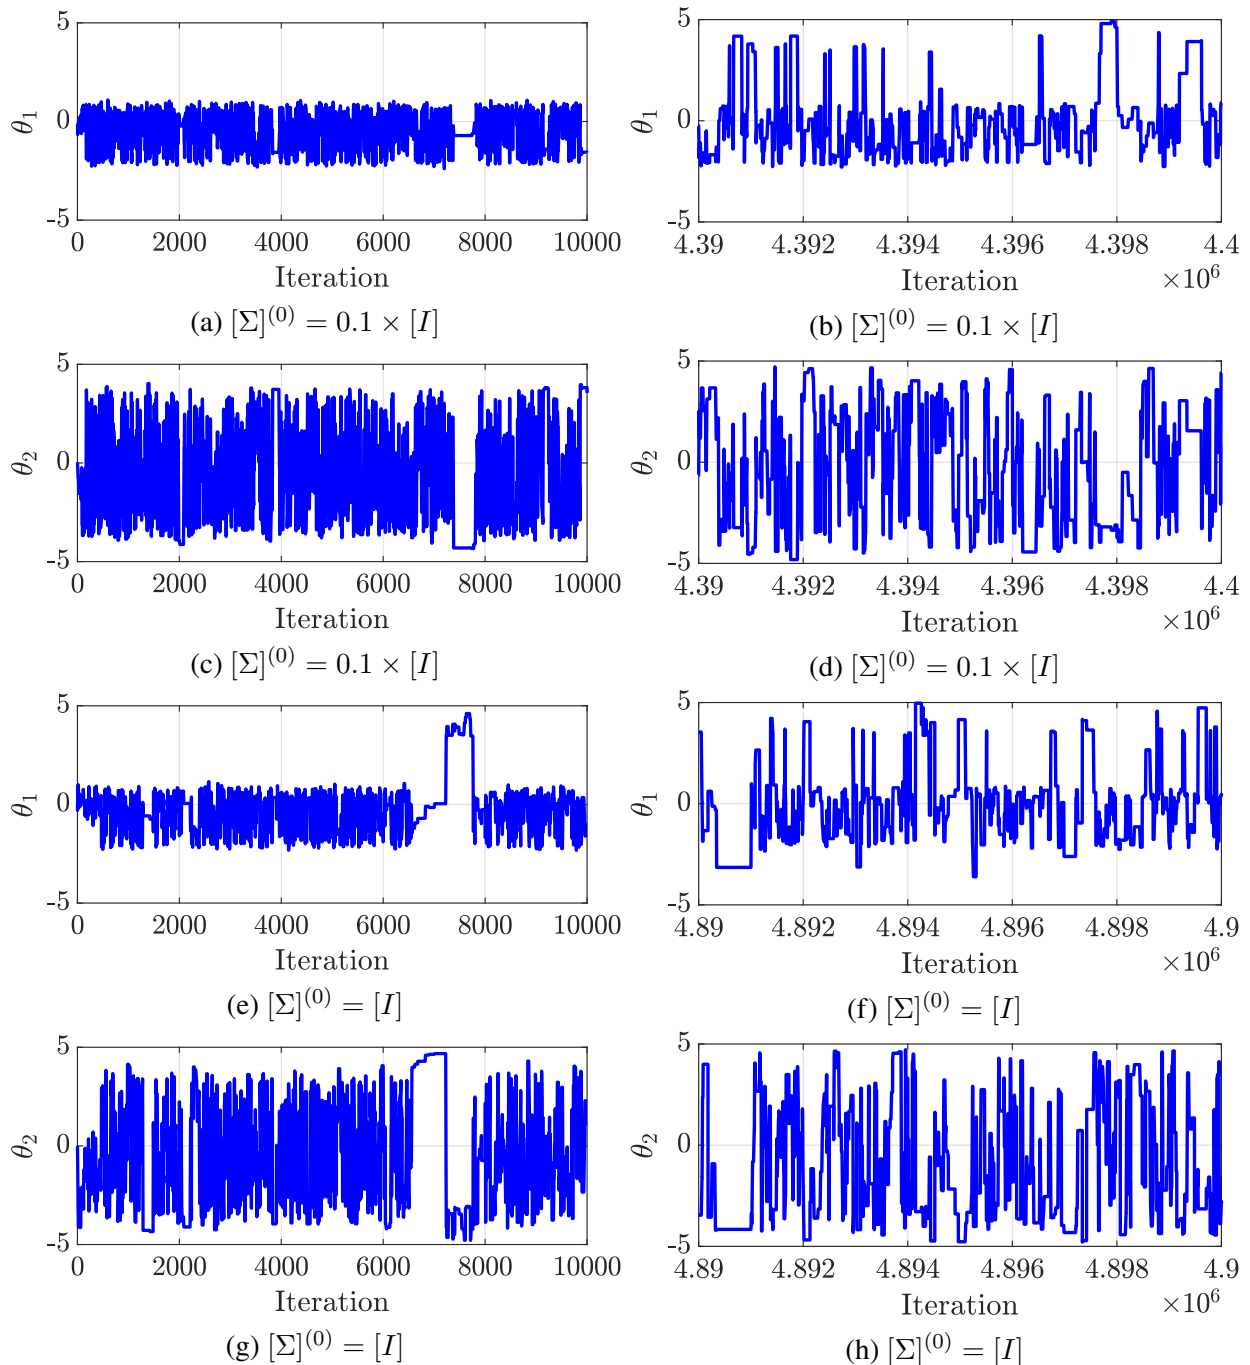


Figure 28: Chain of states for the adaptive Metropolis (AM) algorithm. On the left column, the first ten thousand iterations of the MH algorithm are shown whereas on the right column it is the last ten thousand.

As it can be observed in Fig. 29, the mean states converge towards the same value for both initial covariance matrices. This suggests that, in this case, the stationary distribution to which the

chain converges is not impacted by the initial covariance matrix. However, the convergence of the mean states is faster for $[\Sigma]^{(0)} = [I]$. The convergence property of the chain is thus impacted by the choice of initial covariance.




Figure 29: Mean state $\bar{\theta}$ as a function of the number of states in the chain, using AM algorithm. Grey areas approximate the 95% confidence interval, according to the central limit theorem.

The fact that both initial proposal covariance matrices led to the same stationary distribution is illustrated in Fig. 30. It can be observed that the target posterior is retrieved.




Figure 30: Posterior estimated from converged samples of the AM algorithm. Light blue correspond to density regions where $\pi_{\text{post}}^{\text{AM}}(\boldsymbol{\theta}|u^{\text{obs}}) < 0.01$, dark blue correspond to density regions where $\pi_{\text{post}}^{\text{AM}}(\boldsymbol{\theta}|u^{\text{obs}}) \geq 0.01$. Dashed-dot black curve correspond to the target posterior level curve $\pi_{\text{post}}(\boldsymbol{\theta}|u^{\text{obs}}) = 0.01$.

