https://doi.org/10.5194/egusphere-2026-4761
Preprint. Discussion started: 18 August 2026

An extensible, algorithmic framework to enumerate, ascertain, and compare mechanical properties from common snowpit observations

Mary Kate Connelly[1], Philipp L. Rosendahl[2], Valentin Adam[2,3], and Samuel V. Verplanck[4]

[1]Montana State University, College of Letters and Science, Department of Mathematical Sciences, Bozeman, MT, USA
[2]Technical University of Darmstadt, Institute of Structural Mechanics and Design, Darmstadt, Germany
[3]WSL Institute for Snow and Avalanche Research SLF, Davos Dorf, Switzerland
[4]Montana State University, Norm Asbjornson College of Engineering, Department of Civil Engineering, Bozeman, MT, USA

Correspondence: Philipp Rosendahl (rosendahl@ismd.tu-darmstadt.de)

Abstract. Mechanical properties of snow, such as Young's modulus and Poisson's ratio, are used in avalanche release models but rarely measured directly in snowpits. Researchers chain published parameterizations to estimate these properties, but the combinations and their practical coverage have not been systematically compared. We present SnowPyt-MechParams, an extensible, graph-based framework that encodes parameterizations as a directed acyclic graph, enumerates valid pathways, and propagates measurement uncertainty through each one. For an example dataset, we developed SnowPylot to parse CAAML files from the SnowPilot database, structuring 50,278 snowpits (371,429 layers, 2015–2025 water years) into 14,776 slabs from Extended Column Test results with propagation. As an example application, we implemented pathways for the shear component of slab weight per unit area with and without elastic properties. The highest-coverage shear-weight pathway succeeds for 5,470 slabs (37.0 %), while the highest-coverage shear-weight-with-elasticity pathway succeeds for only 687 slabs (4.6 %), reflecting elasticity's longer calculation chains. Slab bending stiffness, D_{11}, carries a median relative uncertainty of 78 % (IQR: 56–85 %). Because multiple pathways can reach the same target, estimates vary substantially, with a median max/min ratio for D_{11} of 32 (up to 63) per slab. These results demonstrate the framework's utility to systematically calculate coverage, uncertainty, and range in mechanical properties from snowpit observations using different parameterizations.

1 Introduction and background

Recently, modeling the mechanical behavior of snow relevant to avalanche release has advanced significantly (Weißgraeber and Rosendahl, 2023; Siron et al., 2023; Bergfeld et al., 2023). These models rely on input parameters that range from readily measured snowpit observations, such as slope angle and layer thickness, to more difficult-to-obtain mechanical properties, such as Young's modulus and Poisson's ratio. To bridge this gap, researchers use parameterizations from the literature to relate observable snowpack characteristics to these mechanical properties.

Viallon-Galinier et al. (2022) reviewed a range of models used to calculate snowpack stability from simulated snow stratigraphy (Morin et al., 2020), demonstrating that stability predictions are sensitive to inputs from snow cover models. They also emphasized that multiple parameterizations often exist for a single mechanical property. In some cases, these relationships are sequential, with one parameter required to estimate another, resulting in chains of parameterizations.

https://doi.org/10.5194/egusphere-2026-4761
Preprint. Discussion started: 18 August 2026

The objective of our work is to develop an extensible, algorithmic framework to enumerate, ascertain, and compare mechanical properties from snowpit observations. We apply this framework to a dataset of 50,278 snowpits and present example calculations of the shear component of slab weight per unit area with and without elastic properties.

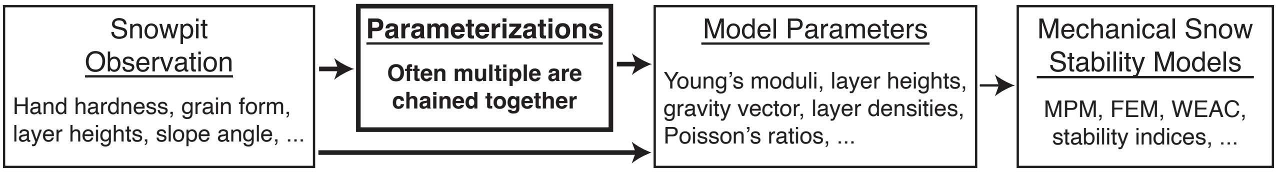

Figure 1. Mechanical models of snow stability employ numerical methods such as the Material Point Method (MPM) and the Finite Element Method (FEM), as well as analytical approaches like the Weak Layer Anticrack Nucleation Model (WEAC); all of these models are sensitive to input parameters. The objective of our work is to develop an extensible, algorithmic framework to connect snowpit observations to model input parameters.

