This manuscript is a non-peer reviewed preprint submitted to EarthArXiv

The manuscript was written by the IGM main developer team that includes
Guillaume Jouvet (UNIL), Brandon Finley (UNIL),
Thomas Gregov (UNIL), Sebastian Rosier (UZH)
from the Universities of Lausanne (UNIL) and Zürich (UZH)

IGM: an accessible, modular, differentiable, and GPU-accelerated high-order ice flow model

The IGM developers*

Abstract. We present the Instructed Glacier Model (IGM, v3.2), an open-source framework for simulating glacier evolution from single-glacier to mountain-range scales. IGM is built on a single design principle: all physical processes, including ice flow, surface mass balance, thermodynamics, and mass conservation, are expressed as short sequences of operations on raster grids. This workflow runs natively on GPUs and scales efficiently to large domains. Within this framework, higher-order ice flow is handled by a differentiable mapping trained by automatic differentiation to satisfy the Blatter–Pattyn equations, reducing the computational cost dramatically compared to traditional CPU-based approaches. The framework is structured around five guiding principles: (i) *accessibility*, through a Python package and a simple YAML-based interface; (ii) *modularity*, enabling modellers to extend or replace individual components without modifying the core code; (iii) *reproducibility*, through automated configuration logging and continuous integration; (iv) *scalability*, enabling large-domain simulations on GPU; and (v) *explorability*, facilitating ensemble runs, parameter sweeps, and model calibration. We describe the physical and numerical foundations of IGM, its software architecture, and illustrate its capabilities through benchmark tests and real-world glacier simulations.

1 Introduction

Glacier evolution models are essential tools for reconstructing past glacier–climate interactions and for projecting future ice loss (Rounce et al., 2023). Over the last two decades, the glaciological community has devoted substantial effort to developing these models (see Zekollari et al., 2022, for a review). These models are designed to represent a wide range of physical processes, including ice flow, thermodynamics, subglacial hydrology, and their interactions with various factors such as atmospheric conditions (e.g., climate-driven surface mass balance), the Earth's lithosphere, and the ocean (e.g., iceberg calving or subsurface melting). Prominent examples of such models include full-Stokes such as Elmer/Ice (Gagliardini et al., 2013) for general use, higher-order mdoels such as PISM (Winkelmann et al., 2011), CISM (Lipscomb et al., 2019), ISSM (Larour et al., 2012), which are popular in the ice-sheet modeling community, or SIA-based models such as OGGM (Maussion et al., 2019) or PyGEM (Rounce et al., 2020), which were designed for global glacier modeling. However, increasing model complexity and the growing volume of observational data to assimilate come with rising computational burdens. Parallel computing and automatic differentiation are promising ways to overcome these limitations.

In recent years, there has been a growing interest in employing Graphics Processing Units (GPUs) to tackle the computational bottleneck in ice flow modeling. GPUs, with many more cores than CPUs, can achieve substantial speed improvements when numerical methods can be subdivided into massively parallel tasks, a particularly challenging requirement for ice flow because of its viscous, diffusive nature. In glaciology, GPU acceleration has been demonstrated through explicit time integration of the Second Order Shallow Ice Approximation (SIA) (Brædstrup et al., 2014), the SIA (Višnjević et al., 2020), and the Stokes equations (Räss et al., 2020), as well as through GPU-native solvers such as MALI (Tezaur et al., 2015; Hoffman et al., 2018; Watkins et al., 2023). The rapid development of libraries such as TensorFlow and PyTorch, driven by the rise of artificial intelligence, has further popularized GPU computing in science. Beyond raw hardware acceleration, machine learning is increasingly being leveraged to aid scientific computing more broadly, and the glaciological community has begun to explore its potential. Examples include data-driven emulators for subglacial hydrology (Brinkerhoff et al., 2021; Verjans and Robel, 2024), basal melt (Burgard et al., 2023; Rosier et al., 2023), and ice flow (Jouvet et al., 2021; Koo and Rahnemoonfar, 2024), but also physics-informed neural networks (Riel et al., 2021;

*The list of developers, their affiliations and contributions appears at the end of the paper.

Bolibar et al., 2023; Howard et al., 2023; Jouvet and Cordonnier, 2023; Cheng et al., 2025; Wang et al., 2025) and hybrid approaches (He et al., 2023).

Building on these advances, the Instructed Glacier Model (IGM), originally introduced by Jouvet et al. (2021), follows a simple design principle. Glacier modeling consists, to a large extent, of operations on large raster fields representing glacier characteristics over a regular spatial grid. Such workloads map naturally onto GPU architectures, but are poorly suited to algorithms that rely on inverting large matrices. IGM therefore organizes each physical component as short sequences of rasterized operations, keeping the sequential, non-parallelizable portion of each update as short as possible. With the exception of ice flow, all components (e.g., thermodynamics, surface mass balance, and mass conservation) are naturally solved raster-wise. Ice flow is the one component that does not fit the raster-wise pattern, since solving the higher-order ice flow equations would normally require inverting a large system. IGM circumvents this by training a physics-informed neural network to satisfy the higher-order ice flow physics (Jouvet and Cordonnier, 2023; Cordonnier et al., 2023), so that the sequential cost reduces to the depth of the network. Advantageously, this neural formulation also provides automatic differentiation, which enables physics-based training and data assimilation by model inversion (Jouvet, 2023).

In this paper, we present the design of IGM, version 3.2.0, as a stable and benchmarked modeling platform for glaciers (including tidewater ones, Fig. 1), with an emphasis on its physical processes, its numerical methods, its software architecture, its modularity, and its integration with the Hydra and Optuna frameworks for reproducible ensemble simulations. This paper is structured as follows. Sections 2 and 3 describe the physical model and its numerical discretization. Data assimilation is briefly discussed in Section 4. We then describe the software design in Section 5, and applications in Section 6. Last, Section 7 discusses community aspects and Section 8 discusses IGM potential for education, outreach, and visualization. Finally, Section 9 outlines our vision for the ongoing development of IGM and its potential for future applications.

2 Physics

In the following, we write $\boldsymbol{x} = (x, y, z)$ for the spatial coordinates and t for the time coordinate. We use ∇ for the gradient operator, acting in two or three dimensions depending on context: when applied to fields defined on (x, y), it denotes (∂_x, ∂_y); when applied to fields defined on (x, y, z), it denotes $(\partial_x, \partial_y, \partial_z)$. We use $b = b(x, y, t)$, $s = s(x, y, t)$, and $h = h(x, y, t)$ to represent the bedrock elevation, the surface elevation, and ice thickness, with $h = s - b$. When considering tidewater glaciers, we denote the lower ice surface by $l = l(x, y, t)$, taking it to coincide with the bed where the ice is grounded and, where the ice is floating, assuming local hydrostatic equilibrium:

$$l = \max\left(b, z_{\mathrm{wl}} - \frac{\rho_{\mathrm{i}}}{\rho_{\mathrm{w}}} h\right), \tag{1}$$

where ρ_{w} and ρ_{i} are the densities of water and ice, respectively. The three-dimensional velocity field of the ice is written as $\boldsymbol{v} = \boldsymbol{v}(\boldsymbol{x}, t)$, with components (u, v, w). We also use $\boldsymbol{u} = (u, v)$ for the horizontal components.

We now describe the physical processes represented in IGM (Fig. 1), each implemented as an independent module in IGM (Section 5), including both essential components (such as the ice thickness, surface mass balance, and ice flow), and other components (such as the enthalpy or Lagrangian particles).

2.1 Ice thickness

The glacier geometry evolves according to mass conservation, relating the variation in ice thickness to ice dynamics and mass balance at the upper and lower surfaces:

$$\frac{\partial h}{\partial t} + \nabla \cdot (\bar{\boldsymbol{u}} h) = \dot{a}_{\mathrm{s}} + \dot{a}_{\mathrm{b}}, \tag{2}$$

where $\bar{\boldsymbol{u}}$ denotes the vertically-averaged horizontal velocity, and $\dot{a}_{\mathrm{s}}$ and $\dot{a}_{\mathrm{b}}$ denote the surface and basal mass balances, respectively. Similarly, when the glacier terminates in water, the calving front Γ_{cf} migrates under the competition between ice advection and calving, with a normal velocity

$$v_{\mathrm{cf}} = \bar{\boldsymbol{u}} \cdot \boldsymbol{n} - \dot{a}_{\mathrm{cf}}, \tag{3}$$

where $\boldsymbol{n}$ is the outward unit normal to Γ_{cf} and $\dot{a}_{\mathrm{cf}}$ is the calving rate.

2.2 Mass balance components

The surface mass balance $\dot{a}_{\mathrm{s}}$ is the net rate of mass gain or loss at the glacier surface per unit time. IGM provides three SMB modules of increasing physical complexity: (i) a simple altitude-dependent parametrization in which the SMB varies piecewise-linearly with elevation around the equilibrium-line altitude (Meier et al., 1971); (ii) a monthly temperature-index model following OGGM (Maussion et al., 2019), calibrated against geodetic mass balance observations (Hugonnet et al., 2021); and (iii) a positive degree-day scheme (Hock, 2003) with explicit snowpack tracking and separate melt factors for snow and ice. The mathematical formulations of these three schemes are given in Appendix A.

When the thermodynamics module is enabled (Section 2.4), $\dot{a}_{\mathrm{b}}$ is derived from the enthalpy field through a heat balance at the ice–bed interface (see Appendix C).

In the presence of water, the calving rate can be computed using standard parameterizations such as water depth (Brown et al., 1982), height above flotation (Van der Veen, 1996; Vieli et al., 2001), or principal strain rates (Levermann et al., 2012).

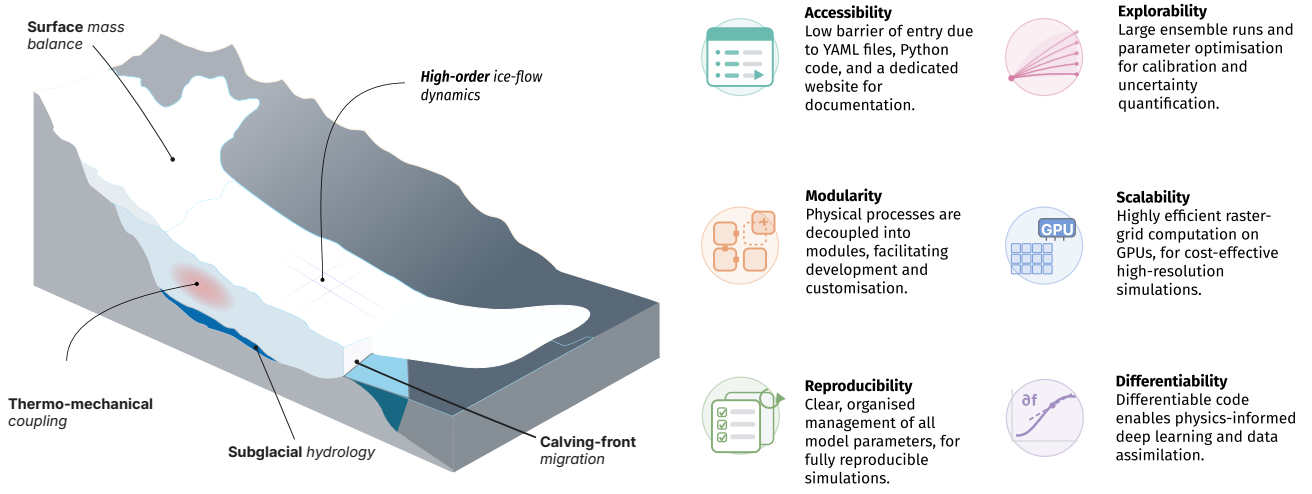

Figure 1. Overview of IGM, with a schematic of a glacier showing the coupled physical processes (left) and the framework's design principles (right) that together enable efficient glacier evolution simulations.

2.3 Ice flow

Over long time scales, on the order of years, ice behaves as an incompressible viscous fluid governed by Stokes flow, since inertia in the momentum-balance equation is negligible. Ice flow is therefore described by the following partial differential equations:

$$\nabla \cdot \boldsymbol{\sigma} + \rho_{\mathrm{i}} \boldsymbol{g} = \boldsymbol{0}, \tag{4a}$$

$$\nabla \cdot \boldsymbol{v} = 0, \tag{4b}$$

where $\boldsymbol{\sigma}$ denotes the Cauchy stress tensor and $\boldsymbol{g}$ is the gravitational acceleration vector, pointing downward. Following Glen's flow law (Glen, 1955), ice is modeled as a generalized Newtonian fluid with a power-law rheology:

$$\boldsymbol{\sigma}(\boldsymbol{v}, p) = -p\boldsymbol{I} + 2\eta \boldsymbol{D}(\boldsymbol{v}), \tag{5a}$$

$$\eta = \frac{1}{2} A^{-\frac{1}{n}} |\boldsymbol{D}(\boldsymbol{v})|^{\frac{1}{n}-1}. \tag{5b}$$

Here, p is the ice pressure, $\boldsymbol{I}$ is the identity tensor, η is the effective viscosity, and $\boldsymbol{D}(\boldsymbol{v}) = [\nabla \boldsymbol{v} + (\nabla \boldsymbol{v})^\top]/2$ is the strain-rate tensor. The notation $|\cdot|$ denotes the Frobenius norm, defined for any second-order tensor $\boldsymbol{X}$ by $|\boldsymbol{X}|^2 = (\boldsymbol{X} : \boldsymbol{X})/2$. The physical parameters A and $n > 1$ correspond to the Arrhenius factor and Glen's flow-law exponent, respectively. The former depends on the temperature of the ice, so that the thermal state can impact the ice-flow dynamics by modifying the effective viscosity; see Appendix C.

For the boundary conditions, the upper surface of the glacier, corresponding to the ice-air interface, is stress-free (neglecting atmospheric pressure):

$$\boldsymbol{\sigma} \cdot \boldsymbol{n} = \boldsymbol{0}, \tag{6}$$

where $\boldsymbol{n}$ denotes the outer normal vector to the surface. Along the lower surface of the glacier, which is in contact with the bedrock, we impose the following conditions:

$$\boldsymbol{v} \cdot \boldsymbol{n} = 0, \tag{7a}$$

$$\boldsymbol{\sigma} \cdot \boldsymbol{n} = -\tau_{\mathrm{ref}} \left(\frac{N}{N_{\mathrm{ref}}} \right)^q \left(\frac{|\boldsymbol{v}|}{u_{\mathrm{ref}}} \right)^m \frac{\boldsymbol{v}}{|\boldsymbol{v}|}, \tag{7b}$$

where $m > 0$ is the power-law exponent. The first equation is a non-penetration constraint, and the second takes the form of a nonlinear friction law relating the basal velocity to the basal shear stress. Here, we adopt a generalized power-law (Weertman–Budd) friction law (Weertman, 1957; Nye, 1969; Kamb, 1970; Fowler, 1981; Budd et al., 1979), written in terms of a reference shear stress τ_{ref} and a reference sliding velocity u_{ref}: by construction, the basal shear stress equals τ_{ref} when the basal speed equals u_{ref} and $N = N_{\mathrm{ref}}$. The quantity N denotes the basal effective pressure (Section 2.5), N_{ref} is a reference effective pressure, and $q \geq 0$ is the effective-pressure exponent that controls the sensitivity of basal friction to subglacial hydrology. Setting $q = 0$ removes the dependence on N and recovers the classical Weertman law, while $q > 0$ yields a Budd-type law. For $q = 0$, this formulation is equivalent to the classical $\boldsymbol{\sigma} \cdot \boldsymbol{n} = -c|\boldsymbol{v}|^{m-1}\boldsymbol{v}$ with $c = \tau_{\mathrm{ref}}/u_{\mathrm{ref}}^m$, but allows working with physical quantities that have a clear interpretation and direct units, rather than with the more abstract sliding coefficient c. IGM also implements other friction laws such as the regularized Coulomb law (Schoof, 2005, 2010; Gagliardini et al., 2007; Minchew and Joughin, 2020).

For marine-terminating glaciers, a stress boundary condition is applied along the calving front Γ_{cf}, i.e. the vertical ice face in contact with the water. Denoting by $\boldsymbol{n}$ the outward

horizontal unit normal to Γ_{cf} and by z_{wl} the water level, the ice is subject to hydrostatic water pressure on its submerged portion:

$$\boldsymbol{\sigma} \cdot \boldsymbol{n} = -p_{\mathrm{w}}\, \boldsymbol{n}, \tag{8a}$$

$$p_{\mathrm{w}} := \rho_{\mathrm{w}} g\, (z_{\mathrm{wl}} - z)_+, \tag{8b}$$

where $(\cdot)_+ := \max(\cdot, 0)$. On land-terminating glaciers, $\Gamma_{\mathrm{cf}} = \emptyset$ and this condition is absent.