9.2 Neural posterior estimation

In this section, samples will not be drawn directly from the posterior, but an neural posterior estimator (NPE) from which we can sample will be created with masked autoregressive flows (MAF). The neural network parameters that led to the transformations are optimized with a stochastic gradient descent algorithm (AdamW) over a training dataset. After training, the NPE that is kept is the one that minimizes the negative log-likelihood over the validation dataset (which is referred to as the validation loss) to avoid overfitting. The final performance of a NPE is assessed by the expected coverage probability (ECP) which is computed over the testing dataset. The sizes of the validation dataset and the testing dataset are respectively $N_{\text{valid}} = 30000$ and $N_{\text{test}} = 3000$. The quality of the NPE will be evaluated as a function of the number of transformations and of the size of the dataset on which the NPE is trained. The architecture of each of the masked MLPs behind each transformation is fixed. Each masked MLP contains 4 hidden layers of 64 units.

9.2.1 Number of transformations

For NPE with MAF to be expressive, several transformations should be stacked. Therefore it is expected for the NPE to be of reduced quality if the number of transformations is too low. In this subsection, the size of the training dataset is fixed to $N_{\text{train}} = 80000$. Figure 31 shows the evolution of the training loss and the validation loss during the training for various number of transformations.

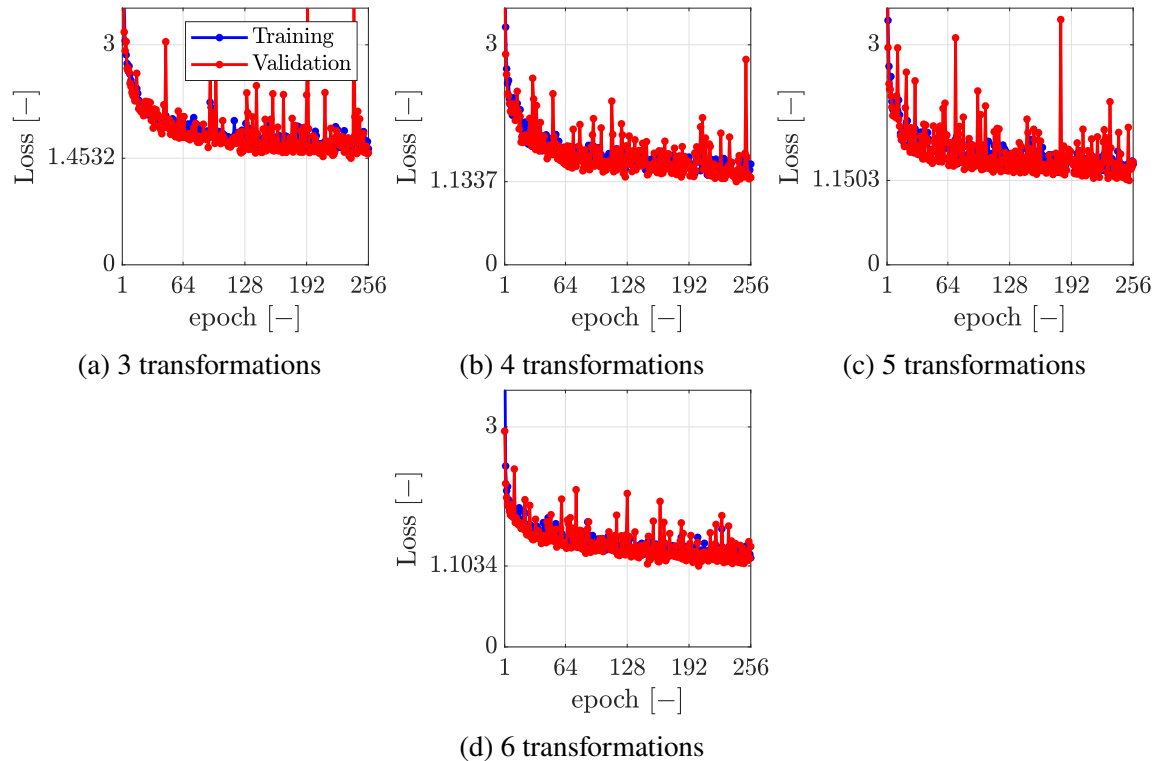


Figure 31: Evolution of the training loss (in blue) and of the validation loss (in red) with respect to the number of epochs, for various number of transformations. The minimum of the validation loss is indicated for each.

As it can be observed, the minimum of the validation loss is significantly higher for NPE with 3

transformations than for the others. Four, five and six transformations give similar results in terms of minimum loss, which suggests that 4 transformations are sufficient for MAF to be expressive enough for the current target posterior. In Fig. 32 are shown the NPEs for these various number of transformations.