1.1 Snowpits and the SnowPilot database

A snowpit is a hole dug into the snowpack to a depth of interest, which may be all the way to the ground. The observer records information about the snowpack (e.g. layering, grain forms, densities), terrain characteristics (e.g. slope angle, aspect, elevation), weather conditions (e.g. cloud cover, precipitation rate), and stability test results. There are a variety of guiding documents that specify standard techniques and protocols to conduct a snowpit (Fierz et al., 2009; Greene et al., 2022; CAA Technical Committee, 2024; Norwegian Water Resources and Energy Directorate, 2022). Snowpit observations are fundamental to avalanche forecasting because they provide critical information about snowpack structure and stability (LaChapelle, 1980).

In this study, we utilize Extended Column Test (ECT) results (Simenhois and Birkeland, 2006). The ECT is a popular stability test that involves isolating a parallelepiped of snow, 30 cm upslope, 90 cm cross-slope, and loading the surface of the isolated block by striking a shovel blade with a gloved hand (Toft et al., 2024). If failure occurs in a weak layer, then the depth of failure, the tap number, and whether the crack propagates across the column are recorded.

SnowPilot (snowpilot.org) is a free, open-source software designed to help users record, visualize, and store snowpit observations (Chabot et al., 2004, 2016). The SnowPilot database contains observations of over 65,000 snowpits, collected by snow recreationists and professionals around the world. It is particularly popular among avalanche professionals in the United States. The dataset used in this analysis (Tables 1 and 2) comprises a decade of snowpit data exported from SnowPilot, spanning the 2015-2025 water years (September 1, 2014 to August 31, 2025). The scale and nature of the SnowPilot dataset makes it uniquely valuable in the context of investigating the realistic distribution of field-observed values, and the completeness of field observations.

https://doi.org/10.5194/egusphere-2026-4761
Preprint. Discussion started: 18 August 2026

Table 1. A selection of summary statistics for the SnowPilot database used in this study. The dataset is comprised of 10 water years, from September 1, 2014 to August 31, 2025.

	Count
Snowpits	
Total	50,278
Unique users	5,381
Unique countries	35
With slope angle observation	45,515 (90.5 %)
Recorded near an avalanche	1,568 (3.1 %)
Snow layers	
Total	371,429
With hand hardness	336,888 (90.7 %)
With primary grain form	303,726 (81.8 %)
With grain size	176,044 (47.4 %)
With direct density measurement[a]	10,468 (2.8 %)
Stability tests	
Compression tests (CT)	51,599
Extended column tests (ECT)	47,684
Propagation saw tests (PST)	6,213

[a] SnowPilot allows density to be entered either as an attribute of a snow profile layer or as a separate density profile. Our analysis uses only density observations whose depth and thickness match a snow profile layer boundary; the count reported here reflects this aligned subset.

Table 2. Primary grain form counts for SnowPilot layers, sorted by frequency. Layers with secondary grain form codes are assigned to their parent class.

Grain form	Layers
Faceted crystals (FC)	79,518
Rounded grains (RG)	64,884
Melt forms (MF)	55,416
Decomposing and fragmented precipitation particles (DF)	43,427
Precipitation particles (PP)	21,526
Ice formations (IF)	19,136
Depth hoar (DH)	12,379
Surface hoar (SH)	7,417
Machine-made snow (MM)	23

https://doi.org/10.5194/egusphere-2026-4761
Preprint. Discussion started: 18 August 2026

1.2 Mechanical slab properties and avalanche release

Dry snow slab avalanche release is a phenomenon that occurs when a "weak layer" layer of snow fails and releases the overlying "slab" from the slope (Schweizer et al., 2016). Mechanical models of snow stability span a large range in complexity, underlying theory, and number of input parameters. A common theme among them is an input of slab properties, hence, this is the focus of our example implementation of the algorithmic framework.