Equations (4)–(8) describe a non-linear Stokes problem in which the unknowns are the 3D velocity field $\boldsymbol{v} = (u, v, w)$ and the pressure p. To simplify the problem, we make the 'cryostatic assumption' and neglect second-order terms in the aspect ratio of the ice domain within the strain-rate tensor (see, e.g., Schoof and Hindmarsh, 2010). By doing so, and invoking the incompressibility equation (4b), both the vertical velocity component w and the pressure p are eliminated from the momentum-conservation equation. The resulting model, commonly referred to as the Blatter–Pattyn model and known as a 'higher-order' model, conveniently transforms into a 3D non-linear elliptic equation solely for the horizontal velocity components $\boldsymbol{u} = (u, v)$ (Herterich, 1987; Blatter, 1995; Pattyn, 2003). This modification makes it easier to solve than the original Stokes model: only two unknowns (u, v) instead of four (u, v, w, p), and no divergence-free constraint.

As any ice flow model, Blatter–Pattyn can be written as an energy-minimization problem (Colinge and Rappaz, 1999; Jouvet, 2016, see appendix B), which means it satisfies $\langle \mathcal{J}'(\boldsymbol{u}), \delta\boldsymbol{u} \rangle = 0$ for all admissible variations $\delta\boldsymbol{u}$, where the associated convex functional $\mathcal{J}$ is defined by:

$$\mathcal{J}(\boldsymbol{u}) := \int_\Omega \frac{2A^{-\frac{1}{n}}}{1+1/n} |\boldsymbol{D}(\boldsymbol{u})|^{1+\frac{1}{n}}\, \mathrm{d}\Omega \tag{9a}$$

$$+ \int_{\Gamma_{\mathrm{b}}} \frac{\tau_{\mathrm{ref}}\, u_{\mathrm{ref}}}{1+m} \left(\frac{N}{N_{\mathrm{ref}}}\right)^q \left(\frac{|\boldsymbol{u}|_{\mathrm{b}}}{u_{\mathrm{ref}}}\right)^{1+m} \mathrm{d}\Gamma \tag{9b}$$

$$+ \int_\Omega \rho_{\mathrm{i}} g\, \nabla s \cdot \boldsymbol{u}\, \mathrm{d}\Omega \tag{9c}$$

$$- \int_{\Gamma_{\mathrm{cf}}} [\rho_{\mathrm{i}} g (s - z) - p_{\mathrm{w}}]\, \boldsymbol{u} \cdot \boldsymbol{n}\, \mathrm{d}\Gamma, \tag{9d}$$

where Ω denotes the glacier volume and $|\boldsymbol{u}|_{\mathrm{b}}$ is the norm of the basal velocity, accounting for bed slope. The four terms represent, respectively: (i) the viscous energy dissipated by internal ice deformation, (ii) the frictional energy dissipated by basal sliding, (iii) the gravitational driving energy, i.e. the work done by gravity as ice flows down the surface slope, and (iv) the calving-front energy, which accounts for the work done at the calving front against the ice overburden pressure and the hydrostatic water back-pressure. Note that the minimization formulation (9) was used in the IcePack model, however, referred as action principle (Shapero et al., 2021).

Minimization of $\mathcal{J}$ yields the horizontal velocity field $\boldsymbol{u}$, but the vertical velocity component w is then recovered by vertical integration of the incompressibility condition (4b).

2.4 Enthalpy and Arrhenius factor

In many modeling applications, ice is treated as isothermal with prescribed, possibly spatially varying, values of the Arrhenius factor A and the sliding parametrization τ_{ref} that control the ice dynamics. However, for glaciers in which the thermal regime plays a significant role, such as polythermal glaciers or ice sheets, IGM can optionally compute the ice thermodynamics based on the enthalpy formulation (Aschwanden et al., 2012) largely inspired from its implementation in the Parallel Ice Sheet Model (PISM; Khroulev and the PISM Authors, 2020). This approach unifies the treatment of cold and temperate ice within a single advection–diffusion equation for the enthalpy field E, avoiding the need to explicitly track the cold–temperate transition surface. In cold ice, the enthalpy is proportional to the temperature $T = T(\boldsymbol{x}, t)$; in temperate ice, it encodes the water content $\omega = \omega(\boldsymbol{x}, t)$. The enthalpy equation accounts for heat advection by the ice flow, vertical diffusion, viscous dissipation, and water drainage. Boundary conditions at the surface are set by the climate forcing, while basal conditions depend on the local thermal and hydrological state. The thermal state feeds back into the ice dynamics by modifying the Arrhenius factor A and the sliding parametrization τ_{ref} through the basal melt rate and the subglacial hydrology (next Section). The full mathematical formulation, including the enthalpy equation, boundary conditions, basal melt rate, subglacial hydrology, and the coupling to ice dynamics, is provided in Appendix C.

2.5 Subglacial hydrology

The effective pressure N is a key variable depending on the subglacial hydrology in the basal friction law, as it controls the coupling between the ice and the bedrock. Physically, N is the difference between the ice overburden pressure and the basal water pressure, $N = \rho_{\mathrm{i}} g h - p_{\mathrm{w}}$, so that it vanishes when the bed is fully lubricated and grows when subglacial water is efficiently evacuated. While the modeling of N can be user-defined, IGM provides several built-in simple parametrizations: (i) *Water-connected closure*: the basal water pressure equals the hydrostatic water pressure at the bed, $p_{\mathrm{w}} = \rho_{\mathrm{w}} g \max(z_{\mathrm{wl}} - b, 0)$; this is the relevant closure for marine-terminating glaciers. (ii) *Fixed fraction of overburden*: $p_{\mathrm{w}} = \delta\, \rho_{\mathrm{i}} g h$, with $\delta \in [0, 1]$ prescribing the basal water pressure as a constant fraction of the ice overburden. (iii) *Till-based closure*: N is set to the till effective pressure N_{till} predicted by the subglacial hydrology model of Bueler and van Pelt (2015) described in Appendix C, which couples N to the basal melt rate and the till water storage.

2.6 Lagrangian particles

While IGM components are formulated in an Eulerian framework, the model also includes a Lagrangian utility that tracks

the trajectories of (physical or virtual) particles advected by the ice flow. This capability is particularly useful for studying the evolution of tracers within the ice such as debris, morainic material, cosmogenic nuclides, or for tracking the age and other properties of ice (see, e.g., Licciulli et al., 2020; Scherler and Egholm, 2020; Rowan et al., 2022; Leger et al., 2026). The trajectory $\boldsymbol{p}(t)$ of a particle that passes through $\boldsymbol{p}_0$ at time t_0 satisfies the following ordinary differential equation:

$$\frac{\mathrm{d}\boldsymbol{p}}{\mathrm{d}t}(t) = \boldsymbol{v}(\boldsymbol{p}, t), \tag{10a}$$

$$\boldsymbol{p}(t_0) = \boldsymbol{p}_0. \tag{10b}$$

Each particle can carry associated state variables (e.g., age or chemical composition) that are advected along the trajectory. Particle seeding is user-defined: for instance, subglacially as a function of ice-flow speed to represent glacial erosion, or supraglacially as a function of surrounding slopes to represent debris accumulation (Leger et al., 2026; Hardmeier et al., 2026).

3 Numerics

All fields in IGM are discretized on a regular horizontal grid of $N_x \times N_y$ nodes (x_i, y_j) with uniform spacing H. This design choice is central to the computational efficiency of IGM: computation is parallelized over the horizontal grid, which can be very large, whereas dimensions that cannot be easily parallelized (time stepping, vertical layers, neural network depth) are kept as shallow as possible. Accordingly, the fields h, b, s, and τ_{ref} are represented as tensor fields:

$$\mathbf{h}_H,\ \mathbf{b}_H,\ \mathbf{s}_H,\ \boldsymbol{\tau}_H \in \mathbb{R}^{N_y \times N_x}, \tag{11}$$

where $(\mathbf{h}_H)_{ji} := h(x_i, y_j)$; the other fields are defined similarly. We now describe the numerical treatment of each physical component with a particular focus on the ice flow.

3.1 Ice thickness

To solve equation (2), we employ an explicit upwind finite-volume scheme. Ice transport is computed from edge-defined fluxes, derived from depth-averaged velocities and ice thickness in the upwind direction. To reduce numerical diffusion while preserving monotonicity, a piecewise-linear reconstruction of the ice thickness at cell edges is performed using the Superbee slope limiter (Roe, 1986), which ensures greater accuracy while satisfying the total variation diminishing property. The scheme is mass-conserving and readily parallelizable by virtue of its explicit structure, but is subject to a Courant–Friedrichs–Lewy (CFL) condition that imposes an upper bound on the time step, ensuring that no cell transfers more than its entire content to a neighbor within a single iteration. The time step Δt is therefore constrained through $\|\bar{\boldsymbol{u}}\|_{L^\infty} \Delta t / H \leq C$, where C is the CFL number that must be chosen below 1.

When the glacier terminates in water, the above scheme is complemented with a dedicated treatment of the calving front to implement (3). A first option follows the sub-grid scheme of Albrecht et al. (2011): water cells adjacent to ice accumulate incoming flux into an area-specific volume until a thickness threshold (set by their icy neighbours and capped at flotation) is reached, upon which the cell is promoted to a full ice cell. This scheme is locally mass-conserving and limits numerical diffusion of the calving front. A second option, based on the level-set method of Bondzio et al. (2016), is presented by Prasad et al. (2026).

3.2 Mass balance components

By construction, computing the SMB amounts to evaluating pointwise functions of local fields, as in equations (A1)–(A2), and are therefore intrinsically parallel on the GPU. The sole exception is the Positive Degree Day (PDD) scheme with explicit snowpack tracking (Appendix A), which must be integrated sequentially in time; its computational cost scales linearly with the number of time steps. In practice, using weekly or monthly time steps enables capturing seasonality while avoiding excessively fine temporal resolutions (e.g., daily) that would create a computational bottleneck offsetting the efficiency gains of GPU parallelism elsewhere in the model. Similarly, the calving rate and basal melt reduce to pointwise formulas evaluated on the horizontal raster grid, and therefore pose no specific computational overload.

3.3 Ice flow

Computing ice velocities is often the most expensive step in glacier evolution simulations, as it requires solving the three-dimensional, non-linear Stokes momentum-balance equation (4). This step therefore offers the greatest potential for acceleration.

3.3.1 Vertical discretization

We introduce the terrain-following coordinate ζ, defined by $\zeta := (z - b)/h$ (Fig. 2). Vertical derivatives must then be transformed accordingly in this coordinate system (e.g., Pattyn, 2003; Shapero et al., 2021). For the velocity field, we use the nodal/modal expansions

$$u(x_i, y_j, z) = \sum_{k=1}^{N_z} (\mathbf{u}_H)_{kji}\, \phi_k(\zeta(z)), \tag{12a}$$

and similarly for v, where $\mathbf{u}_H, \mathbf{v}_H \in \mathbb{R}^{N_z \times N_y \times N_x}$ denote the corresponding degrees of freedom. Several families of basis functions $\{\phi_k\}_{k=1}^{N_z}$ may be employed, including default Lagrange basis functions, for which the degrees of freedom coincide with nodal velocity values, as well as polynomial bases such as Legendre polynomials (Shapero et al., 2021)

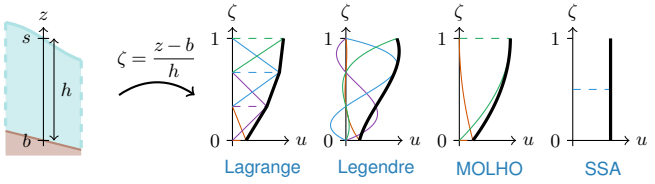

Figure 2. Vertical discretization schematic. The terrain-following coordinate $\zeta = (z-b)/h$ maps the ice column to $[0,1]$. Four vertical basis types are shown: Lagrange (piecewise linear polynomial interpolation), Legendre (polynomial expansion; Shapero et al., 2021), MOLHO (SIA profile; Dias dos Santos et al., 2022), and SSA (SSA profile; depth-averaged).

the MOLHO polynomials (Dias dos Santos et al., 2022), and the depth-averaged SSA basis ($N_z = 1$); see Fig. 2. For Lagrange discretization, the N_z vertical nodes can be distributed in a non-uniform manner to ensure finer discretization near the ice-bedrock interface, where the steepest gradients are expected, and coarser near the ice-surface interface following the strategy proposed in PISM (Khroulev and the PISM Authors, 2020).

3.3.2 Setting the minimization problem

The discretized version of the minimization of the functional $\mathcal{J}$ described in (9) reads as follows:

$$(\mathbf{u}_H^*, \mathbf{v}_H^*) = \arg\min \{\mathcal{J}_H(\mathbf{u}_H, \mathbf{v}_H), (\mathbf{u}_H, \mathbf{v}_H) \in \mathcal{X}_H\} \quad (13)$$

in which $\mathcal{J}_H$ is the numerical approximation of $\mathcal{J}$ and $(\mathbf{u}_H^*, \mathbf{v}_H^*)$ denotes the numerical solution to this optimization problem, and $\mathcal{X}_H$ is the discrete approximation space $\mathcal{X}_H = \mathbb{R}^{2\times N_z \times N_y \times N_x}$. In the spirit of the Ritz–Galerkin method, the key idea is to replace the full (infinite-dimensional) function space $\mathcal{X}$ by a finite-dimensional space $\mathcal{X}_H$, and to seek the minimizer within that subspace. In IGM, the horizontal integrals are evaluated on each grid cell using one of several available quadrature rules: (i) a single cell-centred evaluation point (default); (ii) a 2×2 Gaussian quadrature on bilinear (Q1) elements; (iii) a P1 triangulation that splits each quadrilateral cell into two triangles with centroid evaluation; and (iv) a marker-and-cell (MAC) staggered-grid scheme that evaluates gradients at edge midpoints. The first is the least expensive but may admit checkerboard zero-energy modes, while the latter three eliminate such spurious modes, at the cost of increased computational overhead. The vertical integrals are computed with quadrature points and weights consistent with the chosen vertical basis (Lagrange, Legendre, MOLHO, or SSA).

In IGM, the ice flow is parametrized as follows:

$$(\mathbf{u}_H, \mathbf{v}_H) = \mathcal{M}(\boldsymbol{\theta}, \mathbf{I}_H), \quad (14)$$

in which $\mathcal{M}$ is a differentiable mapping, $\boldsymbol{\theta}$ denotes trainable weights, and $\mathbf{I}_H$ denotes the set of inputs corresponding to the parameters describing ice-flow parameters and glacier geometry, namely $\boldsymbol{\tau}_H$, $\mathbf{A}_H$, $\mathbf{h}_H$, and $\mathbf{s}_H$ (Fig. 3). For a given input $\mathbf{I}_H$, the mapping $\mathcal{M}$ defines an approximation space spanned by weights $\boldsymbol{\theta}$:

$$\mathcal{X}_{\mathcal{M}} = \{\mathcal{M}(\boldsymbol{\theta}, \mathbf{I}_H) \mid \boldsymbol{\theta}\}. \quad (15)$$

The ice-flow problem is then solved by minimizing the energy functional over this approximation space $\mathcal{X}_{\mathcal{M}}$ rather than over the full discrete velocity space $\mathcal{X}_H$ as in (13). Specifically, the optimization problem (13) becomes

$$\boldsymbol{\theta}^* = \arg\min \{\mathcal{J}_H(\mathcal{M}(\boldsymbol{\theta}, \mathbf{I}_H)), \boldsymbol{\theta}\}. \quad (16)$$

and the solution to the ice-flow problem is then given by:

$$(\mathbf{u}_H, \mathbf{v}_H) = \mathcal{M}(\boldsymbol{\theta}^*, \mathbf{I}_H). \quad (17)$$

The framework defined by equations (14)–(17) provides flexibility:

- If $\mathcal{M}$ is the identity mapping with respect to the weights $\boldsymbol{\theta}$, so that the weights coincide with the velocity degrees of freedom themselves, then we recover optimization problem (13). This direct approach is computationally expensive and is used to manufacture reference solutions for verification and testing.
- The mapping $\mathcal{M}$ may also take the form of a neural network parametrized by inputs $\mathbf{I}_H$ (Fig. 3). In this case, the weights $\boldsymbol{\theta}$ over which the cost $\mathcal{J}_H \circ \mathcal{M}$ is minimized correspond to those of the neural network, and we recover the approach described by Jouvet and Cordonnier (2023) and Cordonnier et al. (2023). This approach is known in the literature as the Deep Ritz Method (E and Yu, 2018) because it minimizes the energy associated with the variational problem, in contrast to more classical physics-informed neural-network approaches, which minimize directly the norm of the residuals of the equations. While Jouvet et al. (2021) originally proposed $\mathcal{M}$ as a convolutional neural network, the framework (14)–(17) allows for other architectures, such as Fourier neural operator networks (Li et al., 2021).
- The mapping $\mathcal{M}$ could also take a hybrid form, e.g., polynomial expressions of the input fields inspired by the analytical structure of simplified models such as the SIA, with trainable coefficients optimized within (16). While this approach has not yet been investigated, the framework readily accommodates it.