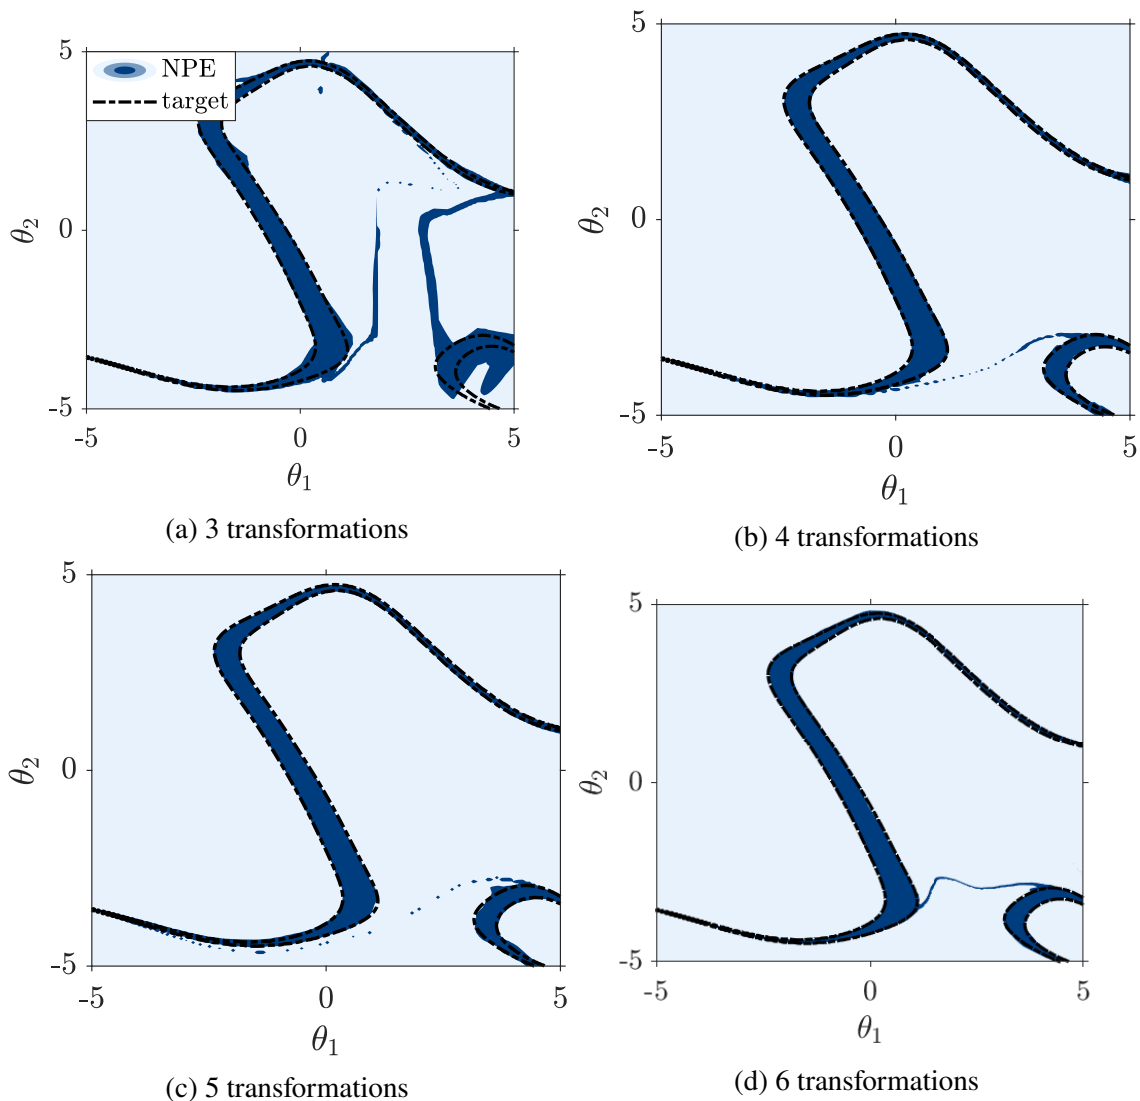


Figure 32: NPE for various number of transformations compared with target posterior level curve. Light blue corresponds to density regions where $\pi_{\text{post}}^{\text{NPE}}(\boldsymbol{\theta}|u^{\text{obs}}) < 0.01$, dark blue corresponds to density regions where $\pi_{\text{post}}^{\text{NPE}}(\boldsymbol{\theta}|u^{\text{obs}}) \geq 0.01$. Dashed-dot black curve corresponds to the target posterior level curve $\pi_{\text{post}}(\boldsymbol{\theta}|u^{\text{obs}}) = 0.01$.

As expected, the NPE with 3 transformations encounters difficulties in representing the target posterior, whereas comparable satisfying results are found with 4, 5 and 6 stacked transformations. The ECP is shown in Fig. 33. It can be seen that the calibration is great in any case, but the estimator seems to be slightly conservative in the case of 3 transformations, which can be observed in Fig.32a.




Figure 33: Expected coverage probability for various number of transformations. ECP below the perfect diagonal indicates an overconfident posterior, whereas an ECP above the diagonal indicates a conservative posterior.

In the following, the number of transformations of the NPE will be fixed to 6 to ensure a sufficient expressivity.

9.2.2 Size of the dataset

In this subsection, the NPE will be created with six different training set sizes. The size of the validation dataset and the testing dataset are respectively fixed to $N_{\text{valid}} = 30000$ and $N_{\text{test}} = 3000$. In Fig. 34 is shown the evolution of both the training loss and the validation loss with respect to the number of epochs for different size of training dataset.

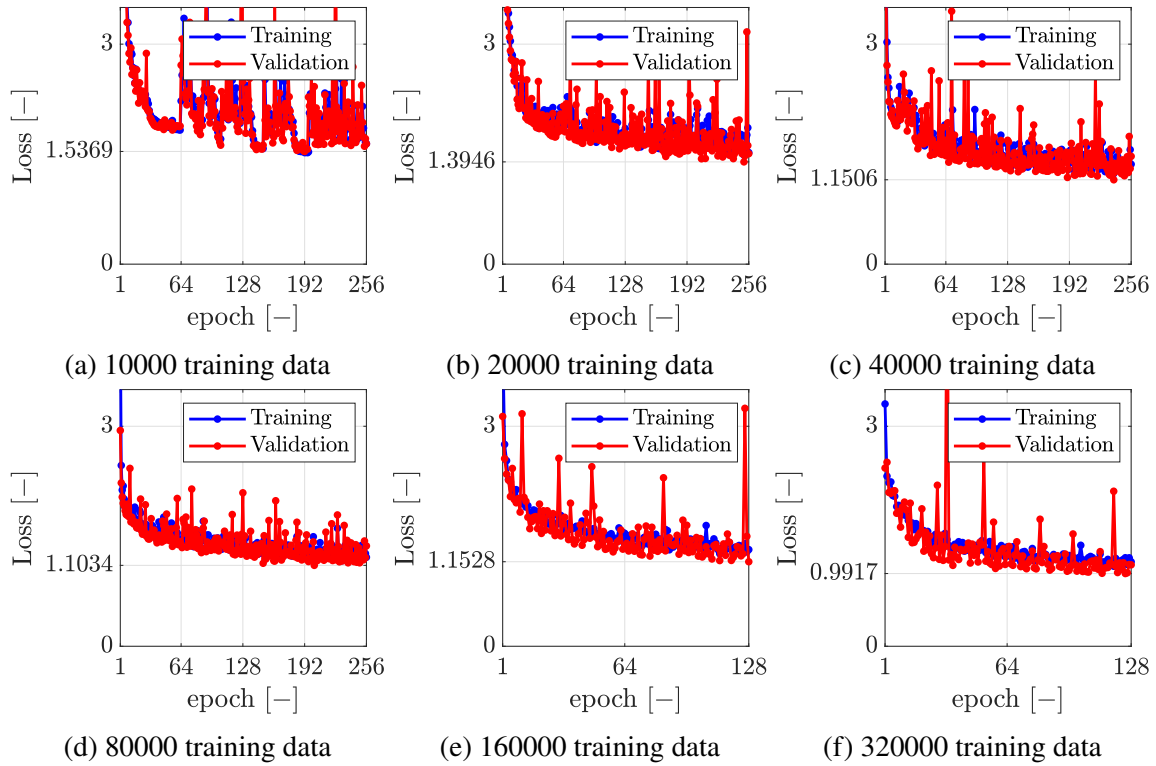


Figure 34: Evolution of the training loss (in blue) and of the validation loss (in red) with respect to the number of epochs. The minimum of the validation loss is indicated for each training dataset size.

As it can be observed, the minimum of the validation loss seems to decrease with the size of the training dataset. However, the relative change in the minimum validation loss seems to decrease which suggests that it eventually converges.

For all sizes of training dataset, the model which minimizes the validation loss can be used to evaluate the posterior estimator and to compare it with the target posterior. Figure 35 shows the posterior estimator for all training dataset sizes.