Slab weight per unit area, W, is calculated using density, ρ_n, and thickness, h_n, for each layer, n, of N total layers comprising a defined slab.

$$W = \sum_{n=1}^{N} \rho_n g h_n, \tag{1}$$

where g is the acceleration due to gravity. The weight per unit area can be decomposed into a normal (slope-perpendicular) component, W_{normal}, and shear (slope-parallel) component, W_{shear}, using the slope angle ψ.

$$W_{normal} = W \cos\psi. \tag{2}$$

$$W_{shear} = W \sin\psi. \tag{3}$$

The slab weight imposes a static stress on the underlying weak layer. Stress-based stability criteria predict weak layer failure when an applied stress exceeds the weak layer strength. Roch (1965) utilized a failure criteria solely based on shear strength, while more recent efforts include mixed-mode failure envelope with shear and compressive components of strength (Reiweger et al., 2015).

Energy-based formulations (Heierli et al., 2008) and combined energy and stress models, such as WEAC (Rosendahl and Weißgraeber, 2020b, a; Weißgraeber and Rosendahl, 2023), consider slab elasticity, in addition to slab weight for stability criterion.

In an elastic material, stress is linearly related to strain by generalized Hooke's Law:

$$\sigma_{ij} = c_{ijkl}\epsilon_{kl} \tag{4}$$

where σ_{ij} is the rank 2 stress tensor, ϵ_{kl} is the rank 2 strain tensor, and c_{ijkl} is the rank 4 stiffness tensor. Note that indicial notation is used here, and the indices i, j, k, and l are each comprised of [1,2,3] in three-dimensional space. Thus, the stiffness tensor is comprised of 81 coefficients, which can be reduced to 21 independent elastic constants by utilizing symmetry of the stress and strain tensors. In the case of an isotropic, elastic material, the stiffness tensor is further reduced to two unique values, Young's modulus, E, and Poisson's ratio, ν.

https://doi.org/10.5194/egusphere-2026-4761
Preprint. Discussion started: 18 August 2026

$$\sigma_{ij} = \frac{E}{1+\nu}(\epsilon_{ij} + \frac{\nu}{1-2\nu}\delta_{ij}\epsilon_{kk}) \tag{5}$$

where δ_{ij} is the Kronecker delta. An isotropic, elastic material is characterized by two, material-dependent parameters. Often these two parameters are presented as Young's modulus and Poisson's ratio, however, they can be also presented as bulk modulus, Lamé's first parameter, P-wave modulus, or shear modulus. For example, shear modulus, G, can be calculated from Young's modulus and Poisson's ratio using

$$G = \frac{E}{2(1+\nu)}. \tag{6}$$

Weißgraeber and Rosendahl (2023) provide closed-form, analytic formulations for extensional stiffness, A_{11}, bending stiffness, D_{11}, bending-extension coupling stiffness B_{11} and shear stiffness κA_{55} for a layered slab.

$$A_{11} = \sum_{n=1}^{N} \frac{E_n}{1-\nu_n^2} h_n, \tag{7}$$

$$B_{11} = \frac{1}{2}\sum_{n=1}^{N} \frac{E_n}{1-\nu_n^2}(z_{n+1}^2 - z_n^2), \tag{8}$$

$$D_{11} = \frac{1}{3}\sum_{n=1}^{N} \frac{E_n}{1-\nu_n^2}(z_{n+1}^3 - z_n^3), \tag{9}$$

and

$$A_{55} = \sum_{n=1}^{N} G_n h_n, \tag{10}$$

where κ is shear correction factor, set to 5/6. The top and bottom, slope-perpendicular coordinates are z_n and z_{n+1}, respectively.