3.3.3 Solving the minimization problem

To solve problem (16), IGM supports first-order methods such as Adam (Kingma and Ba, 2015), as well as second-order methods such as L-BFGS (Liu and Nocedal, 1989) and SOAP (Vyas et al., 2024). The latter are increasingly

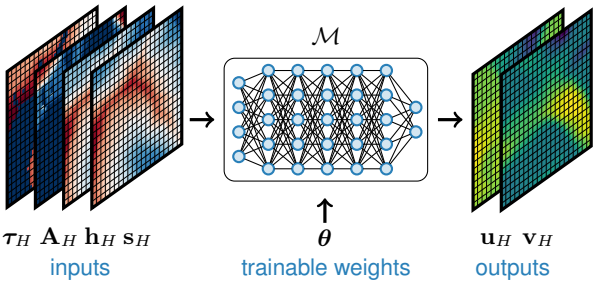

Figure 3. Schematic of the ice-flow computation framework. The degrees of freedom for the horizontal velocity are parametrized as $(\mathbf{u}_H, \mathbf{v}_H) = \mathcal{M}(\boldsymbol{\theta})$, where $\mathcal{M}$ denotes a differentiable mapping and $\boldsymbol{\theta}$ represents the trainable weights. This framework encompasses both the 'classical' approach to solving ice flow, taking $\mathcal{M}$ as the identity mapping $\mathcal{I}$ and $\boldsymbol{\theta} = (\mathbf{u}_H, \mathbf{v}_H)$ so that the trainable weights are given directly by the velocity degrees of freedom, and a neural approach, in which $\mathcal{M}$ is a neural network $\mathcal{N}$ mapping glacier inputs $(\boldsymbol{\tau}_H, \mathbf{A}_H, \mathbf{h}_H, \mathbf{s}_H)$ to $(\mathbf{u}_H, \mathbf{v}_H)$. In the latter case, the weights $\boldsymbol{\theta}$ correspond to the neural-network weights.

favoured in scientific machine learning because they account for the geometry of the loss landscape during training (e.g., Müller and Zeinhofer, 2023, 2024; Jnini et al., 2026). The optimization problem (16) can be solved in two regimes.

- In the *offline* regime, the mapping $\mathcal{M}$ is trained once on a large glacier catalogue in a data-driven fashion (Jouvet et al., 2021), yielding a generic network that is then applied to any geometry and set of ice-flow parameters without further retraining. Generalization typically requires a large network, and accuracy is ultimately bounded by its representational capacity.

- In the *online* regime, $\mathcal{M}$ is retrained continuously in a physics-informed fashion alongside the time-stepping of the glacier evolution model (Jouvet and Cordonnier, 2023; Cordonnier et al., 2023). Retraining can be triggered periodically or adaptively in order to mitigate training costs while maximizing accuracy. Because the network is specialized to a single simulation, a lightweight architecture is usually sufficient, and accuracy can be improved by increasing the retraining effort.

Training (problem (16)) is the computational bottleneck of the framework, with GPU memory being the most critical constraint: automatic differentiation requires storing the activations of the entire forward pass to propagate gradients backwards, so memory scales with both network depth and the size of each training input. In contrast, inference (problem (17)) through $\mathcal{M}$ is a plain forward pass and has a modest memory footprint. This asymmetry between expensive training and inexpensive inference also frames the trade-off between the two regimes: the *offline* regime amortizes a large one-time training cost over many simulations at the price of a sizeable, generic network, whereas the *online* regime keeps the network lightweight and problem-specific but pays a training cost at each retraining step.

3.4 Enthalpy

When active, the thermodynamics module takes as input the ice velocity field computed in Section 3.3 together with boundary conditions (surface temperature and geothermal heat flux), and returns the three-dimensional enthalpy field, from which temperature and water content are deduced, as well as the basal melt rate, thereby coupling the thermal state back to the ice dynamics. Internally, the enthalpy equation (C4) is discretized on a vertically stretched Lagrange. Time integration follows an operator-splitting approach: horizontal advection is treated explicitly with upwind finite differences, while vertical advection–diffusion is solved implicitly. Because the implicit step contains no horizontal coupling, it reduces to $N_x \times N_y$ independent tridiagonal systems of size $\tilde{N}_z$ (one per ice column), each solved in parallel using the Thomas algorithm. Note that the vertical discretization for the enthalpy may differ from the iceflow one, justifying the introduction of $\tilde{N}_z$. The computational cost of this step therefore scales as $\mathcal{O}(N_x \times N_y \times \tilde{N}_z)$ and is fully parallelizable over the horizontal grid. Full details of the discretization are given in Appendix C.

3.5 Lagrangian particles

The particle-trajectory equation (10) is integrated in time using an explicit forward Euler scheme. Because particle positions generally do not coincide with grid nodes, the velocity field must be interpolated to the particle locations. Bilinear interpolation is performed by a dedicated CUDA kernel called from Python, which enables the simulation of large numbers of particles at modest computational and memory cost. As a demonstration, the trajectories of 50 million particles were computed in a high-resolution simulation of the glaciers in the European Alps during the last glacial cycle, with a 25% overhead compared to the full simulation (see Fig. 4 in Leger et al., 2026). The particle-tracking module also provides mapping utilities between the Lagrangian and Eulerian representations; particle-borne quantities can be scattered back onto the grid through bilinear weighting to produce Eulerian fields, as used for instance by Hardmeier et al. (2026) to reconstruct a supraglacial debris-thickness field from the particle ensemble.

4 Data assimilation

Beyond solving the forward ice-flow problem, formulating the velocity as a differentiable mapping $\mathcal{M}$ offers another key advantage for model inversion: since $\mathcal{M}$ is differentiable with respect to its inputs $\mathbf{I}_H$ as well as its weights $\boldsymbol{\theta}$, one can back-propagate through $\mathcal{M}$ not only to assimilate the physics by training $\boldsymbol{\theta}$, but also to assimilate observational data by

optimizing the input fields $\mathbf{I}_H$ (e.g., bedrock topography or ice-flow parameters) so as to minimize the misfit to measurements (Jouvet, 2022), yielding a unified physics- and data-assimilation framework (Jouvet and Cordonnier, 2023). Yet, simultaneously minimizing physics and data misfit poses significant challenges, as the joint optimization landscape can be complex with conflicting gradient directions — a problem that remains under active investigation. An alternative, more tractable approach consists of using an offline-trained ice-flow mapping $\mathcal{M}$ with fixed weights $\boldsymbol{\theta}$, which reduces the inversion to a data-only optimization, substantially simplifying the problem and making it more suitable, e.g., for ice-thickness inversion to match surface ice velocity data.

In addition, IGM integrates a complementary, non-gradient-based inversion framework that uses IGM as a forward model and, within a single forward time loop, nudges one or more control fields—surface mass balance, basal sliding coefficient, ice thickness or bedrock topography—so as to drive the discrepancy between modeled and observed quantities (ice-thickness change rate, surface elevation or surface velocity) toward zero. This generalizes the ice-thickness-change-rate bedrock inversion of Frank and van Pelt (2024) and recovers, as special cases, the broader family of nudging/time-relaxation schemes such as PISM-style force-to-thickness (Khroulev and the PISM Authors, 2020) and basal-sliding inversion from surface misfit (Pollard and DeConto, 2012).

5 Software design

The software design of IGM is built around four pillars: (i) *accessibility*, (ii) *modularity*, (iii) *reproducibility*, (iv) *computational scalability*, and (v) *parameter explorability*. We describe below how each is realized in practice.

5.1 Accessibility

5.1.1 Installation, dependencies, and hardware

A common challenge among scientific computing software is accessibility, starting with installation. To facilitate this, we offer several ways a user can install IGM. First, for researchers who want a stable version, they can install IGM using PyPi servers for the latest tagged release. Secondly, for users who wish to contribute to the development and have access to the latest changes, they can instead install IGM directly from the source repository (Section 5.3).

Importantly, both methods of installing IGM automatically resolve TensorFlow's CUDA requirements, resulting in an effortless GPU installation which only requires a sufficiently up-to-date Nvidia driver. In the case the user does not have access to an Nvidia GPU, IGM will otherwise install the corresponding CPU version.

IGM is built around various Python packages including TensorFlow (Abadi et al., 2015) for GPU-accelerated computations and automatic differentiation, Hydra (Yadan, 2019) for configuration management, Optuna (Ozaki et al., 2025) for hyperparameter optimization, Xarray for handling I/O, and OGGM (Maussion et al., 2019) for glacier data retrieval and preprocessing. As such, we recommend users to have a dedicated virtual environment to contain all dependencies in a consistent manner.

IGM is hosted on GitHub (https://github.com/instructed-glacier-model/igm) and developed under an open-source policy, encouraging community contributions through pull requests and discussion.

5.1.2 Interface

The main interface to IGM is a YAML configuration file that describes the simulation to be performed. This file specifies which input, process, and output components are used, the order in which they are called, and any parameter values that should differ from the model defaults. Separating the experiment configuration from the source code allows standard simulations, parameter perturbations, and component substitutions to be carried out without editing the numerical implementation itself.

A simulation is launched from the command line as

```
igm_run +experiment=[YAML_FILENAME]
```

where `[YAML_FILENAME]` points to the configuration file for the experiment. After this command is executed, Hydra resolves the user-defined configuration together with the default IGM settings. IGM then initializes the selected components, enters the time loop, updates the model state by calling each process component (Fig. 4) in the prescribed order, writes the requested outputs, and finalizes the run. A more detailed description of this execution sequence is given in Appendix F.

This configuration-driven interface is intended to make the coupling between model components explicit. For example, adding a thermo-mechanical coupling, a specific diagnostic output, or modifying a forcing parameter is done by changing the configuration file rather than by editing the main simulation loop. At the same time, the resolved configuration is written to the run directory together with the model output, so that the exact setup of each experiment is retained for later inspection or reproduction (Section 5.3).

An example configuration is shown below. It selects local input and output components, combines several process components, and overrides a small number of default parameter values. The example is intended to illustrate the structure of the interface rather than define a fixed template, since the available components and parameter names may evolve between IGM versions.

parameters.yaml

```yaml
# Required: Specify IGM components
defaults:
  - override /inputs:
      - local
  - override /processes:
      - smb
      - iceflow
      - time
      - thk
  - override /outputs:
      - local

# Optional: Overriding default values
input:
  local:
    filename: greenland.nc
processes:
  time:
    start: 2000.0
    end: 2100.0
    save: 10.0
  smb:
    method: simple
```

5.2 Modularity

5.2.1 Exchangability

For a researcher using scientific computing software, being able to freely explore physical parameters and create experiments is vital. However, many scientific codebases contain highly coupled code, acting as a limitation as one change in the code might inadvertently affect another. This then ultimately restricts the researcher's freedom to explore.

One way to reduce this coupling is to organize the model around self-contained components with clearly defined responsibilities. This is the approach taken, for example, in Landlab (Barnhart et al., 2020; Hutton et al., 2020; Hobley et al., 2017), where model behavior is assembled from independent components that can be combined, reordered, or replaced depending on the application. Such a structure is particularly useful for glacier modeling, where different experiments may require different choices of surface mass balance, ice dynamics, hydrology, forcing, or output diagnostics, without requiring changes to the rest of the code base.

Following this convention, our decoupled components are called "modules", which are organized into four distinct categories: "*inputs*", "*assimilations*", "*processes*", and "*outputs*" (Fig. 4). Input and output modules handle I/O, e.g., utilizing Xarray to load georeferenced rasters or OGGM to retrieve glacier data. Assimilation modules are an optional intermediate step following "Input", including data assimilation methods such as "*field_inversion*" or "*time_relaxation*", which aim to optimize spatial fields (e.g. ice thickness) to match observations (Section 4). Processes, encapsulating a broader definition, represent a physical state change through time. This can be the time itself, the ice thickness, the subglacial hydrology, or the precipitation at any given time through the simulation. As this is dynamically updated at each time step, it allows for low memory usage and efficient exporting through one of the I/O modules. The main process modules shipped with IGM are summarized in Appendix G.

This design decouples the overall complexity of the model: each module encapsulates one physical or technical aspect, with its own parameters, source code, and documentation, and can therefore be developed, tested, and documented independently of the rest. This modularity also facilitates community development, as contributors can add or extend modules (see Section 5.2.2) without needing to understand the entire codebase.

5.2.2 Extensibility

A modular model structure is only useful if individual components can also be replaced or extended in practice. In glacier modeling, this is often necessary because different applications require different process representations. For example, a regional-scale experiment may use a simple temperature-index surface mass-balance scheme, whereas a more targeted study may require additional feedbacks, alternative forcing corrections, or new diagnostic quantities. Directly modifying the core source code for each such experiment would make the model harder to maintain and would also make it more difficult to identify which changes were responsible for a given result.

IGM therefore allows users to add custom components outside the core package. A user-defined component follows the same structure as the default components and can be selected in the YAML configuration file in the same way as any built-in input, process, or output component. For instance, a user who wishes to test an alternative surface mass-balance formulation can copy the default component, modify only the relevant update routine, place the new version in the `user/` directory, and reference it in the experiment configuration. During model execution, this component is then initialized, updated, and finalized through the same mechanism as the standard components.

This approach keeps experiment-specific developments separated from the maintained source code, while still allowing them to use the standard IGM infrastructure for configuration (e.g., time stepping or input/output). It also makes such developments easier to version-control and archive together with the simulations that depend on them. The layout of the IGM source package is given in Appendix F, the project folder layout in Appendix E, and a concrete user-component example in Appendix H. This design supports individual model experimentation, while also providing a clear path for useful extensions to be contributed back through the open-source development workflow described in Section 5.3.

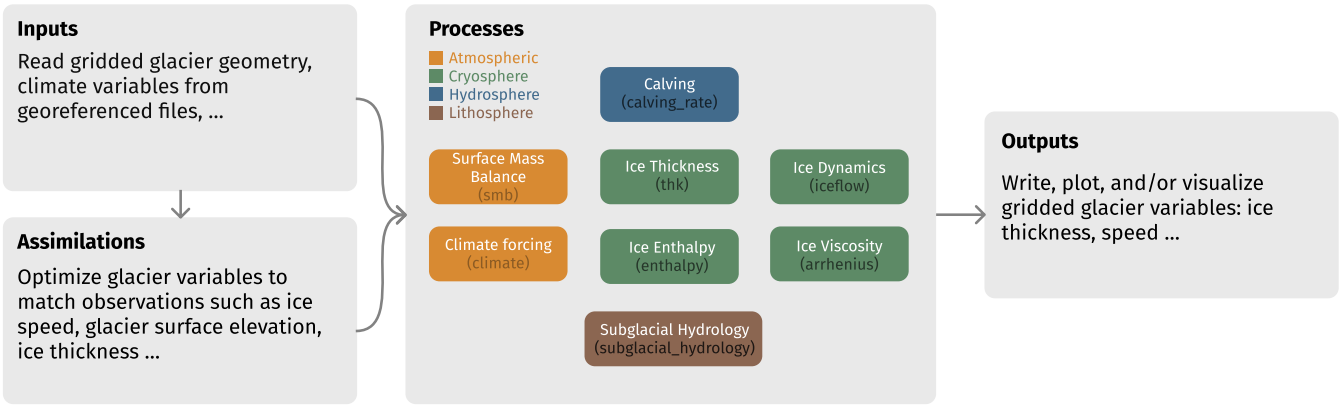

Figure 4. Overview of the IGM modular architecture: Gridded glacier geometry and climate variables are read from georeferenced files in "input" modules and optionally adjusted in "assimilation" modules to match observations (e.g. ice surface speed, surface elevation, thickness). A configurable chain of "process" modules — color-coded by earth-system component — is then evaluated; each box gives the process name and its module identifier (e.g. `smb`, `thk`, `iceflow`). Results are written, plotted, or visualized as gridded glacier variables in "output" modules.

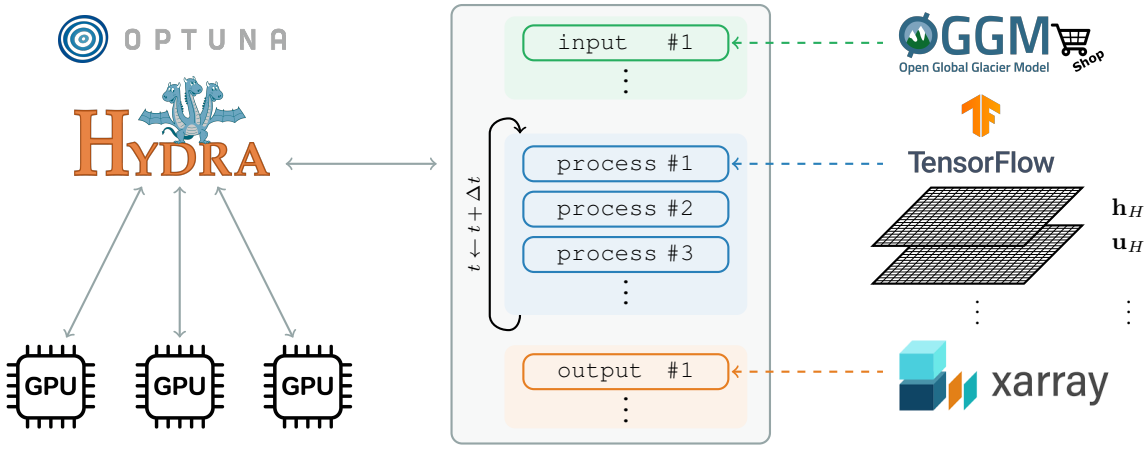

Figure 5. Schematic of the IGM software architecture. Simulations are assembled from input, process, and output components selected in a YAML configuration file. Input components populate the model state from local files or external glacier data sources; process components update state variables on a raster grid during the time loop; and output components write fields, diagnostics, or visualizations. Internally, spatial fields are represented as TensorFlow tensors, while file-based input and output are handled through Xarray.