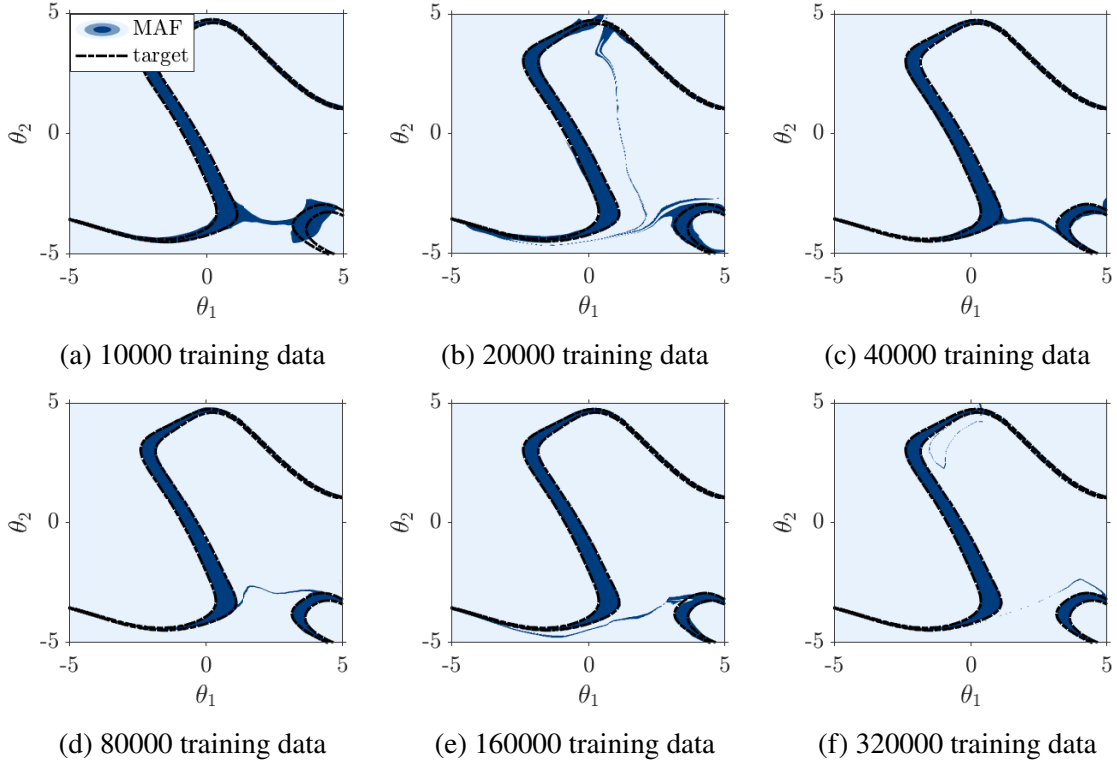


Figure 35: Posterior estimator with MAF $\pi^{\text{MAF}}(\boldsymbol{\theta}|u^{\text{obs}})$ for different training dataset sizes. Light blue corresponds to regions where $\pi^{\text{MAF}}(\boldsymbol{\theta}|u^{\text{obs}}) < 0.01$; dark blue corresponds to regions where $\pi^{\text{MAF}}(\boldsymbol{\theta}|u^{\text{obs}}) \geq 0.01$. The target density level curve $\pi_{\text{post}}(\boldsymbol{\theta}|u^{\text{obs}}) = 0.01$ is shown in black.

One observes that even for a small dataset, the global tendency of the posterior is retrieved. However, as expected, the quality of the posterior estimator seems to increase with the training dataset. In order to quantify the quality of the posterior estimator rigorously, the ECP is shown in Fig. 36. As it can be observed, the ECP is close to the diagonal even for small datasets which indicates that great estimators can be built without having to generate a lot of data. One notices however that the larger is the training dataset, the better is the posterior estimator.




Figure 36: Expected coverage probability (ECP) for several training dataset sizes. ECP below the perfect diagonal indicates an overconfident posterior, whereas an expected coverage above the diagonal indicates a conservative posterior.

10 Testing phase

Now that the methods have been successfully validated, they can be used to solve the five-dimensional problem described in Sec. 8.2.2 to infer the friction field.

10.1 Adaptive Metropolis

The sampling method that will be used is the AM algorithm as it has been shown to be less dependent on the choice of initial proposal covariance matrix, which contains arbitrariness. The initial proposal covariance matrix is set to $[\Sigma]^{(0)} = [I]$. In Fig. 37 is shown the states of the chain with respect to the iterations of the AM algorithm. As it can be observed, the acceptance rate is extremely small for the first few thousands iterations. This is due to the fact that (i) the chain starts at the origin which is likely to be a low density region of the posterior; (ii) the adaptive covariance matrix needs a large number of states to converge. Figure 38 shows that, for the testing problem, ten thousand states are sufficient for the convergence of the AM algorithm. This is one order of magnitude below what was observed during the validation phase despite the increased dimensionality. This suggests that the five-dimensional posterior is less complex than the two-dimensional one, which could be due to the choice of prior. Once the convergence is reached, the states can be considered as samples from the posterior distribution. Results of this sampling are shown in Sec. 10.3 together with neural posterior estimation results.