1.3 Objectives and scope

The objective of our work is to develop an extensible framework for enumerating, ascertaining, and comparing the mechanical properties of snow from common snowpit observations. Accordingly, we do not attempt to implement an exhaustive set of mechanical parameterizations across constitutive models, nor do we incorporate weak layer failure criteria. Instead, we present representative calculations of the shear component slab weight per unit area, W_{shear}, with and without slab elasticity properties assuming each layer is isotropic and elastic. Our focus is on slab-side parameterizations that can be derived from standard snowpit observations. Full implementation of snow stability models is intentionally deferred to emphasize the general utility and adaptability of the presented framework

https://doi.org/10.5194/egusphere-2026-4761
Preprint. Discussion started: 18 August 2026

2 Methods

2.1 From SnowPilot data to Python objects

Despite its widespread use for recording and sharing snowpit data, the SnowPilot database has been underutilized as a research resource, partly due to its structural complexity and the lack of tools for efficiently extracting specific snowpack characteristics. To address these challenges and improve accessibility to the dataset, we developed SnowPylot, an open-source Python library to import and structure data from the SnowPilot dataset within Python. The library is designed to work with files exported in the Canadian Avalanche Association Markup Language (CAAML) format (CAA Technical Committee, 2024). CAAML.xml files are exported from SnowPilot and parsed into structured Python objects using SnowPylot.

The SnowPylot library implements an object-oriented design that mirrors the CAAML data structure. The `snowpit` class serves as the container for all snowpit observations and consists of three main components relevant to this analysis: `core_info`, `snow_profile`, and `stability_tests`. Each CAAML.xml file in the dataset is parsed into a `snowpit` object. Fig. 2 summarizes the subset of `snowpit` components used in this study to assemble `layer` objects, match direct density observations, and identify the slabs of ECT's that resulted in propagation.

The `core_info` class encapsulates observation metadata and pit-level attributes, including the pit's identification details, observer information, details about the time and location of the observation, and weather conditions.

The `snow_profile` class contains basic profile information like measurement direction, profile depth, and height of the snowpack. The class also contains information about the snowpit layers (stored in a list of `layer` objects), temperature observations (stored in a list of `temp_obs` objects), and density observations (stored in a list of `density_obs` objects). Each `layer` object contains information about the layer's thickness, hand hardness, grain form, grain size, and density.

The `stability_tests` class serves as a container for different types of snow stability tests. It contains lists as attributes, each representing a different type of stability test (e.g. Extended Column Test, Compression Test, and Propagation Saw Test). Each `stability_tests` object contains information about the test's score, failure depth, and other relevant details.

SnowPilot allows density to be entered either as an attribute of an individual layer in the snow profile, or as an observation in a separate density profile. When entered as a separate profile, a density observation cannot always be matched to a specific layer in the snow profile if the depth and thickness of the layers in the density profile do not align with the layers in the snow profile. SnowPylot allows the user to access the entire density profile, but the analysis in this study includes only density observations whose depth and thickness match a snow profile layer exactly, density observations that cannot be matched are excluded.

Additionally, SnowPilot allows the user to enter both the primary grain form and secondary grain form of a snowpit layer. SnowPylot parses this information as recorded by the user, but our example implementation uses only the primary grain form for all parameterization calculations. When we refer to the grain form of a layer, we are referring to the primary grain form of that layer as recorded by the user.

Our investigation into slab properties utilizes the snow profile, density profile, stability test, and slope angle information from SnowPilot profiles. SnowPylot was used to access and connect the data.

https://doi.org/10.5194/egusphere-2026-4761
Preprint. Discussion started: 18 August 2026