5.3 Reproducibility

IGM records the information needed to repeat a simulation as part of each run directory. The resolved Hydra configuration, including command-line overrides, is written to a *.hydra* folder. In addition, IGM records the commit hash used for the simulation. This provides a direct link between a model result, the parameter set used to generate it, and the source code version that was executed.

Furthermore, IGM follows a continuous integration and continuous deployment (CI/CD) workflow: GitHub Actions automatically rebuilds and publishes the documentation (https://igm-model.org/) each time a commit is pushed, whether it corresponds to a tagged release or not. This ensures that users can always access documentation corresponding to their installed version of IGM. Together, resolved configurations, version tracking, and continuously updated documentation support reproducible simulation workflows and make parameter studies easier to inspect, rerun, and archive.

5.4 Computational scalability

In glacier modelling, large spatial and gridded domains are commonplace. However, computations on such domains can quickly become limiting on classical hardware, especially when the model state consists of several large raster fields that are updated repeatedly through time. IGM addresses this by representing spatial fields internally as TensorFlow tensors (Abadi et al., 2015), allowing array operations on struc-

tured grids to make use of GPU-optimized kernels where possible. When needed, IGM can also use CPU-based execution, for example for algorithms or environments where GPU execution is not appropriate.

By default, a single IGM simulation currently runs on one GPU. Domain decomposition across multiple GPUs for a single simulation is not yet supported. However, independent simulations can be assigned to different devices, and multiple small-domain simulations can share a GPU when memory permits. This latter capability is useful for ensemble and optimization workflows, but it should be interpreted as task-level parallelism rather than multi-GPU scaling of a single numerical simulation.

As GPU hardware and libraries continue to evolve, IGM can benefit from improvements in accelerator performance and memory capacity, enabling increasingly efficient simulations over large spatial domains.

5.5 Parameter explorability

Beyond individual simulations, glacier modeling often requires systematic exploration of uncertain parameters, forcing choices, or process representations. Such workflows are needed for sensitivity analysis, uncertainty quantification, and calibration against observations. IGM supports these workflows through two complementary mechanisms: ensemble simulations using Hydra library and optimization-based calibration using Optuna library.

5.5.1 Ensemble simulations

IGM leverages Hydra's *multirun* feature to run ensembles of simulations over user-defined parameter combinations. This allows users to vary selected parameters from the command line while keeping the rest of the configuration fixed and recorded. For example, the following command runs a 3×3 ensemble in which only the precipitation and temperature factors vary from the original parameter file:

```
igm_run +experiment=params \
  processes.precip.factor=100,200,300 \
  processes.temp.factor=0.5,1.0,1.5 \
  hydra/launcher=joblib --multirun
```

Each member of the ensemble is executed as a separate simulation with its own resolved configuration and output directory. This makes ensemble results easier to inspect and reproduce, since the parameter values used for each run are stored together with the corresponding outputs. Depending on the available hardware and launcher configuration, ensemble members can be run sequentially or distributed across available CPU or GPU resources.

5.5.2 Optimization-based calibration

For calibration problems in which the aim is to minimize an objective function rather than enumerate a fixed parameter grid, IGM provides an interface to Optuna (Ozaki et al., 2025). In this workflow, the user defines the parameters to be sampled in a dedicated parameter file and specifies an objective function, for example through a user-defined function as described in Section 5.2.2. After each simulation, the objective function evaluates the model result, and Optuna proposes new parameter combinations based on the history of previous trials.

This workflow allows users to perform single- or multi-objective calibration experiments while relying on IGM for model execution, configuration handling, and output generation. Optuna studies can also be stored and resumed, allowing failed or interrupted calibration experiments to continue without discarding completed simulations. Progress and results can be monitored interactively through the Optuna dashboard, which provides real-time visualization of trial histories, parameter importances, and Pareto fronts for multi-objective problems. A practical demonstration is presented in Section 6.

6 Applications

In this section, we demonstrate the capabilities of IGM through real-world applications. We first show how IGM efficiency, Hydra and Optuna can be used to carry large ensembles of simulations, and perform a calibration of Aletsch Glacier. Note that IGM also ships with a test suite that validates each part of the code against reference solutions; results are shown in Appendix D.

6.1 Ensemble calibration simulations of Aletsch Glacier from 1880 to 2017

Building on an early modeling of the Great Aletsch Glacier (Jouvet et al., 2011), we showcase IGM's multi-objective calibration capability in a 200-trial Optuna sweep over the 1880–2017 transient at 100 m of resolution – an experiment that can be launched in a single command line with IGM. The surface mass balance and ice-flow parameters are tuned simultaneously against two observational constraints: (i) seven historical surface DEMs spanning 1880–2017, and (ii) the present-day satellite-derived surface velocity field. The forward model is forced by a temperature-index mass-balance scheme driven by daily meteorological data. Three parameters among the most uncertain ones are optimized: the accumulation and ablation scaling factors $w_{\mathrm{acc}}, w_{\mathrm{abl}} \in [0.5, 2.0]$, and the reference basal shear stress $\tau_{\mathrm{ref}} \in [0.08, 0.5]\,\mathrm{MPa}$ of a Weertman sliding law. Two scalar misfits are evaluated: the RMSE of the seven modeled minus observed DEMs, and the RMSE of modeled vs. observed surface speeds in 2017. The two objectives are minimized jointly with the NSGA-II optimizer. The 200 trials are distributed over ten parallel jobs on a single NVIDIA RTX 4090 GPU, made possible by the rela-

tively modest grid size ($\sim 200 \times 150$ cells), and the sweep converges in roughly 1 h of wall-clock time ($\sim$2.6 trials/min) to a tight 13-point Pareto front (Fig. 6a). Across the entire front, the two melt multipliers lock onto a single value ($w_{\mathrm{abl}} \approx 1.33$, $w_{\mathrm{acc}} \approx 1.17$), whereas τ_{ref} spans the whole search range from 0.13 to 0.41 MPa: the available surface observations alone cannot discriminate a slippery, plug-flow regime from a sticky, deformation-dominated one. We therefore retain the knee-point configuration ($\tau_{\mathrm{ref}} = 0.3\,\mathrm{MPa}, w_{\mathrm{acc}} = 1.20, w_{\mathrm{abl}} = 1.33$) sitting in the middle of the τ_{ref} range; the corresponding centerline sliding fraction $u_{\mathrm{b}}/u_{\mathrm{s}} \approx 0.3$–$0.5$ is consistent with observational expectations on Aletsch Glacier. The calibrated configuration recovers RMSE for DEMs and surface speeds of $\approx 30\,\mathrm{m}$ and $\approx 20\,\mathrm{m\,yr^{-1}}$ (Fig. 6c), a good simultaneous fit of the 140-year geometric evolution and of the present-day largest glacier of the Alps, at a relatively modest computational cost.

7 Community building

The long-term sustainability of IGM depends on an active user and developer community that can maintain, extend, and critically evaluate the model. To this end, several channels for collaboration and engagement have been established:

- The code is openly hosted on GitHub (https://github.com/instructed-glacier-model/igm) and welcomes contributions via pull requests, with continuous integration ensuring that proposed changes pass the existing test suite before merging (Sect. 5.3).
- A public website (https://igm-model.org/) provides user-oriented documentation, including parameter descriptions and configuration guides, alongside more detailed material such as examples and tutorials.
- A dedicated Discord server offers a low-barrier space for discussion and troubleshooting among users and developers.

These resources aim to lower the entry barrier for new users and coordinate development. Community interactions have also catalyzed scientific exchange, including regular IGM meetings, organized online, at EGU, or a joint OGGM–IGM workshop in Switzerland (https://oggm.org/2024/01/19/8th-workshop-announcement/), in 2024, which brought together researchers from complementary glacier-modeling communities.

8 Education, outreach, and 3D visualization

Beyond research, IGM has been used in teaching and outreach. At the Université de Lausanne, it supports a Master-level glaciology course in which students explore glacier dynamics through hands-on numerical experiments (https://jouvetg.github.io/glamod). Thanks to the OGGM Shop (Maussion et al., 2019), which provides pre-processed glacier data for every glacier in the world, each student is assigned a different glacier and carries out a complete modeling exercise, ranging from a Little Ice Age reconstruction to future projections under various climate scenarios. IGM has also been used to run simulations that communicate how glaciers may change under climate warming to broader audiences, for example within the IceAgeCam (https://swiss-time-machine.org/) or FROST projects (https://www.glacier-evolution.nat.fau.de/). The latter makes use of a browser-based 3D visualization tool (Plotly/Dash) bundled with IGM, which converts model output into an interactive web application with a time slider and user-selectable fields displayed as a textured bedrock overlain by the glacier surface.

9 Conclusions, ongoing and future work

We have presented the fundamentals of IGM, an accessible, open-source, and GPU-accelerated glacier evolution model. IGM's physics-informed deep-learning approach to high-order ice-flow computation on GPU delivers the computational efficiency required to simulate large glacier ensembles and high-resolution domains over long time scales. Its modular software architecture keeps the model accessible, with a simple YAML-based configuration built on the Hydra framework and custom modules added via relatively simple Python coding. The model has been evaluated against established benchmarks and demonstrated on real-world applications ranging from individual glaciers to large icefields. An active community, open development workflow, and comprehensive documentation support the long-term sustainability of the project. Several directions for future development are on-going. On the methodological side:

- **Ice-flow solving.** Several questions remain open: the trade-off between small, glacier-specific online-trained and large, generalizable, offline-trained neural network surrogates; the systematic assessment of resulting errors; the design of a posteriori error estimators that give users a proxy for the error without requiring a reference solution; the optimal choice of the mapping $\mathcal{M}$ (neural architecture or otherwise) for a given problem; efficient strategies that maximize accuracy while reducing training effort; and parallelization of training across multiple GPUs.
- **Data assimilation and inversion.** IGM already provides modules for inverting selected parameters, such as basal sliding coefficients (Jouvet, 2022), as well as for training the ice-flow neural network simultaneously with the inversion of these parameters (Jouvet and Cordonnier, 2023). However, jointly assimilating observations and the ice-flow physics within a single differ-

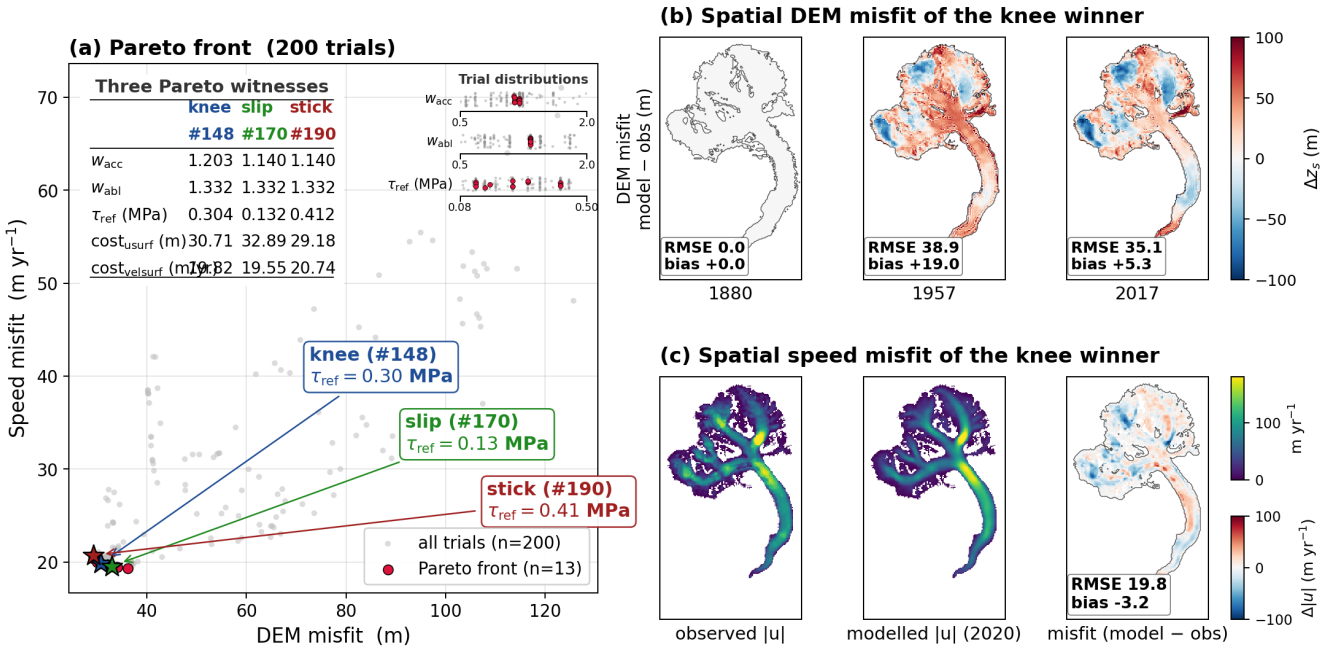

Figure 6. (a) Pareto front of the 200-trial NSGA-II sweep (grey: all trials, red: 13 non-dominated) for the simulation of Aletsch Glacier from 1880 to 2017. The three witnesses (`knee`, `slip`, `stick`) span the τ_{ref} axis of the front; the inset table gives their parameters and scores, and the small top-right inset shows the per-parameter trial distribution. (b) Modeled − observed surface elevation at three DEM dates for the knee winner. (c) Surface speed at 2020 observed, modeled, and (model − observed) for the same winner.

entiable end-to-end framework still lacks robustness in practice, and making such unified inversions reliable is the subject of ongoing work.

- **Implicit time integration.** The current explicit time-stepping scheme limits the maximum allowable time step, which becomes a critical bottleneck for very long simulations such as palaeo-glaciation over glacial cycles or coupled landscape-evolution runs spanning millions of years (Cordonnier et al., 2023). Implicit or semi-implicit schemes could relax this constraint, but their efficient implementation on GPUs is hampered by memory limitations and remains an active area of research (e.g. Aissa et al., 2017; Gaudreault et al., 2018; Kashi and Nadarajah, 2025), including recent work applied to ice flow (Ahlkrona et al., 2026).

On the application side, several promising directions are emerging:

- **Global-scale modeling and uncertainty quantification.** IGM has the efficiency to be deployed at regional to global scales. In this context, its tight integration with the Hydra/Optuna workflow opens the door to intensive ensemble-based parameter calibration and uncertainty quantification directly against remote-sensing observations.
- **Paleo-glaciation and landscape evolution.** IGM's scalability to long paleo-glacier simulations has already been demonstrated through Last Glacial Maximum reconstructions across the Alps (Leger et al., 2025). Building on this foundation, ongoing work focuses on coupling IGM with glacial erosion and sediment-transport models to address landscape evolution over geological time (> 1 My) scales.
- **Marine ice sheets.** Now equipped with a calving-front term, IGM holds strong potential for simulating marine-terminating glaciers and ice sheets at the scale and resolution required for continental Antarctic or Greenland applications. However, challenges remain in capturing the complex grounding-line dynamics, which require high resolution, together with the design of appropriate calving laws and sub-shelf melt parametrizations.
- **Lagrangian particle tracking.** GPU acceleration allows IGM to advect millions of Lagrangian particles at essentially little additional cost, opening the door to a new class of large-ensemble applications that would be intractable on CPU. Such massive particle populations can be leveraged to trace sediment provenance (Leger et al., 2026), reconstruct englacial layer stratigraphy, track debris (Hardmeier et al., 2026), and more generally to couple glacier dynamics with isotopic or geochemical models.

Code and data availability. IGM is available as an open-source software at https://github.com/instructed-glacier-model/igm. The associated documentation website can be found at https://igm-model.org/.

Appendix A: Surface mass balance schemes

The simplest parametrization assumes a piecewise-linear dependence of the SMB on altitude z, as supported by observations (Meier et al., 1971), controlled by the equilibrium-line altitude z_{ELA}:

$$\dot{a}_{\mathrm{s}}(z) = \begin{cases} \beta_{\mathrm{abl}}(z - z_{\mathrm{ELA}}), & z \le z_{\mathrm{ELA}}, \quad \text{(A1a)} \\ \min\left(\beta_{\mathrm{acc}}(z - z_{\mathrm{ELA}}), \dot{a}_{\max}\right), & z > z_{\mathrm{ELA}}, \quad \text{(A1b)} \end{cases}$$

where $\beta_{\mathrm{abl}} > 0$ and $\beta_{\mathrm{acc}} > 0$ are the ablation and accumulation gradients, and $\dot{a}_{\mathrm{s,max}}$ is a cap on the maximum accumulation rate. The four parameters z_{ELA}, β_{abl}, β_{acc}, and $\dot{a}_{\max}$ may vary in time and are linearly interpolated from a user-supplied time series.

IGM also provides two climate-driven SMB modules that compute the mass balance from temperature and precipitation fields.