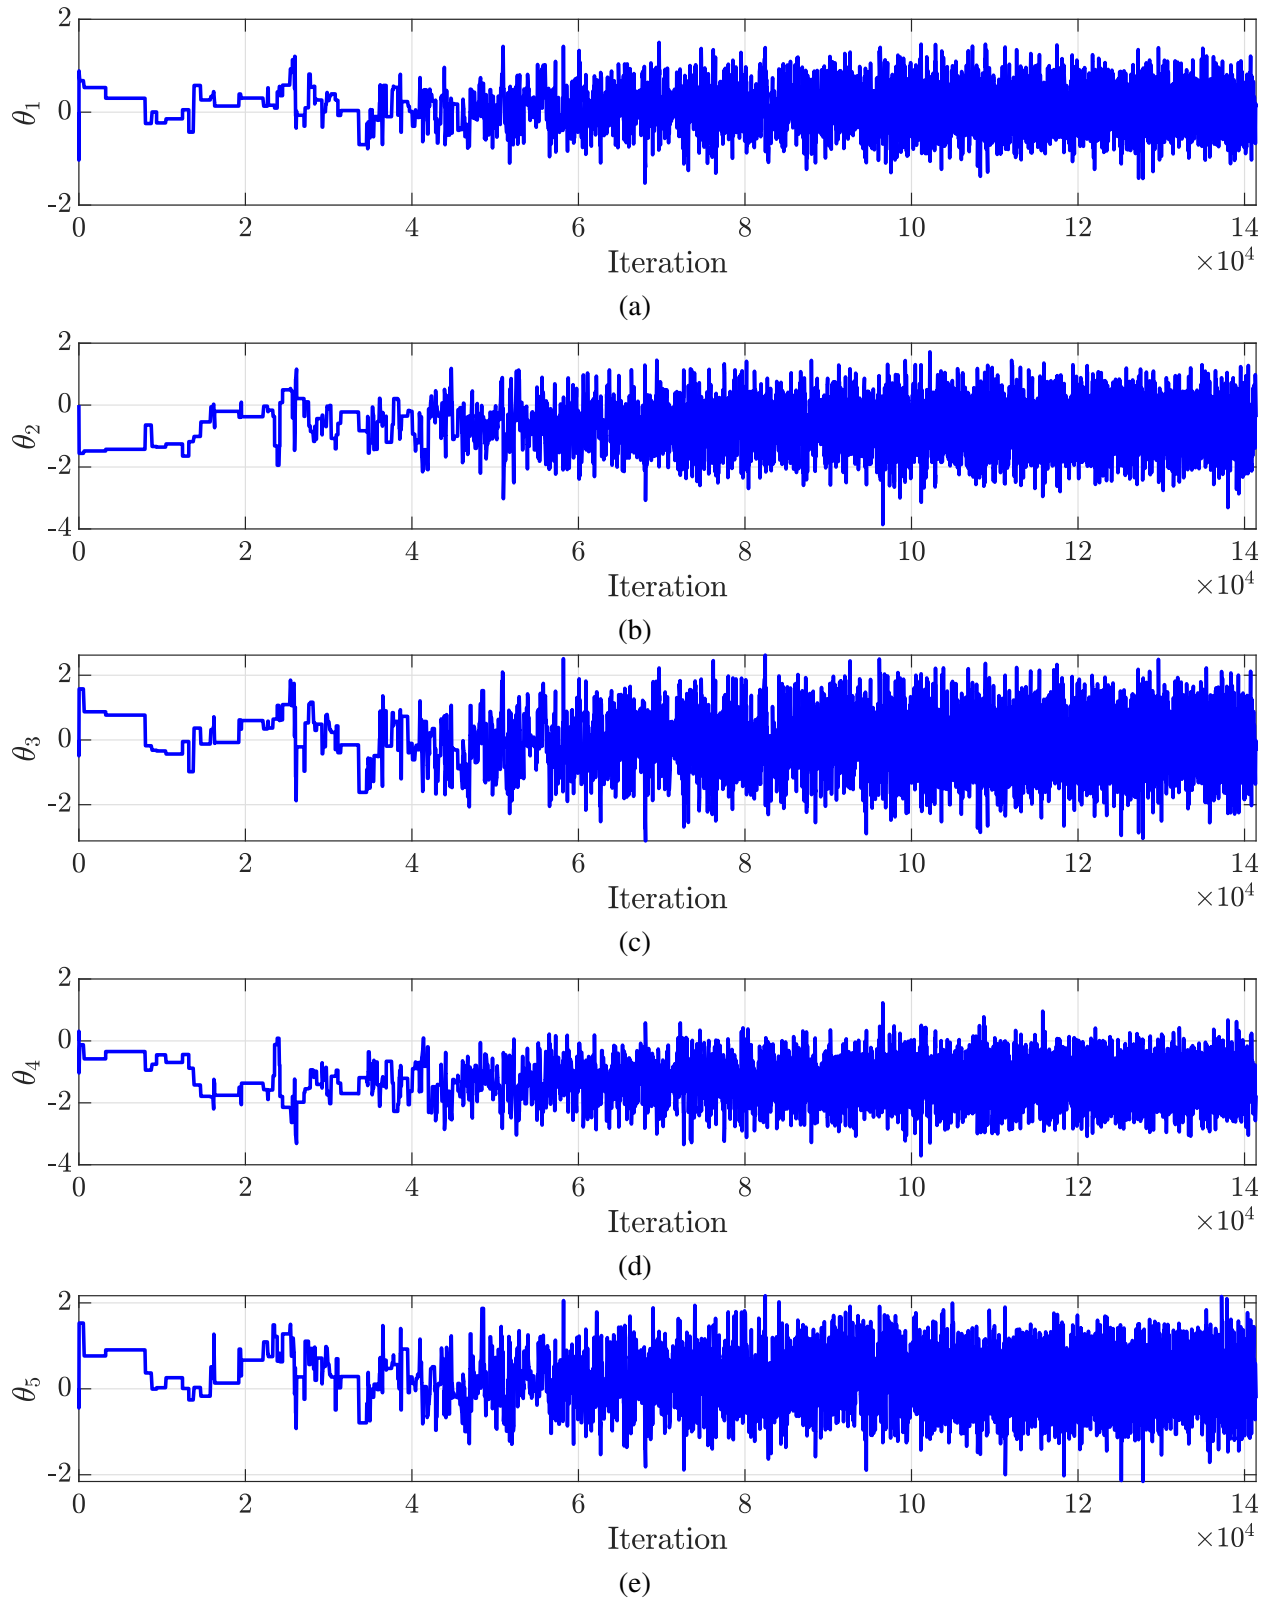


Figure 37: Evolution of the states with the iterations of the AM algorithm for the five-dimensional problem, with $[\Sigma]^{(0)} = [I]$.




Figure 38: Mean statistics of the states in the chain generated by the AM algorithm applied to the five-dimensional testing problem. Grey area estimate the 95% confidence interval, based on the central limit theorem.

10.2 Neural posterior estimation

The five-dimensional testing problem being totally different from the validation problem, the training has to be reiterated. The same procedure will be followed, starting from the minimization of the validation loss for several size of training set, towards the evaluation of the ECP. The number of transformations in the MAF is fixed to 6 as it has been shown during the validation phase that such number of transformations are sufficient to render the NPE highly expressive. Figure 39 shows the minimization of the loss during the training of the neural networks. Here again, the minimum reached value of the validation loss seems to decrease as the size of the training set increases. However, the relative decrements of the minimum validation loss between two successive training set sizes decreases, which suggests that it eventually converges. The models that minimize the validation loss for all training set sizes can be retrieved and their quality can be assessed with the ECP. In Fig. 40 is shown the ECP for all the training set sizes that have been considered. As it can be seen, the posterior estimator is slightly overconfident for small training set, but converge towards perfectly calibrated estimator as the training set size increases. In the following, the MAFs that correspond to the training dataset of 10000 data and of 320000 data will be sampled from. They will be compared together with the sampling resulting from the AM algorithm.

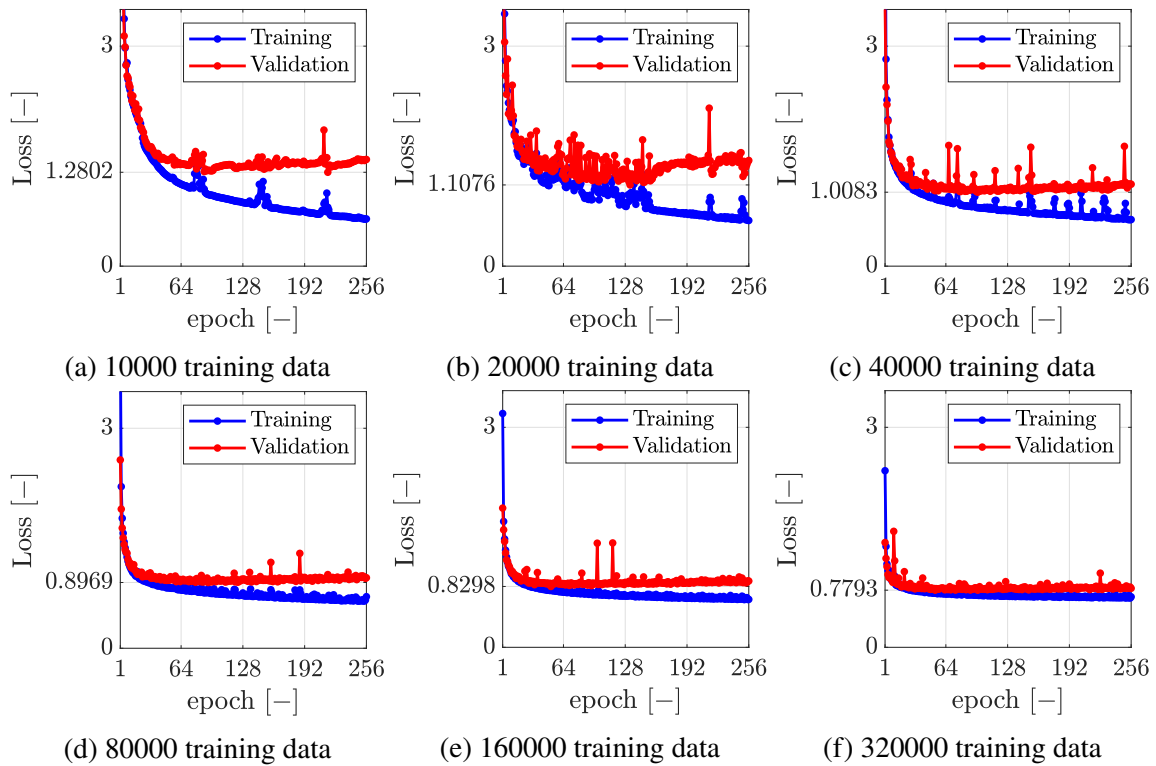


Figure 39: Evolution of the training loss (in blue) and of the validation loss (in red) with respect to the number of epochs for the five-dimensional problem. The minimum of the validation loss is indicated for each training dataset size.