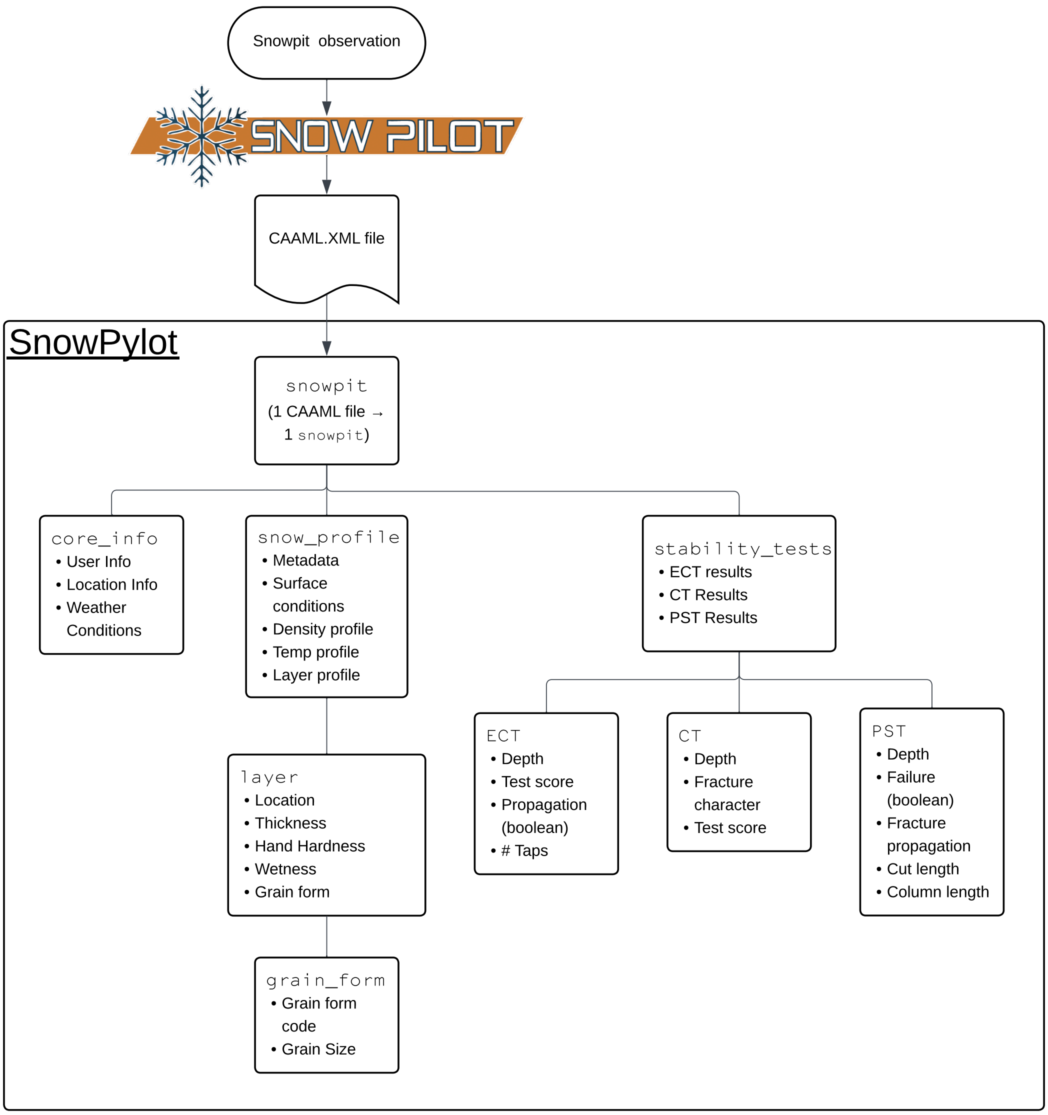

Figure 2. SnowPylot object hierarchy used in this analysis. Each CAAML.xml file exported from SnowPilot is parsed into a `snowpit` object with three main components: `core_info` (pit identification, observer, location, and weather metadata), `snow_profile` (layer records with grain form, grain size, hand hardness, and thickness; plus separate density and temperature observation lists), and `stability_tests` (ECT, CT, and Rutschblock result lists). Grain size and grain form are stored within each `layer` object as a nested `grain_form` sub-object.

https://doi.org/10.5194/egusphere-2026-4761
Preprint. Discussion started: 18 August 2026

2.2 Uncertainty estimation and propagation

Every field observation carries associated uncertainty due to observation method and spatial variability within the snowpit, we consider this to be measurement uncertainty. SnowPilot does not have the capability for users to record uncertainty to their observations, yet, uncertainty is inherent in any observation.