The first scheme implements the monthly temperature-index model of OGGM (Maussion et al., 2019), whose parameters are calibrated against geodetic mass balance observations (Hugonnet et al., 2021). The annual SMB is computed as

$$\dot{a}_{\mathrm{s}} = \frac{\rho_{\mathrm{w}}}{\rho_{\mathrm{i}}} \sum_{m=1}^{12} \left[P_m^{\mathrm{sol}} - d_{\mathrm{f}} \max(T_m - T_{\mathrm{melt}}, 0) \right], \quad \text{(A2)}$$

where T_m and P_m^{sol} are the monthly air temperature and solid precipitation, T_{melt} is the melting-point threshold, d_{f} is the melt factor, and $\rho_{\mathrm{w}}/\rho_{\mathrm{i}}$ converts water to ice equivalent. Solid precipitation is derived from total precipitation assuming full solid fraction below a lower temperature threshold and zero above an upper threshold, with a linear transition in between.

The second scheme extends the model in (A2) with a more elaborate positive degree-day (PDD) approach (Hock, 2003). It explicitly tracks the snowpack depth and applies distinct melt factors for snow and ice. The PDD integral is evaluated using the semi-analytical expectation-based formulation of Calov and Greve (2005), which accounts for sub-monthly temperature variability through a prescribed standard deviation. Additionally, a fixed fraction of the computed melt is assumed to refreeze within the snowpack.

Appendix B: Derivation of the Blatter–Pattyn energy

Given the importance of the Blatter–Pattyn energy in our ice-flow solver, we describe here how it can be derived from the original Stokes problem (4)–(8). We closely follow Jouvet (2016) and refer interested readers to that work for additional technical details. The derivation is carried out in two steps: first, we derive a minimization formulation for the Stokes problem, and then we eliminate the vertical velocity component in order to obtain the Blatter–Pattyn energy.

B1 Minimization formulation of the Stokes problem

To derive a minimization formulation, we first recast the Stokes problem in a variational formulation. We therefore introduce the following space of admissible velocities:

$$\mathcal{X} := \{ \boldsymbol{v} \in [W^{1,1+1/n}(\Omega)]^3, \nabla \cdot \boldsymbol{v} = 0, \ \boldsymbol{v} \cdot \boldsymbol{n}|_{\Gamma_{\mathrm{b}}} = 0 \}, \quad \text{(B1)}$$

where the Sobolev space $W^{1,1+1/n}$ guarantees that the problem is well-defined. Note that we assume $0 < m \le 1/n$.

Multiplying the momentum-balance equation (4a) by a test function $\boldsymbol{\varphi} \in \mathcal{X}$ yields

$$\int_\Omega [\nabla \cdot \boldsymbol{\sigma}(\boldsymbol{v}, p)] \cdot \boldsymbol{\varphi} \, \mathrm{d}\Omega + \int_\Omega \rho_{\mathrm{i}} \boldsymbol{g} \cdot \boldsymbol{\varphi} \, \mathrm{d}\Omega = 0, \quad \text{(B2)}$$

from which integration by parts gives

$$\int_\Omega \boldsymbol{\sigma} : \nabla \boldsymbol{\varphi} \, \mathrm{d}\Omega - \int_{\partial\Omega} \boldsymbol{\varphi} \cdot \boldsymbol{\sigma}(\boldsymbol{v}, p) \cdot \boldsymbol{n} \, \mathrm{d}\Gamma = \int_\Omega \rho_{\mathrm{i}} \boldsymbol{g} \cdot \boldsymbol{\varphi} \, \mathrm{d}\Omega, \quad \text{(B3)}$$

in which $\partial\Omega$ denotes the boundary of the domain Ω.

The first and second terms can be particularized to our problem. First, using (5), we get

$$\boldsymbol{\sigma} : \nabla \boldsymbol{\varphi} = A^{-\frac{1}{n}} |\boldsymbol{D}(\boldsymbol{v})|^{\frac{1}{n}-1} \boldsymbol{D}(\boldsymbol{v}) : \boldsymbol{D}(\boldsymbol{\varphi}), \quad \text{(B4)}$$

in which we have exploited the fact that $\boldsymbol{\varphi} \in \mathcal{X}$ and that $\boldsymbol{D}(\boldsymbol{v})$ is symmetric. Second, we apply the boundary conditions (6)–(8). The only remaining contributions are the ones over the basal boundary and the calving front, which gives

$$\int_{\partial\Omega} \boldsymbol{\varphi} \cdot \boldsymbol{\sigma}(\boldsymbol{v}, p) \cdot \boldsymbol{n} \, \mathrm{d}\Gamma = - \int_{\Gamma_{\mathrm{b}}} c |\boldsymbol{v}|^{m-1} \boldsymbol{v} \cdot \boldsymbol{\varphi} \, \mathrm{d}\Gamma \quad \text{(B5a)}$$

$$- \int_{\Gamma_{\mathrm{cf}}} p_{\mathrm{w}} \boldsymbol{\varphi} \cdot \boldsymbol{n} \, \mathrm{d}\Gamma, \quad \text{(B5b)}$$

where we have introduced $c := (\tau_{\mathrm{ref}}/u_{\mathrm{ref}}^m)(N/N_{\mathrm{ref}})^q$ to simplify the notations. Combining these results in (B3) leads to the following variational formulation: find $\boldsymbol{v} \in \mathcal{X}$ such that

$$\int_\Omega A^{-\frac{1}{n}} |\boldsymbol{D}(\boldsymbol{v})|^{\frac{1}{n}-1} \boldsymbol{D}(\boldsymbol{v}) : \boldsymbol{D}(\boldsymbol{\varphi}) \, \mathrm{d}\Omega \quad \text{(B6a)}$$

$$+ \int_{\Gamma_{\mathrm{b}}} c |\boldsymbol{v}|^{m-1} \boldsymbol{v} \cdot \boldsymbol{\varphi} \, \mathrm{d}\Gamma \quad \text{(B6b)}$$

$$+ \int_{\Gamma_{\mathrm{cf}}} p_{\mathrm{w}} \boldsymbol{\varphi} \cdot \boldsymbol{n} \, \mathrm{d}\Gamma - \int_\Omega \rho_{\mathrm{i}} \boldsymbol{g} \cdot \boldsymbol{\varphi} \, \mathrm{d}\Omega = 0, \quad \text{(B6c)}$$

for all $\boldsymbol{\varphi} \in \mathcal{X}$. A key result is that this formulation is equivalent to the minimization of the following functional (Jouvet

and Rappaz, 2011; Jouvet, 2016):

$$\mathcal{J}(\boldsymbol{v}) := \int_\Omega \frac{2A^{-\frac{1}{n}}}{1+1/n} |\boldsymbol{D}(\boldsymbol{v})|^{1+\frac{1}{n}} \,\mathrm{d}\Omega \tag{B7a}$$

$$+ \int_{\Gamma_\mathrm{b}} \frac{c}{1+m} |\boldsymbol{v}|^{1+m} \,\mathrm{d}\Gamma \tag{B7b}$$

$$+ \int_{\Gamma_\mathrm{cf}} p_\mathrm{w}\, \boldsymbol{v}\cdot\boldsymbol{n} \,\mathrm{d}\Gamma - \int_\Omega \rho_\mathrm{i} \boldsymbol{g}\cdot\boldsymbol{v} \,\mathrm{d}\Omega. \tag{B7c}$$

B2 Elimination of the vertical velocity

The functional $\mathcal{J}$ can be written as a function of $\boldsymbol{u}$ only by applying suitable assumptions. We consider each of its components separately.

The first term involves the strain-rate tensor $\boldsymbol{D}(\boldsymbol{v})$. It can be simplified using the small aspect ratio of ice sheets, which leads to strain rates $\partial_x w$ and $\partial_y w$ that are much smaller than $\partial_z u$ and $\partial_z v$. Consequently, $\boldsymbol{D}(\boldsymbol{v})$ can be written as

$$\begin{pmatrix} \partial_x u & \frac{1}{2}(\partial_y u + \partial_x v) & \frac{1}{2}(\partial_z u + \partial_x w) \\ \frac{1}{2}(\partial_y u + \partial_x v) & \partial_y v & \frac{1}{2}(\partial_z v + \partial_y w) \\ \frac{1}{2}(\partial_z u + \partial_x w) & \frac{1}{2}(\partial_z v + \partial_y w) & \partial_z w \end{pmatrix}$$

$$\approx \begin{pmatrix} \partial_x u & \frac{1}{2}(\partial_y u + \partial_x v) & \frac{1}{2}\partial_z u \\ \frac{1}{2}(\partial_y u + \partial_x v) & \partial_y v & \frac{1}{2}\partial_z v \\ \frac{1}{2}\partial_z u & \frac{1}{2}\partial_z v & -\partial_x u - \partial_y v \end{pmatrix}, \tag{B8}$$

in which we have also used the incompressibility condition (4b) for $\partial_z w$. As this expression depends only on the horizontal velocity, we can abuse notation and write it as $\boldsymbol{D}(\boldsymbol{u})$. For the second term, we rely on the non-penetration condition (7a) at the base, which yields $w = \nabla b \cdot \boldsymbol{u}$. Hence, over Γ_b,

$$|\boldsymbol{v}|^2 = |\boldsymbol{u}|^2 + w^2 = |\boldsymbol{u}|^2 + (\partial_x b)^2 u^2 + (\partial_y b)^2 v^2, \tag{B9}$$

which defines a norm for $\boldsymbol{u}$ that we write $|\boldsymbol{u}|_\mathrm{b}$. Finally, for the last two terms, we successively compute

$$\int_\Omega (\nabla s \cdot \boldsymbol{u} - w) \,\mathrm{d}\Omega$$

$$= \int_\Omega \nabla(s-z)\cdot\boldsymbol{v} \,\mathrm{d}\Omega \tag{B10a}$$

$$= \int_{\partial\Omega} (s-z)(\boldsymbol{v}\cdot\boldsymbol{n}) \,\mathrm{d}\Gamma - \int_\Omega (s-z)(\nabla\cdot\boldsymbol{v}) \,\mathrm{d}\Omega \tag{B10b}$$

$$= \int_{\Gamma_\mathrm{cf}} (s-z)(\boldsymbol{v}\cdot\boldsymbol{n}) \,\mathrm{d}\Gamma, \tag{B10c}$$

where the last equality follows from the incompressibility condition (4b), the equality $z = s$ on the upper surface, and the non-penetration condition at the lower surface. Hence, we have

$$\int_\Omega \rho_\mathrm{i} \boldsymbol{g}\cdot\boldsymbol{v} \,\mathrm{d}\Omega = -\int_\Omega \rho_\mathrm{i} g \nabla s \cdot \boldsymbol{u} \,\mathrm{d}\Omega \tag{B11a}$$

$$+ \int_{\Gamma_\mathrm{cf}} \rho_\mathrm{i} g (s-z)(\boldsymbol{v}\cdot\boldsymbol{n}) \,\mathrm{d}\Gamma. \tag{B11b}$$

Because the calving front is vertical, its outward normal has no vertical component, and therefore $\boldsymbol{v}\cdot\boldsymbol{n} = \boldsymbol{u}\cdot\boldsymbol{n}$ on Γ_cf. Using the previous identity, the contributions in $\mathcal{J}$ associated with the hydrostatic water pressure and the gravity term (the last two terms in (B7)) are given by

$$\int_\Omega \rho_\mathrm{i} g \nabla s \cdot \boldsymbol{u} \,\mathrm{d}\Omega - \int_{\Gamma_\mathrm{cf}} [\rho_\mathrm{i} g (s-z) - p_\mathrm{w}]\, \boldsymbol{u}\cdot\boldsymbol{n} \,\mathrm{d}\Gamma. \tag{B12}$$

The term in brackets is the local pressure imbalance between the cryostatic ice pressure and the hydrostatic water pressure at the calving front.

Overall, it follows that the functional $\mathcal{J}$ can be written as a function of $\boldsymbol{u}$ only:

$$\mathcal{J}(\boldsymbol{u}) := \int_\Omega \frac{2A^{-\frac{1}{n}}}{1+1/n} |\boldsymbol{D}(\boldsymbol{u})|^{1+\frac{1}{n}} \,\mathrm{d}\Omega \tag{B13a}$$

$$+ \int_{\Gamma_\mathrm{b}} \frac{c}{1+m} |\boldsymbol{u}|_\mathrm{b}^{1+m} \,\mathrm{d}\Gamma \tag{B13b}$$

$$+ \int_\Omega \rho_\mathrm{i} g \nabla s \cdot \boldsymbol{u} \,\mathrm{d}\Omega \tag{B13c}$$

$$- \int_{\Gamma_\mathrm{cf}} [\rho_\mathrm{i} g (s-z) - p_\mathrm{w}]\, \boldsymbol{u}\cdot\boldsymbol{n} \,\mathrm{d}\Gamma. \tag{B13d}$$

Appendix C: Thermodynamics

C1 Enthalpy formulation

Modeling the thermal state of glaciers is complex, with some regions at the pressure-melting point while others are not (Fowler and Larson, 1978; Hutter, 1982; Blatter and Hutter, 1991; Hewitt and Schoof, 2017). We here use an enthalpy formulation (Aschwanden et al., 2012) implemented in the Parallel Ice Sheet Model (Khroulev and the PISM Authors, 2020). Two ice regimes are typically defined:

- cold ice: $T < T_\mathrm{pmp}$ and $\omega = 0$;
- temperate ice: $T = T_\mathrm{pmp}$ and $\omega > 0$.

The pressure-melting-point temperature T_pmp depends on pressure through the Clausius–Clapeyron relation:

$$T_\mathrm{pmp} = T^0_\mathrm{pmp} - \beta\, p, \tag{C1}$$

with $p = \rho g (s-z)$ the ice pressure (approximated as cryostatic), $\beta > 0$ the Clausius–Clapeyron constant and T^0_pmp the pressure-melting-point temperature at standard pressure. To avoid dealing with the two variables (T, ω), it is convenient to introduce the enthalpy $E = E(\boldsymbol{x}, t)$ as follows:

$$E = \begin{cases} E_\mathrm{pmp} + c_\mathrm{i}(T - T_\mathrm{pmp}), & \text{for cold ice,} \quad \text{(C2a)} \\ E_\mathrm{pmp} + L\omega, & \text{for temperate ice,} \quad \text{(C2b)} \end{cases}$$

where c_i and L are the specific heat capacity and latent heat of ice, respectively; see Fig C1. Importantly, each value of E

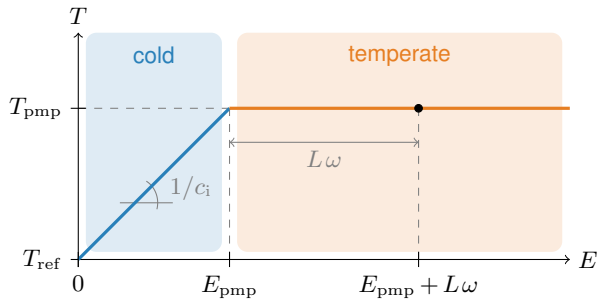

Figure C1. Rationale behind the introduction of the enthalpy. When the ice is cold (blue region), meaning it is below the melting point, the enthalpy is simply proportional to the temperature minus a reference temperature. When the ice is temperate (orange region), the enthalpy continues to increase. In this case, the additional component $L\omega$ accounts for the creation of water content through energy transfer.

corresponds to a unique state (T, ω), so the values of these variables can be inferred once E is known. Since enthalpy can be defined up to an additive constant, we can choose the offset value for E_{pmp}; here, we choose it in a such a way that the enthalpy is zero at a reference temperature T_{ref}:

$$E_{\mathrm{pmp}} = c_{\mathrm{i}}(T_{\mathrm{pmp}} - T_{\mathrm{ref}}). \tag{C3}$$

In the enthalpy formulation, energy conservation takes the form of an advection-diffusion equation for the enthalpy (Aschwanden et al., 2012):

$$\rho_{\mathrm{i}}\left(\frac{\partial E}{\partial t} + \boldsymbol{v}\cdot\nabla E\right) = \frac{\partial}{\partial z}\left(K\frac{\partial E}{\partial z}\right) + \Phi - \rho_{\mathrm{w}} L D_{\mathrm{w}}, \tag{C4}$$

where K is the effective diffusivity, Φ is the rate of viscous dissipation, and D_{w} is a water drainage function. The effective diffusivity depends on the thermal state of the ice, and is typically chosen as $K = k_{\mathrm{i}}/c_{\mathrm{i}}$ for cold ice, and $K = \epsilon\, k_{\mathrm{i}}/c_{\mathrm{i}}$ for temperate ice, with k_{i} the thermal conductivity of (cold) ice and in which $\epsilon \ll 1$. The rate of viscous dissipation is given by

$$\Phi = \boldsymbol{\sigma} : \boldsymbol{D} = 2\eta \boldsymbol{D} : \boldsymbol{D}. \tag{C5}$$

The water drainage function is typically parametrized as a simple function of ω (e.g., Greve, 1997; Aschwanden et al., 2012).