Figure 40: Expected coverage probability (ECP) for several training dataset sizes for the five-dimensional problem. An ECP below the perfect diagonal indicates an overconfident posterior, whereas an ECP above the diagonal indicates a conservative posterior.

10.3 AM and NPE results

In this subsection, the samples resulting from the AM and the NPE methods will be shown and compared. This can be done by representing the bivariate plot, that is, samples in each bidimensional parameter space $\{\theta_i, \theta_j\}$. This is shown in Fig. 41. Plots that are in the diagonal correspond to the marginal probability estimations ($\pi_{\text{post}}(\theta_i|\mathbf{u}^{\text{obs}})$, $i = 1, \dots, 5$). Lower diagonal plots are the bivariate ($\pi_{\text{post}}(\theta_i, \theta_j|\mathbf{u}^{\text{obs}})$, $i = 1, \dots, 5, j \neq i$). As it can be observed, the sampling leads to very similar distributions for the three cases. However, one could notice that with only 10000 training data, some of the marginals are slightly overconfident (see $\pi_{\text{post}}(\theta_3|\mathbf{u}^{\text{obs}})$ and $\pi_{\text{post}}(\theta_5|\mathbf{u}^{\text{obs}})$), and some of the bivariate are slightly shifted (see $\pi_{\text{post}}(\theta_2, \theta_3|\mathbf{u}^{\text{obs}})$ and $\pi_{\text{post}}(\theta_2, \theta_5|\mathbf{u}^{\text{obs}})$).

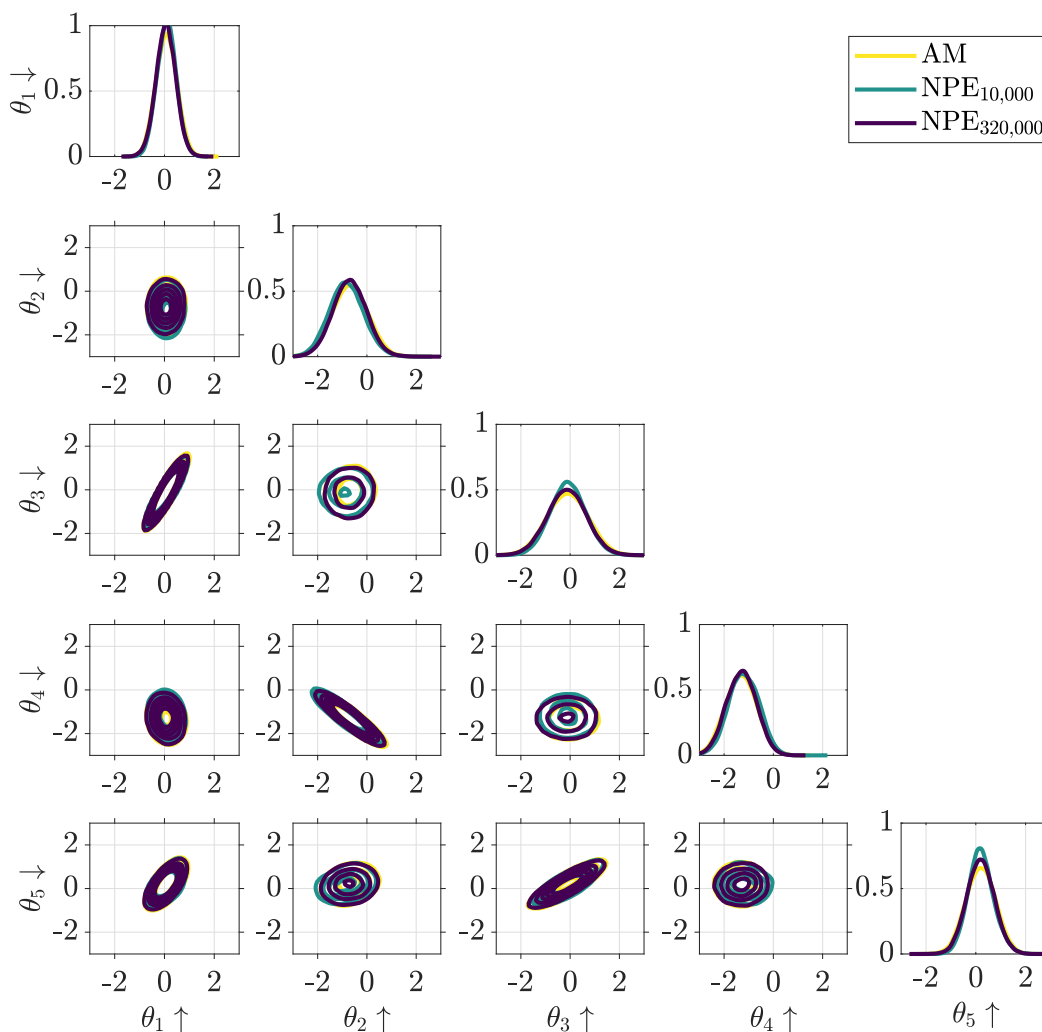




Figure 41: Estimated bivariate distributions from samples drawn with AM and NPE.

All of these samples are directly related with friction fields $\hat{\nu}_b$. During the AM algorithm, the posterior is evaluated at the samples. The maximum a posteriori (MAP) can therefore be estimated by the state which maximizes the posterior among all the samples. Similarly, NPEs can not only


be sampled from, but also their density can be evaluated. The MAP can therefore be estimated by the sample which maximizes the NPE density. In Fig. 42 is shown the MAP friction field $\hat{\nu}_b^{\text{MAP}}$ for AM and NPE, with confidence intervals estimated from Monte Carlo samples.



(a) AM



(b) NPE (10000 training data)



(c) NPE (320000 training data)

Figure 42: Inverted friction field from samples drawn with AM and NPE. Grey areas represent respectively (from the darker to the lighter) 33%, 50% and 95% confidence intervals, estimated from 100000 Monte carlo samples.

As it can be seen, the inversion of the friction field gives similar results for AM than for NPE. Moreover, NPE related to the training dataset of 10000 data leads to similar results than for 320000 data. One could notice that the friction field is uncertain for small x . This is because the noise is

absolute, that is, the noise is independent of the velocity. Since the velocity is small for small x , the noise destroys the informations of observations at small x .

Conclusion

This work was concerned with the inference of a space dependent friction field from partial observations of the ice velocity. In this context, an inverse problem has been formulated in a probabilistic framework that allows to quantify uncertainty.