In our framework (SnowPyt-MechParams), uncertain values are represented with the Python uncertainties package (Lebigot, 2010) as `ufloat` objects, that contain nominal values with associated uncertainty. The package propagates uncertainty automatically through the chained calculations of a first-order Taylor expansion.

For example implementation, we estimated reasonable values for snowpit observation uncertainty based on available literature and observational experience (Birkeland, 2025).

For density, we assume directly observed values to have an uncertainty of $\pm 10\%$ based on Conger and McClung (2009) who estimated $\pm 11\%$ and Proksch et al. (2016) who estimated $\pm 9\%$.

For slope angle, McCammon (2023) observed users making profile measurements of slope angle have an accuracy of $\pm 4°$ across a variety of inclinometers commonly used by avalanche practitioners. In his study, participants were estimating the slope of lines projected on a screen. In a snowpit, slope angle observations are often made by laying an inclinometer directly on the slope, a method expected to have less uncertainty, thus we assume $\pm 2°$.

Höller and Fromm (2010) provide perhaps the most detailed study of the hand hardness observation, yet they do not include a recommended uncertainty for hand hardness index within a snowpit. The primary hand hardness classes of gloved fist (F), gloved four fingers (4F), gloved one figner (1F), blunt end of pencil (P), knife blade (K), and ice (I) are sub-classified into 16 categories of F-, F, F+, 4F-, 4F, 4F+, and so on. In our study we assume an uncertainty of $\pm$ two subclasses, meaning, a nominal observation of 4F- would range from F to 4F+. Grain size is assumed to be ± 0.5 mm and layer thickness is assumed to be $\pm 5\,\%$.

In addition to measurement uncertainty, there is uncertainty associated with each parameterization, referred to as parameterization uncertainty. These uncertainties are derived from published regressions or reported between-sample variability. Not every implemented parameterization has available or comparable information needed to propagate parameterization uncertainty. To enable consistent comparison across pathways, only measurement uncertainty was propagated in the results reported here and parameterization uncertainty was excluded. The framework is coded to optionally include parameterization uncertainty, and the relevant values are documented in the code within each method's implementation wherever available.

2.3 Example parameterizations

Example parameterizations are implemented in the framework to ascertain density, Young's modulus, and Poisson's ratio for each layer in a snowpit. We limit our implementation to methods proposed in the original studies and do not develop new parameterizations from published datasets. Our goal is not to present an exhaustive review of published mechanical parameterizations, but instead to include a representative set of methods that are recent, widely cited, and/or can be implemented from commonly recorded snowpit observations.

https://doi.org/10.5194/egusphere-2026-4761
Preprint. Discussion started: 18 August 2026

2.3.1 Density

Geldsetzer and Jamieson (2000) provide an empirically-derived estimate of snow density based on an observed grain form and hand hardness. Kim and Jamieson (2014) then extended the work of Geldsetzer and Jamieson (2000) to update the regressions with larger sample sizes and a wider range of grain forms and hand hardness values. Furthermore, Kim and Jamieson (2014) developed alternate regressions that include grain size as an additional regressor.

Machine made snow, surface hoar, and ice formations are the primary grain forms (Fierz et al., 2009) not included in Kim and Jamieson (2014) nor Geldsetzer and Jamieson (2000). The implemented density method support domains across grain form and hand hardness are shown in Fig. 3. We use the short labels "Geldsetzer" for Geldsetzer and Jamieson (2000) and "Kim T2" and "Kim T6" for the Kim and Jamieson (2014) Table 2 regression and the Equation 5 regression with Table 6 coefficients, respectively, in figures and compact pathway labels. Melt forms (MF) are included in the latter method because they are included in the Table 6 multivariable density regression.

Others have estimated snow density from measurements such as the Snow MicroPenetrometer (King et al., 2020), however, since it is not a common snow pit observation, it is not included in our analysis.