Equation (C4) needs to be accompanied by suitable boundary conditions. At the top ice surface, the enthalpy equation is constrained by the surface temperature (or equivalently, the enthalpy) provided by the climate forcing, which is enforced as a Dirichlet condition:

$$T = T_{\mathrm{s}}. \tag{C6}$$

The conditions at the base are more complex, and result from the interactions between the heat energy balance at the base and the hydrological state (whether there is liquid water at the base). There are four categories (Aschwanden et al., 2012; Wang et al., 2020):

(a) Cold base and dry bed: $E < E_{\mathrm{pmp}}$ and $h_{\mathrm{w}} = 0$;

(b) Cold base and wet bed: $E < E_{\mathrm{pmp}}$ and $h_{\mathrm{w}} > 0$;

(c) Temperate base and no basal temperate layer: $E \geq E_{\mathrm{pmp}}$ at the base but $E < E_{\mathrm{pmp}}$ just above it;

(d) Temperate base and basal temperate layer: $E < E_{\mathrm{pmp}}$ at the base and just above it.

This results in the following Dirichlet or Neumann boundary conditions for the enthalpy at the base:

$$K\partial_z E = q_{\mathrm{geo}} + \phi, \tag{C7a}$$
$$E = E_{\mathrm{pmp}}, \tag{C7b}$$
$$E = E_{\mathrm{pmp}}, \tag{C7c}$$
$$\partial_z E = 0. \tag{C7d}$$

Here, ϕ is the frictional dissipation at the ice-bed interface, given by $\phi = c|\boldsymbol{u}|_{\mathrm{b}}^{m+1}$ for the Weertman friction law.

The thermal state in the ice impact ice flows at two levels: through the modification of the Arrhenius factor in the effective viscosity, and through a change in the basal conditions.

C1.1 Arrhenius factor

The effect of temperature and water content on the Arrhenius factor follows Glen–Paterson–Budd–Lliboutry–Duval law (Glen, 1955; Paterson and Budd, 1982; Lliboutry and Duval, 1985):

$$A(T,\omega) = A_0(1 + 181.25\,\omega)\exp(-Q/RT_{\mathrm{pa}}), \tag{C8}$$

in which A_0 is a prefactor, Q is an activation energy, R is the universal gas constant, and $T_{\mathrm{pa}} := T + \beta p$ is the pressure-adjusted temperature, defined in a such a way that the pressure-melting-point is at T^0_{pmp}. Both A_0 and Q take different values as a function of the temperature, based on empirical relations (see, e.g., Cuffey and Paterson, 2010).

C1.2 Melt rate

At the bed, the melt rate is computed as follows:

$$\dot{a}_{\mathrm{b}} = \frac{1}{\rho_{\mathrm{i}}(1-\omega)L}\left(q_{\mathrm{geo}} + \phi - K\frac{\partial E}{\partial z}\right) + \frac{\rho_{\mathrm{w}}}{\rho_{\mathrm{i}}}\int_b^s D_{\mathrm{w}}(\omega)\,\mathrm{d}z. \tag{C9}$$

The first term stems from a heat balance at the ice-bedrock interface, and the second one results from the drainage through the full ice column.

C1.3 Till effective pressure

Assuming a simple local storage model, water then accumulates into the till as follows (Bueler and van Pelt, 2015):

$$\partial_t h_\mathrm{w} = \frac{\rho_\mathrm{i}}{\rho_\mathrm{w}} \dot{a}_\mathrm{b} - \dot{a}_\mathrm{dr}, \tag{C10}$$

in which h_w is the effective thickness of water in the till, and a_dr is an aquifer drainage rate. The thickness h_w is further capped to a maximum value $h_\mathrm{w} = h_\mathrm{w}^\mathrm{max}$. The effect of till water content on ice sliding can be modeled by assuming that the friction coefficient c follows a Mohr–Coulomb behavior (Hooke et al., 1997; Tulaczyk et al., 2000):

$$c = \tan\phi \, N_\mathrm{till}, \tag{C11}$$

in which ϕ denotes a till friction angle and the till effective pressure N_till is obtained from the following equations (Bueler and van Pelt, 2015):

$$N_\mathrm{till} = N_0 \left(\delta \rho_\mathrm{i} g h / N_0\right)^{s_\mathrm{till}} 10^{(e_0/C_\mathrm{c})(1-s_\mathrm{till})}, \tag{C12a}$$

$$s_\mathrm{till} = h_\mathrm{w} / h_\mathrm{w}^\mathrm{max}. \tag{C12b}$$

Here, δ is a fixed fraction corresponding to the lower bound for N_till with respect to the ice overburden pressure, N_0 is a reference effective pressure, e_0 is the reference void ratio at N_0, and C_c is the till compressibility.

C2 Numerics

A different vertical discretization is adopted for the enthalpy field, because its vertical structure is typically qualitatively different from that of the velocity. While the velocity profile is generally smooth and well approximated by low-order polynomials, enthalpy often exhibits stronger variations near the bed and a discontinuity in its vertical gradient at the cold–temperate transition surface. We therefore write

$$E(x_i, y_j, z) = \sum_{k=1}^{\tilde{N}_z} (\mathbf{E}_H)_{kji} \, \tilde{\phi}_k(\zeta(z)), \tag{C13}$$

where $\mathbf{E}_H \in \mathbb{R}^{\tilde{N}_z \times N_y \times N_x}$ denotes the enthalpy degrees of freedom. For the basis functions $\{\tilde{\phi}_k\}_{k=1}^{\tilde{N}_z}$, we use Lagrange shape functions defined on a vertically stretched grid whose node density increases quadratically towards the bed in order to better resolve this region. In practice, $N_z, \tilde{N}_z \ll N_x, N_y$.

Finally, the Arrhenius factor is represented as an effective horizontal tensor field $\mathbf{A}_H \in \mathbb{R}^{N_y \times N_x}$, obtained by vertically averaging its value according to the rheological exponent in equation (9):

$$(\mathbf{A}_H)_{ji} := \left(\frac{1}{h} \int_b^s A^{-\frac{1}{n}} \, \mathrm{d}z \right)^{-n} \Bigg|_{(x,y)=(x_i,y_j)}. \tag{C14}$$

The enthalpy field, discretized as described in subsection 3.3, is updated by solving the advection–diffusion equation (C4) subject to the boundary conditions (C6)–(C7). This equation is solved with an operator-splitting approach: horizontal advection is treated explicitly in the first step, and vertical advection-diffusion implicitly in the second. Denoting by $E^n(\boldsymbol{x}) := E(\boldsymbol{x}, t^n)$ the enthalpy at time t^n, the first step reads:

$$E^{n+\frac{1}{2}} = E^n - \Delta t \left(u \frac{\partial E^n}{\partial x} + v \frac{\partial E^n}{\partial y} \right), \tag{C15}$$

where horizontal derivatives are approximated by upwind finite differences at every grid node. The second step reads:

$$\rho_\mathrm{i} \left(\frac{E^{n+1} - E^{n+\frac{1}{2}}}{\Delta t} + w \frac{\partial E^{n+1}}{\partial z} \right) \tag{C16a}$$

$$= \frac{\partial}{\partial z} \left(K \frac{\partial E}{\partial z} \right) + \Phi - \rho_\mathrm{w} L \, D_\mathrm{w}(\omega^{n+\frac{1}{2}}), \tag{C16b}$$

where vertical advection is discretized with an upwind scheme and diffusion with a centered scheme. Since equation (C16) contains no horizontal coupling, it decouples into $N_x \times N_y$ independent tridiagonal systems of size $\tilde{N}_z$ (one per ice column), solved in parallel using the Thomas algorithm. Following the enthalpy solve, excess water content in the ice column is drained and added to the basal melt rate via equation (C9); the till water thickness and sliding coefficient are then updated according to equations (C10) and (C11), respectively.

Appendix D: Benchmarks

To verify the correctness of the core ice-flow and thermodynamics (enthalpy) modules, IGM has been successfully run against standard community benchmarks. Specifically, we reproduce experiment B of the ISMIP-HOM intercomparison (Pattyn et al., 2008), which tests higher-order ice-flow solvers against a full-Stokes reference (Fig. D1), as well as experiments A and B of the Kleiner et al. (2015) enthalpy benchmark, which assess the transient thermodynamics against numerical and analytical reference solutions, respectively (Figs. D2, D3). We also reproduce the MISMIP+ initialization set-up of Cornford et al. (2020); the resulting steady-state velocity pattern and grounding-line position are consistent with the expected benchmark behavior (Fig. D4). In all cases, IGM's results agree closely with the reference solutions or expected benchmark behavior, confirming the correctness of the implementation.

Appendix E: Project folder layout

A typical IGM project lives in a single directory containing four subdirectories:

```
experiment/params.yaml # config. file
data/...               # input data
```

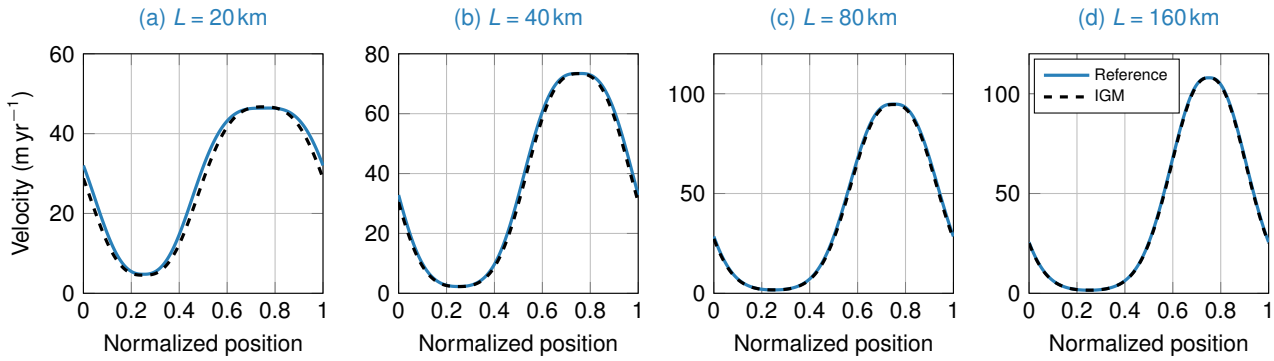

Figure D1. ISMIP-HOM test case, experiment B (Pattyn et al., 2008). Reference is 'oga', corresponding to the full-Stokes submission from the finite-element code Elmer/Ice (Gagliardini and Zwinger, 2008).

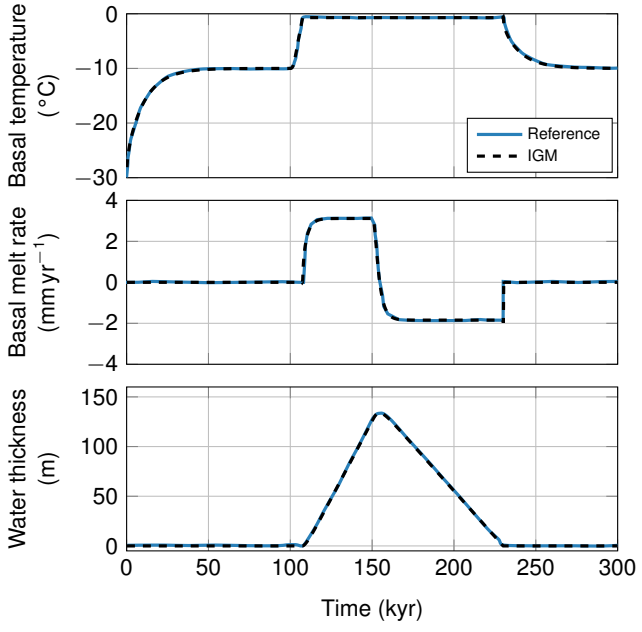

Figure D2. Experiment A of Kleiner et al. (2015). Reference correspond to the numerical solutions shown in Kleiner et al. (2015).

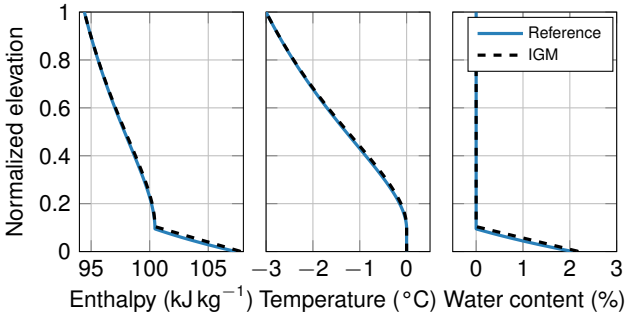

Figure D3. Experiment B of Kleiner et al. (2015). Reference is the analytical solution of the problem.

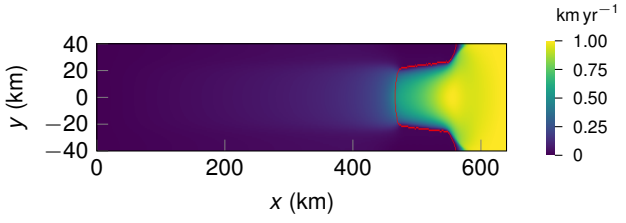

Figure D4. MISMIP+ initialization experiment following the setup of Cornford et al. (2020), computed with IGM at a horizontal grid resolution of 1 km. Colors show the steady-state surface speed, highlighting the transition from slow grounded flow to fast flow across the marine outlet. The red contour marks the grounding-line position.

```
output/                 # output directory
user/...                # opt. user modules
```

`experiment/` Holds one or more YAML configuration files (typically `params.yaml`). A single file fully specifies the simulation: which modules to run, their parameters, and any Hydra overrides. Several files can coexist to represent different experiments or sensitivity runs.

`data/` Holds the input data required by the selected input modules, such as NetCDF files with bed topography, ice thickness, or surface elevation. The path and filenames are referenced from the YAML configuration.

`output/` Created by Hydra at run time. By default, Hydra names each run subfolder after the launch date and time (e.g. `output/2026-04-15/14-32-07/`), so successive runs do not overwrite each other; this pattern can be customized via Hydra's `hydra.run.dir` setting. Each run folder contains the outputs produced by the output modules together with a `.hydra/` subfolder that records the full, resolved configuration used for that run (Section 5.3).

`user/` Optional. Contains user-written modules that extend or replace core modules without modifying the IGM source tree. The internal layout of a user module is detailed in Appendix H.

The simulation is launched from the project root with `igm_run +experiment=params`; Hydra then composes the final configuration by combining built-in defaults, the YAML file, and any command-line overrides (Appendix I).

Appendix F: Core code structure

The `igm_run.py` script is the command-line entry point: it is invoked as `igm_run` after installation and delegates execution to the runner in `common/`, which composes the module pipeline from the resolved Hydra configuration. Schematically, the runner `igm_run.py` executes the following loop:

```
# build module lists from the Hydra cfg
inputs, proc., outputs = setup(cfg, state)

# input modules: load data once
for m in inputs:
    m.run(cfg, state)

# init output modules: e.g., open files
for m in outputs:
    m.initialize(cfg, state)

# init process modules
for m in assimilations+processes:
    m.initialize(cfg, state)

# time loop
while state.continue_run:
    for m in assimilations+processes:
        m.update(cfg, state) # upt state
    for m in outputs:
        m.run(cfg, state) # write results

# finalize, e.g. write score
for m in assimilations+processes:
    m.finalize(cfg, state)
```

This interface defines the contract that each module category must respect:

- an *input* module must provide a `run(cfg, state)` function (called once to load data into `state`),
- a *assimilation* or *process* module must provide three functions `initialize(cfg, state)`, `update(cfg, state)`, and `finalize(cfg, state)` (called once before, at every time step, and once after the time loop, respectively),
- an *output* module must provide `initialize(cfg, state)` (called once before the loop, e.g. to open files) and `run(cfg, state)` (called at every time step to write the current state).

The shared `state` object carries the glacier state as TensorFlow tensors: for instance, `state.thk` is the ice thickness, `state.smb` the surface mass balance, `state.t` the current time, and so on. These fields can be read or assigned freely within a module, and changes propagate to all subsequent modules in the pipeline. The `cfg` object exposes the resolved Hydra configuration as a hierarchically structured parameter tree (e.g. `cfg.processes.time.end`; see Appendix I). A *user module* (Appendix H) reproduces this same layout on a smaller scale inside the project's `user/` directory, so that user-defined modules can be picked up by the runner exactly as if they were part of the core package, without modifying the IGM source tree.

Appendix G: Core process modules

Table G1 lists the main core process modules (Fig. 4) shipped with IGM, together with the physical or numerical process they handle and their primary I/O fields on the shared `state` object. Each module is referenced by the name that appears in the Hydra configuration (e.g. `cfg.processes.iceflow`). An up-to-date list of all modules, with complete parameter documentation, is maintained online at https://igm-model.org.

Appendix H: Creating a user module

IGM users can easily add custom physical processes without modifying the core source code (Section 5, Appendix F). A user module consists of a Python file (the code) and a YAML file (its default parameters), placed inside the `user/` directory of the project (Appendix E):

```
user/
  code/
    processes/
      mysmb.py      # module code
  conf/
    processes/
      mysmb.yaml    # default params
```

As described in Appendix F, every process module must define three functions `initialize(cfg, state)`, `update(cfg, state)`, and `finalize(cfg, state)`, called once before the time loop, at every time step, and once after the time loop, respectively. All three must be present, even if empty (`pass`). The `cfg` and `state` objects are the same as described in Appendix F: `cfg` provides hierarchical access to the module's parameters (e.g. `cfg.processes.mysmb.meanela`), and `state` carries the glacier fields as TensorFlow tensors that can be read or assigned freely. As a concrete illustration, we imple-

Table G1. Overview of the main core process modules in IGM (Fig. 4), with their role and key I/O variables on the shared `state` object. For brevity, only the primary fields are listed.