The inversion of a continuous space dependent field such as the friction field leads to an infinite-dimensional inverse problem which cannot be solved in practice. This infinite dimensional inverse problem has been turned into a low dimensional one by considering the friction field to be a random field parametrized by the Karhunen-Loève expansion parameters.

Both sampling and deep-learning based methods have been used. Metropolis-Hastings (MH) and adaptive Metropolis (AM) algorithms, as well as neural posterior estimators (NPE) based on normalizing flows have been studied. These methods require the forward model to be evaluated a large number of times, which is challenging as the forward model is computationally expensive despite being based on the shallow-shelf approximation. An inexpensive polynomial chaos surrogate model has been constructed to render the methods computationally tractable.

A validation phase has shown that the AM algorithm provides a better robustness than MH. The performances of the NPE has been shown to be dependent on the size of the dataset, as well as on its intrinsic architecture.

A testing phase, in which a five-dimensional inverse problem has been formulated to infer the friction field from observations of the ice velocity, has been solved with AM and NPE.

The AM algorithm required $O(10^5)$ calls to the forward model to converge. This suggests that MCMC-based methods such as AM are intractable if the forward model is computationally expensive.

The construction of a NPE requires calls to the forward model exclusively during the creation of the dataset. As it has been observed, satisfying NPEs can be constructed with small datasets ($O(10^4)$ data), and can then be sampled from without the forward model. Such estimators can therefore be constructed directly with the SSA forward model, without having to build a polynomial chaos surrogate. Moreover, such estimators are robust to any changes in the observations as they are amortized estimators.

In terms of inverted friction fields, the NPE trained with 10000 training samples, NPE trained with 320000 training samples and AM gave similar and satisfying results. NPE has therefore required less calls to the forward model than AM to attain satisfying results, in addition of allowing to draw an arbitrarily large number of samples once trained, irrespective of the complexity of the forward model.

It is concluded that the current state-of-the-art of inverse problem in glaciology can greatly benefit from deep-learning tools such as NPE.

Future work: A non exhaustive list of future work that could be considered to improve the current

work is listed below:

- Extend the current work to realistic large-scale implementations of ice sheet dynamics such as f.etch (Pattyn (2017)). This would require to extend the friction model to two dimensions;
- Assess the impact of mismatch between the ground truth correlation length (the correlation length of ν_b^*) and the correlation length of the truncated field $\hat{\nu}_b$ that is inferred;
- Study the effect of the choice of friction law;
- Explore other normalizing flow architectures.

A Parameters in the SSA equations

The parameters in the SSA equations are listed in Tab. 1.

Physical quantity	Symbol	Value	Units
Ice density	ρ	910	kg m^{-3}
Water density	ρ_w	1028	kg m^{-3}
Exponent in Glen and Weertman's laws	m, n	1/3	–
Ice accumulation rate	a	0.3	m y^{-1}
Rate factor in Glen's law	A	4.9×10^{-25}	$\text{Pa}^{-3} \text{s}^{-1}$
Gravity acceleration	g	9.81	m s^{-2}

Table 1: Physical quantities in SSA equations.

Figure 43 shows the geometry of the bedrock.






Figure 43: Geometry of the bedrock in the interval $x \in [0, 1000]$ km (Schoof (2007)).

B Complementary results


Figure 44 shows the complete evolution of the states with respect to the MH algorithm iterations, for two different proposal covariance matrices.




(a) $[\Sigma] = 0.1 \times [I]$



(b) $[\Sigma] = 0.1 \times [I]$




(c) $[\Sigma] = [I]$




(d) $[\Sigma] = [I]$

Figure 44: Chain of states for the Metropolis-Hastings (MH) algorithm. (a) θ_1 with $[\Sigma] = 0.1 \times [I]$; (b) θ_2 with $[\Sigma] = 0.1 \times [I]$; (c) θ_1 with $[\Sigma] = [I]$; (d) θ_2 with $[\Sigma] = [I]$.


Figure 45 shows the complete evolution of the states with respect to the AM algorithm iterations, for two different initial proposal covariance matrices.




(a) $[\Sigma] = 0.1 \times [I]$



(b) $[\Sigma] = 0.1 \times [I]$



(c) $[\Sigma] = [I]$



(d) $[\Sigma] = [I]$

Figure 45: Chain of states for the adaptive Metropolis (AM) algorithm. (a) θ_1 with $[\Sigma] = 0.1 \times [I]$; (b) θ_2 with $[\Sigma] = 0.1 \times [I]$; (c) θ_1 with $[\Sigma] = [I]$; (d) θ_2 with $[\Sigma] = [I]$.

References

- Armstrong McKay, D. I., Staal, A., Abrams, J. F., Winkelmann, R., Sakschewski, B., Loriani, S., Fetzner, I., Cornell, S. E., Rockström, J., and Lenton, T. M. (2022). Exceeding 1.5 c global warming could trigger multiple climate tipping points. *Science*, 377(6611):eabn7950.
- Bosten, A. et al. (2019). Towards a contact formulation for efficient numerical simulation of marine ice-sheet instabilities.
- Budd, W., Jenssen, D., and Smith, I. (1984). A three-dimensional time-dependent model of the west antarctic ice sheet. *Annals of Glaciology*, 5:29–36.
- Bulthuis, K., Arnst, M., Sun, S., and Pattyn, F. (2019). Uncertainty quantification of the multi-centennial response of the antarctic ice sheet to climate change. *The Cryosphere*, 13(4):1349–1380.
- Gelman, A., Roberts, G., and Gilks, W. (1996). Efficient metropolis jumping rules. *Bayesian Statistics*, 5:599–608.
- Germain, M., Gregor, K., Murray, I., and Larochelle, H. (2015). Made: Masked autoencoder for distribution estimation. In *International conference on machine learning*, pages 881–889. PMLR.
- Ghanem, R. G. and Spanos, P. D. (2003). *Stochastic finite elements: a spectral approach*. Courier Corporation.
- Glen, J. W. (1955). The creep of polycrystalline ice. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 228(1175):519–538.
- Goldberg, D. N. and Sergienko, O. V. (2011). Data assimilation using a hybrid ice flow model. *The Cryosphere*, 5(2):315–327.
- Gregov, T., Pattyn, F., and Arnst, M. (09 June 2022). A primal-dual formulation for numerical simulations of marine ice sheets with various friction laws.
- Greve, R. and Blatter, H. (2009). *Dynamics of ice sheets and glaciers*. Springer Science & Business Media.
- Haario, H., Saksman, E., and Tamminen, J. (1999). Adaptive proposal distribution for random walk metropolis algorithm. *Computational statistics*, 14:375–395.
- Haario, H., Saksman, E., and Tamminen, J. (2001). An adaptive metropolis algorithm. *Bernoulli*, pages 223–242.
- Hastings, W. K. (1970). Monte carlo sampling methods using markov chains and their applications.
- Hermans, J., Delaunoy, A., Rozet, F., Wehenkel, A., and Louppe, G. (2021). Averting a crisis in simulation-based inference. *arXiv preprint arXiv:2110.06581*.
- Huybrechts, P. (2009). West-side story of antarctic ice. *Nature*, 458(7236):295–296.