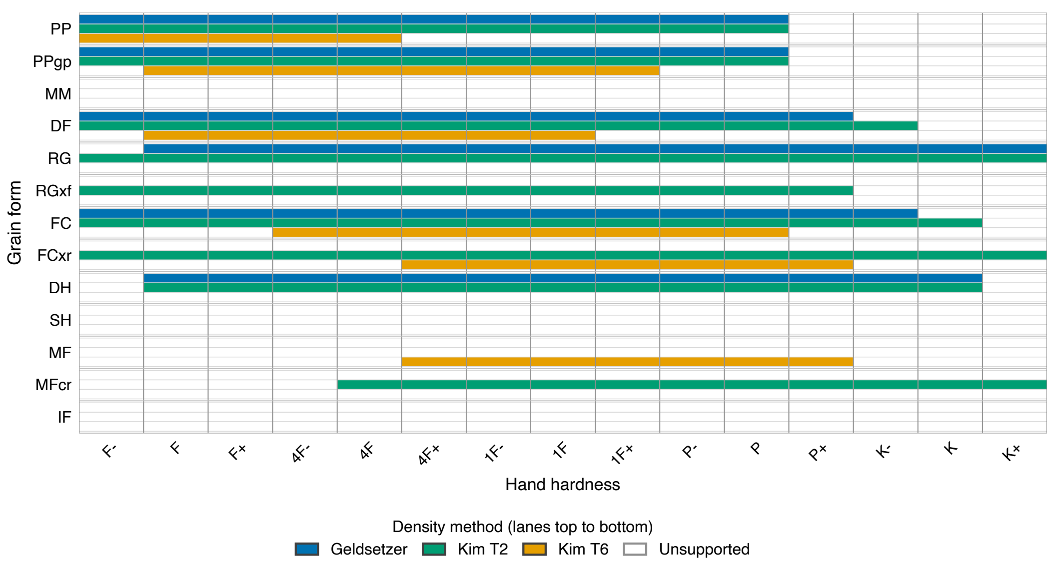

Figure 3. Implemented density method support as a function of primary grain form and hand hardness. Colored lanes indicate supported domains for Geldsetzer, Kim T2, and Kim T6; white cells indicate unsupported combinations.

2.3.2 Effective, isotropic Young's modulus

Although the anisotropic, inelastic behavior of snow has long been recognized; models often assume isotropic, elastic behavior of the slab. Thus an input into such models is an effective, isotropic Young's modulus for each snow layer.

https://doi.org/10.5194/egusphere-2026-4761
Preprint. Discussion started: 18 August 2026

Köchle and Schneebeli (2014) used X-ray micro-computed tomography (μ-CT) scans of rounded grains, faceted crystals, melt forms, and depth hoar to provide the geometry input into finite element analysis (FEA) of simulated elastic behavior to determine density-based parameterizations of Young's modulus. They suggested two different empirical relationships with density.

$$E_{n,Köchle-low} = 0.0061e^{0.0396\rho_n}, \tag{11}$$

for density values of 150-250 kg m^{-3}.

$$E_{n,Köchle-high} = 6.0457e^{0.011\rho_n}, \tag{12}$$

for density values of 250-450 kg m^{-3}. The formulations in Eqs. 11 and 12 also assume input densities in kg m^{-3} and output Young's moduli in MPa.

Wautier et al. (2015) determined orthotropic elastic stiffness tensor values based on FEA of μ-CT scans. A wide range of grain forms were studied (Table 3). They found normalized averages of Young's moduli

$$\frac{E_{n,Wautier}}{E_{ice}} = 0.78(\frac{\rho_n}{\rho_{ice}})^{2.34} \tag{13}$$

valid for snow densities from 103 to 544 kg m^{-3} utilizing the Young's modulus of ice, E_{ice}, and density of ice, ρ_{ice}. We assume E_{ice} to be 10 GPa (Nanthikesan and Sunder, 1994) and ρ_{ice} to be 917 kg m^{-3}.