Module	Process modeled	Inputs	Outputs
`climate`	Climate forcing	`usurf`	`air_temp`, `air_temp_sd`, `precipitation`
`smb`	Surface mass balance	`air_temp`, `precipitation`	`smb`
`enthalpy`	Englacial enthalpy and basal melt	`thk`, `air_temp`, `arrhenius`, `U`, `V`, `W`	`E`, `basal_melt_rate`
`arrhenius`	Arrhenius factor	`E`	`arrhenius`
`subglacial_hydrology`	Subglacial hydrology	`thk`, `topg`, `water_level`,	`effective_pressure`
`iceflow`	Ice-flow solver	`thk`, `usurf`, `topg`, `arrhenius`	`U`, `V`, `W`, `ubar`, `vbar`
`calving_rate`	Calving rate	`thk`, `topg`, `water_level`, `ubar`, `vbar`	`calving_rate`
`time`	Time stepping	`ubar`, `vbar`	`t`, `dt`
`thk`	Ice-thickness evolution	`thk`, `topg`, `ubar`, `vbar`, `smb`	`thk`, `usurf`, `divflux`
`particles`	Lagrangian particle tracking	`U`, `V`, `W`, `thk`	`particle`

ment a simple SMB module. The Python file `mysmb.py` in `user/code/processes/smb/` reads:

```
def initialize(cfg, state):
    pass

def update(cfg, state):
    ELA = cfg.processes.mysmb.meanela
    state.smb = state.usurf - ELA
    state.smb *= 0.005

def finalize(cfg, state):
    pass
```

The companion parameter file `mysmb.yaml` sets the default mean ELA:

```
meanela: 3000.0
```

Finally, the experiment configuration file `params.yaml` registers the user module and inserts it into the process pipeline:

```
defaults:
  - /user/conf/processes@processes: mysmb

  - override /processes:
    - mysmb
    - iceflow
    - time
    - thk
  - override /outputs:
    - local
```

The ordering of the process list matters: here, `mysmb` is called before `iceflow` so that the updated SMB is available to subsequent modules.

Although the example above illustrates a process module, user-defined *input* and *output* modules can be added in exactly the same way: they simply have to be placed under `user/code/inputs/` or `user/code/outputs/` (with a matching YAML in `user/conf/inputs/` or `user/conf/outputs/`) and to implement the functions required by their category, as specified in Appendix F (`run` for an input module; `initialize` and `run` for an output module).

Appendix I: Parameter management

IGM relies on the Hydra configuration framework (Yadan, 2019) to manage all simulation parameters. At run time, Hydra collects the per-module YAML parameter files shipped with the core package as well as those provided by any user modules, applies several levels of overrides (see below), and exposes the resulting hierarchical parameter tree as the `cfg` object passed to every module function (Appendix F). Parameters follow a three-tier override hierarchy, from lowest to highest priority:

1. **Built-in defaults.** Each IGM module ships with a YAML file defining its default parameter values. For example, the built-in defaults for the surface mass balance module include:

```
smb:
  method: simple
  simple:
    update_freq: 1.0
    file: param.txt
    array: []
```

2. **Experiment file.** The user-written YAML file (e.g., `experiment/params.yaml`) selects which modules to activate via a `defaults` list, and overrides specific parameter values:

```
# @package _global_

defaults:
```

```
    - override /inputs:
      - local
    - override /processes:
      - smb
      - iceflow
      - time
      - thk
    - override /outputs:
      - local

processes:
   smb:
      method: simple
      simple:
        update_freq: 1.0
```

3. **Command-line overrides.** Any parameter can be overridden at runtime using dot notation:

```
igm_run +experiment=params \
        processes.time.end=2200
```

Every IGM run automatically creates a `.hydra/` directory containing the complete, merged configuration:

```
output/
  .hydra/
    config.yaml    # final/merged params
    overrides.yaml # CLI overrides applied
```

This record is sufficient to reproduce the simulation exactly: it captures the selected modules, and all parameter final values (whether set by defaults, the experiment file, or the command line) settings.

Author contributions. The IGM developer team includes Guillaume Jouvet (UNIL), Brandon Finley (UNIL), Thomas Gregov (UNIL) and Sebastian Rosier (UZH) from the Universities of Lausanne (UNIL) and Zürich (UZH). G.J. conceived and created IGM, was the main applicant to acquire funding to develop it, and was the primary developer through 2025, writing the bulk of the code modules and documentation, which were subsequently extended and enriched by all co-authors. G.J. wrote most of the manuscript. B.F. led IGM's software engineering from 2024, creating the core code design framework, integrated Hydra, created the IGM website, and implemented the automated workflows (CI/CD) such as the PyPI packaging, website documentation, and unit tests. B.F. made significant performance optimizations and code refactoring for the iceflow and particle modules. He wrote most of the software design section and contributed to figures. From 2025 onward, T.G. significantly enhanced the iceflow and enthalpy modules, improving robustness, code readability, documentation and test coverage. T.G. revised the manuscript (Physics and Numerics Sections) to reflect the latest code updates (v3.0.1 → v3.2.0), wrote the Ice flow subsection, Appendix B1, ran the benchmark tests and contributed to figures. From 2025 onward, S.R. fully redesigned the data assimilation and pretraining modules, improving robustness, code readability, testing and benchmarking. S.R. reviewed the manuscript and provided revisions and comments. All four authors are core developers of IGM, and contributed to the paper writing. Additional specific contributions to the IGM codebase have been made by Flavio Calvo, Samuel Cook, Guillaume Cordonnier, Oskar Hermann, Thomas Frank, Andreas Henz, Tancrède Leger, Fabien Maussion, Jürgen Mey, Margot Sirdey, Gilliam Smith, Claire-Mathilde Stücki, Ethan Welty.

Acknowledgements. This research has been supported by the Swiss National Science Foundation (SNSF) grant no. 200021-231427 (PINNACLE).

References

Abadi, M., Agarwal, A., Barham, P., Brevdo, E., Chen, Z., Citro, C., Corrado, G. S., Davis, A., Dean, J., Devin, M., Ghemawat, S., Goodfellow, I., Harp, A., Irving, G., Isard, M., Jia, Y., Jozefowicz, R., Kaiser, L., Kudlur, M., Levenberg, J., Mané, D., Monga, R., Moore, S., Murray, D., Olah, C., Schuster, M., Shlens, J., Steiner, B., Sutskever, I., Talwar, K., Tucker, P., Vanhoucke, V., Vasudevan, V., Viégas, F., Vinyals, O., Warden, P., Wattenberg, M., Wicke, M., Yu, Y., and Zheng, X.: TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems, software available from tensorflow.org, 2015.

Ahlkrona, J., Henry, A. C. J., and Löfgren, A.: A fully implicit second order method for viscous free surface Stokes flow – application to glacier simulations, Geoscientific Model Development, 19, 2333–2348, https://doi.org/10.5194/gmd-19-2333-2026, 2026.

Aissa, M., Verstraete, T., and Vuik, C.: Toward a GPU-aware comparison of explicit and implicit CFD simulations on structured meshes, Computers & Mathematics with Applications, 74, 201–217, https://doi.org/10.1016/j.camwa.2017.03.003, 2017.

Albrecht, T., Martin, M., Haseloff, M., Winkelmann, R., and Levermann, A.: Parameterization for subgrid-scale motion of ice-shelf calving fronts, The Cryosphere, 5, 35–44, 2011.

Aschwanden, A., Bueler, E., Khroulev, C., and Blatter, H.: An enthalpy formulation for glaciers and ice sheets, Journal of Glaciology, 58, 441–457, https://doi.org/10.3189/2012jog11j088, 2012.

Barnhart, K. R., Hutton, E. W., Tucker, G. E., Gasparini, N. M., Istanbulluoglu, E., Hobley, D. E., Lyons, N. J., Mouchene, M., Nudurupati, S. S., Adams, J. M., et al.: Landlab v2. 0: a software package for Earth surface dynamics, Earth Surface Dynamics, 8, 379–397, https://doi.org/10.5194/esurf-8-379-2020, 2020.

Blatter, H.: Velocity and stress fields in grounded glaciers: a simple algorithm for including deviatoric stress gradients, Journal of Glaciology, 41, 333–344, https://doi.org/10.3189/s002214300001621x, 1995.

Blatter, H. and Hutter, K.: Polythermal conditions in arctic glaciers, Journal of Glaciology, 37, 261–269, https://doi.org/10.3189/s0022143000007279, 1991.

Bolibar, J., Sapienza, F., Maussion, F., Lguensat, R., Wouters, B., and Pérez, F.: Universal differential equations for glacier ice flow modelling, Geoscientific Model Development, 16, 6671–6687, https://doi.org/10.5194/gmd-16-6671-2023, 2023.

Bondzio, J. H., Seroussi, H., Morlighem, M., Kleiner, T., Rückamp, M., Humbert, A., and Larour, E. Y.: Modelling calving front dynamics using a level-set method: application to Jakobshavn Isbræ, West Greenland, The Cryosphere, 10, 497–510, 2016.

Brædstrup, C. F., Damsgaard, A., and Egholm, D. L.: Ice-sheet modelling accelerated by graphics cards, Computers & Geosciences, 72, 210–220, https://doi.org/10.1016/j.cageo.2014.07.019, 2014.

Brinkerhoff, D., Aschwanden, A., and Fahnestock, M.: Constraining subglacial processes from surface velocity observations using surrogate-based Bayesian inference, Journal of Glaciology, 67, 385–403, https://doi.org/10.1017/jog.2020.112, 2021.

Brown, C. S., Meier, M. F., and Post, A.: Calving speed of Alaska tidewater glaciers, with application to Columbia Glacier, U.S. Geological Survey Professional Paper, 1258-C, C1–C13, 1982.

Budd, W. F., Keage, P. L., and Blundy, N. A.: Empirical studies of ice sliding, Journal of Glaciology, 23, 157–170, https://doi.org/10.3189/S0022143000029804, 1979.

Bueler, E. and van Pelt, W.: Mass-conserving subglacial hydrology in the Parallel Ice Sheet Model version 0.6, Geoscientific Model Development, 8, 1613–1635, https://doi.org/10.5194/gmd-8-1613-2015, 2015.

Burgard, C., Jourdain, N. C., Mathiot, P., Smith, R. S., Schäfer, R., Caillet, J., Finn, T. S., and Johnson, J. E.: Emulating Present and Future Simulations of Melt Rates at the Base of Antarctic Ice Shelves With Neural Networks, Journal of Advances in Modeling Earth Systems, 15, https://doi.org/10.1029/2023ms003829, 2023.

Calov, R. and Greve, R.: A semi-analytical solution for the positive degree-day model with stochastic temperature variations, Journal of Glaciology, 51, 173–175, https://doi.org/10.3189/172756505781829601, 2005.

Cheng, G., Krishna, M., and Morlighem, M.: A Python library for solving ice sheet modeling problems using physics-informed neural networks, PINNICLE v1.0, Geoscientific Model Development, 18, 5311–5327, https://doi.org/10.5194/gmd-18-5311-2025, 2025.

Colinge, J. and Rappaz, J.: A strongly nonlinear problem arising in glaciology, M2AN Math. Model. Numer. Anal., 33, 395–406, 1999.

Cordonnier, G., Jouvet, G., Peytavie, A., Braun, J., Cani, M.-P., Benes, B., Galin, E., Guérin, E., and Gain, J.: Forming terrains by glacial erosion, ACM Transactions on Graphics (TOG), 42, 1–14, 2023.

Cornford, S. L., Seroussi, H., Asay-Davis, X. S., Gudmundsson, G. H., Arthern, R., Borstad, C., Christmann, J., Dias dos Santos, T., Feldmann, J., Goldberg, D., Hoffman, M. J., Humbert, A., Kleiner, T., Leguy, G., Lipscomb, W. H., Merino, N., Durand, G., Morlighem, M., Pollard, D., Rückamp, M., Williams, C. R., and Yu, H.: Results of the third Marine Ice Sheet Model Intercomparison Project (MISMIP+), The Cryosphere, 14, 2283–2301, https://doi.org/10.5194/tc-14-2283-2020, 2020.

Cuffey, K. M. and Paterson, W. S. B.: The Physics of Glaciers, Academic Press, Burlington, MA, 4 edn., ISBN 9780123694614, 2010.

Dias dos Santos, T., Morlighem, M., and Brinkerhoff, D.: A new vertically integrated MOno-Layer Higher-Order (MOLHO) ice flow model, The Cryosphere, 16, 179–195, https://doi.org/10.5194/tc-16-179-2022, 2022.

E, W. and Yu, B.: The Deep Ritz Method: A Deep Learning-Based Numerical Algorithm for Solving Variational Problems, Communications in Mathematics and Statistics, 6, 1–12, https://doi.org/10.1007/s40304-018-0127-z, 2018.

Fowler, A. C.: A Theoretical Treatment of the Sliding of Glaciers in the Absence of Cavitation, Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences, 298, 637–681, https://doi.org/10.1098/rsta.1981.0003, 1981.

Fowler, A. C. and Larson, D. A.: On the flow of polythermal glaciers - I. Model and preliminary analysis, Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences, 363, 217–242, https://doi.org/10.1098/rspa.1978.0165, 1978.

Frank, T. and van Pelt, W. J. J.: Ice volume and thickness of all Scandinavian glaciers and ice caps, Journal of Glaciology, 70, e11, 2024.

Gagliardini, O. and Zwinger, T.: The ISMIP-HOM benchmark experiments performed using the Finite-Element code Elmer, The Cryosphere, 2, 67–76, https://doi.org/10.5194/tc-2-67-2008, 2008.

Gagliardini, O., Cohen, D., Råback, P., and Zwinger, T.: Finite-element modeling of subglacial cavities and related friction law, Journal of Geophysical Research: Earth Surface, 112, https://doi.org/10.1029/2006jf000576, 2007.

Gagliardini, O., Zwinger, T., Gillet-Chaulet, F., Durand, G., Favier, L., de Fleurian, B., Greve, R., Malinen, M., Martín, C., Råback, P., Ruokolainen, J., Sacchettini, M., Schäfer, M., Seddik, H., and Thies, J.: Capabilities and performance of Elmer/Ice, a new-generation ice sheet model, Geoscientific Model Development, 6, 1299–1318, https://doi.org/10.5194/gmd-6-1299-2013, 2013.

Gaudreault, S., Rainwater, G., and Tokman, M.: KIOPS: A fast adaptive Krylov subspace solver for exponential integrators, Journal of Computational Physics, 372, 236–255, https://doi.org/10.1016/j.jcp.2018.06.026, 2018.

Glen, J. W.: The creep of polycrystalline ice, Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences, 228, 519–538, https://doi.org/10.1098/rspa.1955.0066, 1955.

Greve, R.: A continuum–mechanical formulation for shallow polythermal ice sheets, Philosophical Transactions of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences, 355, 921–974, 1997.

Hardmeier, F., Miles, E., Muñoz-Hermosilla, J. M., Jouvet, G., and Vieli, A.: The effect of debris supply on glacier evolution: sensitivities and challenges, in: EGU General Assembly 2026, Vienna, Austria, https://doi.org/10.5194/egusphere-egu26-21688, 2026.

He, Q., Perego, M., Howard, A. A., Karniadakis, G. E., and Stinis, P.: A hybrid deep neural operator/finite element method for ice-sheet modeling, Journal of Computational Physics, 492, 112 428, https://doi.org/10.1016/j.jcp.2023.112428, 2023.

Herterich, K.: On the Flow within the Transition Zone between Ice Sheet and Ice Shelf, p. 185–202, Springer Netherlands, ISBN 9789400937451, https://doi.org/10.1007/978-94-009-3745-1_11, 1987.

Hewitt, I. J. and Schoof, C.: Models for polythermal ice sheets and glaciers, The Cryosphere, 11, 541–551, https://doi.org/10.5194/tc-11-541-2017, 2017.

Hobley, D. E., Adams, J. M., Nudurupati, S. S., Hutton, E. W., Gasparini, N. M., Istanbulluoglu, E., and Tucker, G. E.: Creative computing with Landlab: an open-source toolkit for building, coupling, and exploring two-dimensional numerical models of Earth-surface dynamics , Earth Surface Dynamics, 5, 21–46, https://doi.org/10.5194/esurf-5-21-2017, 2017.

Hock, R.: Temperature index melt modelling in mountain areas, Journal of Hydrology, 282, 104–115, https://doi.org/10.1016/s0022-1694(03)00257-9, 2003.

Hoffman, M. J., Perego, M., Price, S. F., Lipscomb, W. H., Zhang, T., Jacobsen, D., Tezaur, I., Salinger, A. G., Tuminaro, R., and Bertagna, L.: MPAS-Albany Land Ice (MALI): a variable-resolution ice sheet model for Earth system modeling using Voronoi grids, Geoscientific Model Development, 11, 3747–3780, https://doi.org/10.5194/gmd-11-3747-2018, 2018.

Hooke, R. L., Hanson, B., Iverson, N. R., Jansson, P., and Fischer, U. H.: Rheology of till beneath Storglaciären, Sweden, Journal of Glaciology, 43, 172–179, https://doi.org/10.3189/S0022143000002938, 1997.