- IPCC (2019). Summary for policymakers. In Portner, H.-O., Roberts, D., Masson-Delmotte, V., Zhai, P., Tignor, M., Poloczanska, E., Mintenbeck, K., Alegria, A., Nicolai, M., Okem, A., Petzold, J., Rama, B., and Weyer, N., editors, *IPCC Special Report on the Ocean and Cryosphere in a Changing Climate*. In press.
- Isaac, T., Petra, N., Stadler, G., and Ghattas, O. (2015). Scalable and efficient algorithms for the propagation of uncertainty from data through inference to prediction for large-scale problems, with application to flow of the antarctic ice sheet. *Journal of Computational Physics*, 296:348–368.
- Jaynes, E. T. (1957). Information theory and statistical mechanics. *Physical review*, 106(4):620.
- Kingma, D. P., Salimans, T., Jozefowicz, R., Chen, X., Sutskever, I., and Welling, M. (2016). Improved variational inference with inverse autoregressive flow. *Advances in neural information processing systems*, 29.
- Kobyzev, I., Prince, S. J., and Brubaker, M. A. (2020). Normalizing flows: An introduction and review of current methods. *IEEE transactions on pattern analysis and machine intelligence*, 43(11):3964–3979.
- Larour, E., Schiermeier, J., Rignot, E., Seroussi, H., Morlighem, M., and Paden, J. (2012). Sensitivity analysis of pine island glacier ice flow using issm and dakota. *Journal of Geophysical Research: Earth Surface*, 117(F2).
- Larour, E., Utke, J., Csatho, B., Schenk, A., Seroussi, H., Morlighem, M., Rignot, E., Schlegel, N., and 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.
- Le Maître, O. and Knio, O. M. (2010). *Spectral methods for uncertainty quantification: with applications to computational fluid dynamics*. Springer Science & Business Media.
- Lévy, P. (1965). *Processus stochastiques et mouvement brownien*. Gauthier-Villars.
- 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.
- MacAyeal, D. R. (1993). A tutorial on the use of control methods in ice-sheet modeling. *Journal of Glaciology*, 39(131):91–98.
- MacAyeal, D. R., Bindshadler, R. A., and Scambos, T. A. (1995). Basal friction of ice stream e, west antarctica. *Journal of Glaciology*, 41(138):247–262.
- Marzouk, Y. M., Najm, H. N., and Rahn, L. A. (2007). Stochastic spectral methods for efficient bayesian solution of inverse problems. *Journal of Computational Physics*, 224(2):560–586.
- Metropolis, N., Rosenbluth, A. W., Rosenbluth, M. N., Teller, A. H., and Teller, E. (1953). Equation of state calculations by fast computing machines. *The journal of chemical physics*, 21(6):1087–1092.

- Morland, L. (1987). Unconfined ice-shelf flow. In *Dynamics of the West Antarctic Ice Sheet: Proceedings of a Workshop held in Utrecht, May 6–8, 1985*, pages 99–116. Springer.
- Morlighem, M., Rignot, E., Seroussi, H., Larour, E., Ben Dhia, H., and 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).
- Morlighem, M., Seroussi, H., Larour, E., and Rignot, E. (2013). Inversion of basal friction in antarctica using exact and incomplete adjoints of a higher-order model. *Journal of Geophysical Research: Earth Surface*, 118(3):1746–1753.
- Papamakarios, G., Nalisnick, E., Rezende, D. J., Mohamed, S., and Lakshminarayanan, B. (2021). Normalizing flows for probabilistic modeling and inference. *The Journal of Machine Learning Research*, 22(1):2617–2680.
- Papamakarios, G., Pavlakou, T., and Murray, I. (2017). Masked autoregressive flow for density estimation. *Advances in neural information processing systems*, 30.
- Pattyn, F. (2017). Sea-level response to melting of antarctic ice shelves on multi-centennial timescales with the fast elementary thermomechanical ice sheet model (f. etish v1. 0). *The Cryosphere*, 11(4):1851–1878.
- Pattyn, F., Schoof, C., Perichon, L., Hindmarsh, R., Bueler, E., De Fleurian, B., Durand, G., Gagliardini, O., Gladstone, R., Goldberg, D., et al. (2012). Results of the marine ice sheet model intercomparison project, mismip. *The Cryosphere Discussions*, 6(1):267–308.
- Pavliotis, G. A. (2014). *Stochastic processes and applications: diffusion processes, the Fokker-Planck and Langevin equations*, volume 60. Springer.
- Petra, N., Zhu, H., Stadler, G., Hughes, T. J., and Ghattas, O. (2012). An inexact gauss-newton method for inversion of basal sliding and rheology parameters in a nonlinear stokes ice sheet model. *Journal of Glaciology*, 58(211):889–903.
- Pollard, D., DeConto, R. M., and Alley, R. B. (2015). Potential antarctic ice sheet retreat driven by hydrofracturing and ice cliff failure. *Earth and Planetary Science Letters*, 412:112–121.
- Potsdam (2022). Risk of passing multiple climate tipping points escalates above 1.5°C global warming.
- Rezende, D. and Mohamed, S. (2015). Variational inference with normalizing flows. In *International conference on machine learning*, pages 1530–1538. PMLR.
- Rozet, F. (2022). Zuko: Normalizing flows in pytorch. Available at Zenodo: <https://doi.org/10.5281/zenodo.7625672>.
- Schoof, C. (2005). The effect of cavitation on glacier sliding. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 461(2055):609–627.
- Schoof, C. (2007). Ice sheet grounding line dynamics: Steady states, stability, and hysteresis. *Journal of Geophysical Research: Earth Surface*, 112(F3).

- Schoof, C. and Hewitt, I. (2013). Ice-sheet dynamics. *Annual Review of Fluid Mechanics*, 45:217–239.
- Stuart, A. M. (2010). Inverse problems: a bayesian perspective. *Acta numerica*, 19:451–559.
- Tabak, E. G. and Turner, C. V. (2013). A family of nonparametric density estimation algorithms. *Communications on Pure and Applied Mathematics*, 66(2):145–164.
- Tabak, E. G. and Vanden-Eijnden, E. (2010). Density estimation by dual ascent of the log-likelihood. *Communications in Mathematical Sciences*, 8(1):217–233.
- Tarantola, A. (2005). *Inverse problem theory and methods for model parameter estimation*. SIAM.
- Tikhonov, A. N. (1963). On the solution of ill-posed problems and the method of regularization. In *Doklady akademii nauk*, volume 151, pages 501–504. Russian Academy of Sciences.
- Vieli, A. and Payne, A. J. (2003). Application of control methods for modelling the flow of pine island glacier, west antarctica. *Annals of Glaciology*, 36:197–204.
- Villa, U., Petra, N., and Ghattas, O. (2021). Hippylib: an extensible software framework for large-scale inverse problems governed by pdes: part i: deterministic inversion and linearized bayesian inference. *ACM Transactions on Mathematical Software (TOMS)*, 47(2):1–34.
- Weertman, J. (1957). On the sliding of glaciers. *Journal of glaciology*, 3(21):33–38.
- Xiu, D. and Karniadakis, G. E. (2002). The wiener–askey polynomial chaos for stochastic differential equations. *SIAM journal on scientific computing*, 24(2):619–644.