Howard, A. A., Perego, M., Karniadakis, G. E., and Stinis, P.: Multifidelity deep operator networks for data-driven and physics-informed problems, Journal of Computational Physics, 493, 112 462, https://doi.org/10.1016/j.jcp.2023.112462, 2023.

Hugonnet, R., McNabb, R., Berthier, E., Menounos, B., Nuth, C., Girod, L., Farinotti, D., Huss, M., Dussaillant, I., Brun, F., and Kääb, A.: Accelerated global glacier mass loss in the early twenty-first century, Nature, 592, 726–731, https://doi.org/10.1038/s41586-021-03436-z, 2021.

Hutter, K.: A mathematical model of polythermal glaciers and ice sheets, Geophysical & Astrophysical Fluid Dynamics, 21, 201–224, https://doi.org/10.1080/03091928208209013, 1982.

Hutton, E., Barnhart, K., Hobley, D., Tucker, G., Nudurupati, S., Adams, J., Gasparini, N., Shobe, C., Strauch, R., Knuth, J., Mouchene, M., Lyons, N., Litwin, D., Glade, R., Giuseppecipolla95, Manaster, A., Abby, L., Thyng, K., and Rengers, F.: landlab, https://doi.org/10.5281/zenodo.595872, software, 2020.

Jnini, A., Vella, F., and Zeinhofer, M.: Gauss-Newton Natural Gradient Descent for Physics-informed Computational Fluid Dynamics, Computers & Fluids, 307, 106 955, https://doi.org/10.1016/j.compfluid.2025.106955, 2026.

Jouvet, G.: Mechanical error estimators for shallow ice flow models, Journal of Fluid Mechanics, 807, 40–61, https://doi.org/10.1017/jfm.2016.593, 2016.

Jouvet, G.: Inversion of a Stokes glacier flow model emulated by deep learning, Journal of Glaciology, 69, 13–26, https://doi.org/10.1017/jog.2022.41, 2022.

Jouvet, G.: Inversion of a Stokes glacier flow model emulated by deep learning, Journal of Glaciology, 69, 13–26, https://doi.org/10.1017/jog.2022.41, 2023.

Jouvet, G. and Cordonnier, G.: Ice-flow model emulator based on physics-informed deep learning, Journal of Glaciology, p. 1–15, https://doi.org/10.1017/jog.2023.73, 2023.

Jouvet, G. and Rappaz, J.: Analysis and Finite Element Approximation of a Nonlinear Stationary Stokes Problem Arising in Glaciology, Advances in Numerical Analysis, 2011, 1–24, https://doi.org/10.1155/2011/164581, 2011.

Jouvet, G., Huss, M., Funk, M., and Blatter, H.: Modelling the retreat of Grosser Aletschgletscher, Switzerland, in a changing climate, Journal of Glaciology, 57, 1033–1045, 2011.

Jouvet, G., Cordonnier, G., Kim, B., Lüthi, M., Vieli, A., and Aschwanden, A.: Deep learning speeds up ice flow modelling by several orders of magnitude, Journal of Glaciology, 68, 651–664, https://doi.org/10.1017/jog.2021.120, 2021.

Kamb, W. B.: Sliding Motion of Glaciers: Theory and Observation, Reviews of Geophysics and Space Physics, 8, 673–728, https://doi.org/10.1029/RG008i004p00673, 1970.

Kashi, A. and Nadarajah, S.: On the effectiveness of fine-grain parallel linear iterations for computational aerodynamics on structured grids for graphics processing units, Computers & Fluids, 299, 106 714, https://doi.org/10.1016/j.compfluid.2025.106714, 2025.

Khroulev, C. and the PISM Authors: PISM, a Parallel Ice Sheet Model v1.2: User's Manual, User's manual, University of Alaska Fairbanks, Fairbanks, Alaska, https://www.pism-docs.org, 2020.

Kingma, D. P. and Ba, J.: Adam: A Method for Stochastic Optimization, in: International Conference on Learning Representations (ICLR), 2015.

Kleiner, T., Rückamp, M., Bondzio, J. H., and Humbert, A.: Enthalpy benchmark experiments for numerical ice sheet models, The Cryosphere, 9, 217–228, https://doi.org/10.5194/tc-9-217-2015, 2015.

Koo, Y. and Rahnemoonfar, M.: Graph convolutional network as a fast statistical emulator for numerical ice sheet modeling, Journal of Glaciology, 71, https://doi.org/10.1017/jog.2024.93, 2024.

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

Leger, T. P. M., Jouvet, G., Kamleitner, S., Mey, J., Herman, F., Finley, B. D., Ivy-Ochs, S., Vieli, A., Henz, A., and Nussbaumer, S. U.: A data-consistent model of the last glaciation in the Alps achieved with physics-driven AI, Nature Communications, 16, https://doi.org/10.1038/s41467-025-56168-3, 2025.

Leger, T. P. M., Jouvet, G., Kamleitner, S., Finley, B. D., Bernard, M., Allegri, B., Herman, F., Vieli, A., Henz, A., and Nussbaumer, S. U.: First Alps-wide reconstruction of LGM glacial sediment transport enabled by GPU-accelerated particle tracking, Earth Surface Dynamics, 14, 361–389, https://doi.org/10.5194/esurf-14-361-2026, 2026.

Levermann, A., Albrecht, T., Winkelmann, R., Martin, M. A., Haseloff, M., and Joughin, I.: Kinematic first-order calving law implies potential for abrupt ice-shelf retreat, The Cryosphere, 6, 273–286, https://doi.org/10.5194/tc-6-273-2012, 2012.

Li, Z., Kovachki, N., Azizzadenesheli, K., Liu, B., Bhattacharya, K., Stuart, A., and Anandkumar, A.: Fourier Neural Operator for Parametric Partial Differential Equations, in: International Conference on Learning Representations, https://openreview.net/forum?id=c8P9NQVtmnO, 2021.

Licciulli, C., Bohleber, P., Lier, J., Gagliardini, O., Hoelzle, M., and Eisen, O.: A full Stokes ice-flow model to assist the interpretation of millennial-scale ice cores at the high-Alpine drilling site Colle Gnifetti, Swiss/Italian Alps, Journal of Glaciology, 66, 35–48, 2020.

Lipscomb, W. H., Price, S. F., Hoffman, M. J., Leguy, G. R., Bennett, A. R., Bradley, S. L., Evans, K. J., Fyke, J. G., Kennedy, J. H., Perego, M., Ranken, D. M., Sacks, W. J., Salinger, A. G., Vargo, L. J., and Worley, P. H.: Description and evaluation of the Community Ice Sheet Model (CISM) v2.1, Geoscientific Model Development, 12, 387–424, https://doi.org/10.5194/gmd-12-387-2019, 2019.

Liu, D. C. and Nocedal, J.: On the limited memory BFGS method for large scale optimization, Mathematical programming, 45, 503–528, 1989.

Lliboutry, L. and Duval, P.: Various isotropic and anisotropic ices found in glaciers and polar ice caps and their corresponding rheologies, Annales Geophysicae, 3, 207–224, https://doi.org/10.1016/0148-9062(85)90267-0, 1985.

Maussion, F., Butenko, A., Champollion, N., Dusch, M., Eis, J., Fourteau, K., Gregor, P., Jarosch, A. H., Landmann, J., Oesterle, F., Recinos, B., Rothenpieler, T., Vlug, A., Wild, C. T., and Marzeion, B.: The Open Global Glacier Model (OGGM) v1.1, Geoscientific Model Development, 12, 909–931, https://doi.org/10.5194/gmd-12-909-2019, 2019.

Meier, M. F., Tangborn, W. V., Mayo, L. R., and Post, A.: Combined ice and water balances of Gulkana and Wolverine Glaciers, Alaska, and South Cascade Glacier, Washington, 1965 and 1966 hydrologic years, Tech. rep., US Geological Survey, 1971.

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

Müller, J. and Zeinhofer, M.: Achieving High Accuracy with PINNs via Energy Natural Gradient Descent, in: Proceedings of the 40th International Conference on Machine Learning, edited by Krause, A., Brunskill, E., Cho, K., Engelhardt, B., Sabato, S., and Scarlett, J., vol. 202 of *Proceedings of Machine Learning Research*, pp. 25 471–25 485, PMLR, 2023.

Müller, J. and Zeinhofer, M.: Position: Optimization in SciML Should Employ the Function Space Geometry, in: Proceedings of the 41st International Conference on Machine Learning, edited by Salakhutdinov, R., Kolter, Z., Heller, K., Weller, A., Oliver, N., Scarlett, J., and Berkenkamp, F., vol. 235 of *Proceedings of Machine Learning Research*, pp. 36 705–36 722, PMLR, 2024.

Nye, J. F.: A Calculation on the Sliding of Ice over a Wavy Surface using a Newtonian Viscous Approximation, Proceedings of the Royal Society of London. Series A, Mathematical and Physical Sciences, 311, 445–467, https://doi.org/10.1098/rspa.1969.0127, 1969.

Ozaki, Y., Watanabe, S., and Yanase, T.: OptunaHub: A Platform for Black-Box Optimization, arXiv preprint arXiv:2510.02798, 2025.

Paterson, W. S. B. and Budd, W. F.: Flow parameters for ice sheet modeling, Cold Regions Science and Technology, 6, 175–177, https://doi.org/10.1016/0165-232X(82)90010-6, 1982.

Pattyn, F.: A new three-dimensional higher-order thermomechanical ice sheet model: Basic sensitivity, ice stream development, and ice flow across subglacial lakes, Journal of Geophysical Research: Solid Earth, 108, https://doi.org/10.1029/2002jb002329, 2003.

Pattyn, F., Perichon, L., Aschwanden, A., Breuer, B., de Smedt, B., Gagliardini, O., Gudmundsson, G. H., Hindmarsh, R. C. A., Hubbard, A., Johnson, J. V., Kleiner, T., Konovalov, Y., Martin, C., Payne, A. J., Pollard, D., Price, S., Rückamp, M., Saito, F., Souček, O., Sugiyama, S., and Zwinger, T.: Benchmark experiments for higher-order and full-Stokes ice sheet models (ISMIP–HOM), The Cryosphere, 2, 95–108, https://doi.org/10.5194/tc-2-95-2008, 2008.

Pollard, D. and DeConto, R. M.: A simple inverse method for the distribution of basal sliding coefficients under ice sheets, applied to Antarctica, The Cryosphere, 6, 953–971, https://doi.org/10.5194/tc-6-953-2012, 2012.

Prasad, V., Herrmann, O., Tabone, I., K C, M., Groos, A. R., Jouvet, G., Jordan, J. R., and Fürst, J. J.: A numerical framework for modeling iceberg calving and ice-front migration of grounded glacier tongues, EGUsphere, 2026, 1–21, https://doi.org/10.5194/egusphere-2026-508, 2026.

Räss, L., Licul, A., Herman, F., Podladchikov, Y. Y., and Suckale, J.: Modelling thermomechanical ice deformation using an implicit pseudo-transient method (FastICE v1.0) based on graphical processing units (GPUs), Geoscientific Model Development, 13, 955–976, https://doi.org/10.5194/gmd-13-955-2020, 2020.

Riel, B., Minchew, B., and Bischoff, T.: Data-driven inference of the mechanics of slip along glacier beds using physics-informed neural networks: Case study on Rutford ice stream, Antarctica, Journal of Advances in Modeling Earth Systems, 13, e2021MS002 621, 2021.

Roe, P. L.: Characteristic-based schemes for the Euler equations, Annual Review of Fluid Mechanics, 18, 337–365, https://doi.org/10.1146/annurev.fl.18.010186.002005, 1986.

Rosier, S. H. R., Bull, C. Y. S., Woo, W. L., and Gudmundsson, G. H.: Predicting ocean-induced ice-shelf melt rates using deep learning, The Cryosphere, 17, 499–518, https://doi.org/10.5194/tc-17-499-2023, 2023.

Rounce, D. R., Hock, R., and Shean, D. E.: Glacier Mass Change in High Mountain Asia Through 2100 Using the Open-Source Python Glacier Evolution Model (PyGEM), Frontiers in Earth Science, 7, https://doi.org/10.3389/feart.2019.00331, 2020.

Rounce, D. R., Hock, R., Maussion, F., Hugonnet, R., Kochtitzky, W., Huss, M., Berthier, E., Bolch, T., Compagno, L., Copland, L., et al.: Global Glacier Change in the 21st Century: Every Increase in Temperature Matters, Science, 379, 78–83, https://doi.org/10.1126/science.abo1324, 2023.

Rowan, A. V., Egholm, D. L., and Clark, C. D.: Forward modelling of the completeness and preservation of palaeoclimate signals recorded by ice-marginal moraines, Earth Surface Processes and Landforms, 47, 2198–2208, 2022.

Scherler, D. and Egholm, D. L.: Production and Transport of Supraglacial Debris: Insights From Cosmogenic ^{10}Be and Numerical Modeling, Chhota Shigri Glacier, Indian Himalaya, Journal of Geophysical Research: Earth Surface, 125, e2020JF005 586, https://doi.org/10.1029/2020JF005586, 2020.

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

Schoof, C.: Coulomb Friction and Other Sliding Laws in a Higher-Order Glacier Flow Model, Mathematical Models and Methods in Applied Sciences, 20, 157–189, https://doi.org/10.1142/s0218202510004180, 2010.

Schoof, C. and Hindmarsh, R. C. A.: Thin-Film Flows with Wall Slip: An Asymptotic Analysis of Higher Order Glacier Flow Models, The Quarterly Journal of Mechanics and Applied Mathematics, 63, 73–114, https://doi.org/10.1093/qjmam/hbp025, 2010.

Shapero, D. R., Badgeley, J. A., Hoffman, A. O., and Joughin, I. R.: icepack: a new glacier flow modeling package in Python, version 1.0, Geoscientific Model Development, 14, 4593–4616, https://doi.org/10.5194/gmd-14-4593-2021, 2021.

Tezaur, I. K., Perego, M., Salinger, A. G., Tuminaro, R. S., and Price, S. F.: Albany/FELIX: a parallel, scalable and robust, finite element, first-order Stokes approximation ice sheet solver built for advanced analysis, Geoscientific Model Development, 8, 1197–1220, https://doi.org/10.5194/gmd-8-1197-2015, 2015.

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

Van der Veen, C. J.: Tidewater calving, Journal of Glaciology, 42, 375–385, https://doi.org/10.3189/S0022143000004226, 1996.

Verjans, V. and Robel, A.: Accelerating Subglacial Hydrology for Ice Sheet Models With Deep Learning Methods, Geophysical Research Letters, 51, https://doi.org/10.1029/2023gl105281, 2024.

Vieli, A., Funk, M., and Blatter, H.: Flow dynamics of tidewater glaciers: a numerical modelling approach, Journal of Glaciology, 47, 595–606, https://doi.org/10.3189/172756501781831747, 2001.

Višnjević, V., Herman, F., and Prasicek, G.: Climatic patterns over the European Alps during the LGM derived from inversion of the paleo-ice extent, Earth and Planetary Science Letters, 538, 116 185, https://doi.org/10.1016/j.epsl.2020.116185, 2020.

Vyas, N., Morwani, D., Zhao, R., Kwun, M., Shapira, I., Brandfonbrener, D., Janson, L., and Kakade, S.: SOAP: Improving and Stabilizing Shampoo using Adam, https://doi.org/10.48550/ARXIV.2409.11321, 2024.

Wang, Y., Zhang, T., Xiao, C., Ren, J., and Wang, Y.: A two-dimensional, higher-order, enthalpy-based thermomechanical ice flow model for mountain glaciers and its benchmark experiments, Computers & Geosciences, 141, 104 526, https://doi.org/10.1016/j.cageo.2020.104526, 2020.

Wang, Y., Lai, C.-Y., Prior, D. J., and Cowen-Breen, C.: Deep learning the flow law of Antarctic ice shelves, Science, 387, 1219–1224, https://doi.org/10.1126/science.adp3300, 2025.

Watkins, J., Carlson, M., Shan, K., Tezaur, I., Perego, M., Bertagna, L., Kao, C., Hoffman, M. J., and Price, S. F.: Performance portable ice-sheet modeling with MALI, The International Journal of High Performance Computing Applications, 37, 600–625, https://doi.org/10.1177/10943420231183688, 2023.

Weertman, J.: On the Sliding of Glaciers, Journal of Glaciology, 3, 33–38, https://doi.org/10.3189/S0022143000024709, 1957.

Winkelmann, R., Martin, M. A., Haseloff, M., Albrecht, T., Bueler, E., Khroulev, C., and Levermann, A.: The Potsdam Parallel Ice Sheet Model (PISM-PIK) – Part 1: Model description, The Cryosphere, 5, 715–726, https://doi.org/10.5194/tc-5-715-2011, 2011.

Yadan, O.: Hydra - A framework for elegantly configuring complex applications, Github, 2019.

Zekollari, H., Huss, M., Farinotti, D., and Lhermitte, S.: Ice-Dynamical Glacier Evolution Modeling—A Review, Reviews of Geophysics, 60, https://doi.org/10.1029/2021rg000754, 2022.