Gerling et al. (2017) compared three methods of ascertaining Young's moduli of rounded grains in the laboratory: acoustic wave propagation, FEA of μ-CT scans, and Snow MicroPenetrometer (SMP) measurements. The acoustic wave and FEA methods calculated a P-wave modulus, and the SMP method calculated a Young's modulus. They then developed density-based parameterizations for these moduli which Bergfeld et al. (2023) utilized in the analysis of slab bending during PST experiments. Bergfeld et al. (2023) then proposed the parameterization

$$E_{n,Bergeld} = C_0(\frac{\rho_n}{\rho_{ice}})^{C_1} \tag{14}$$

where $C_0 = 6.5 \times 10^3$ MPa and $C_1 = 4.4$. Although Gerling et al. (2017) was limited to rounded grains, Bergfeld et al. (2023) included precipitation particles, decomposing and fragmented, and rounded grains.

Lastly, Schöttner et al. (2026b) developed three groups of density parameterizations of Young's moduli based on stress-strain data from laboratory testing and Digital Image Correlation (DIC).

$$\frac{E_{n,Schottner,DF\&RG}}{E_{ice}} = 0.4(\frac{\rho_n}{\rho_{ice}})^{4.6} \tag{15}$$

https://doi.org/10.5194/egusphere-2026-4761
Preprint. Discussion started: 18 August 2026

$$\frac{E_{n,Schottner,FC\&DH}}{E_{ice}} = 1.8(\frac{\rho_n}{\rho_{ice}})^{5.1} \tag{16}$$

$$\frac{E_{n,Schottner,SH}}{E_{ice}} = 0.011(\frac{\rho_n}{\rho_{ice}})^{1.7} \tag{17}$$

2.3.3 Effective, isotropic Poisson's ratio

Poisson's ratio can be estimated from grain form according to Köchle and Schneebeli (2014), where rounded grains have a Poisson's ratio of 0.171 ± 0.026, faceted crystals have 0.130 ± 0.040, and depth hoar has 0.087 ± 0.063. Although melt-freeze grain forms were included in the study, the authors did not provide mean values and uncertainties for those grains, so they are excluded from our example framework implementation. Although only a finite range of densities was tested in their study, the authors concluded that Poisson's ratio showed no correlation with density.

Srivastava et al. (2016) utilized μ-CT scans to build finite element models and determine orthotropic stiffness tensors for a variety of grain forms. From these tensors, effective Poisson's ratios were derived based on the Voigt and Reuss bounds. They found rounded grains to be 0.191 $\pm$ 0.008, precipitation particles and decomposing and fragmented particles to be 0.132 $\pm$ 0.053, and faceted crystals and depth hoar to be 0.17 $\pm$ 0.02. These values are valid for densities greater than 200 kg m^{-3}. Rounded grains, specifically, have an upper limit of 580 kg m^{-3}. Although Srivastava et al. (2016) also investigated effective Young's moduli, they did not provide a parameterization, thus their work was not implemented as a method to calculate E.

Although Wautier et al. (2015) calculated orthotropic Poisson's ratio values, they did not provide a parameterization for Poisson's ratio explicitly, nor did they propose reasonable values to assume for different grain forms. Thus, we do not include a method for estimating effective, isotropic Poisson's ratio from their study.

https://doi.org/10.5194/egusphere-2026-4761
Preprint. Discussion started: 18 August 2026

Table 3. Applicability of implemented parameterizations for Young's modulus (E) and Poisson's ratio (ν) by grain form. A dash indicates the method is not applicable for that grain form.

	Köchle and Schneebeli (2014)	Wautier et al. (2015)	Bergfeld et al. (2023)	Schöttner et al. (2026)	Srivastava et al. (2016)
Density range (kg m^{-3})	150–450	103–544	110–363	Varies[a]	>200
Precipitation particles (PP, PPgp)	–	–	E	–	ν
Decomposing and fragmented (DF)	–	E	E	E	ν
Rounded grains (RG, RGxf)	E, ν	E	E	E	ν
Faceted crystals (FC, FCxr)	E, ν	E	–	E	ν
Depth hoar (DH)	E, ν	E	–	E	ν
Surface hoar (SH)	–	–	–	E	–
Melt forms (MF, MFcr)	E	E	–	–	–
Ice forms (IF)	–	–	–	–	–
Machine made (MM)	–	–	–	–	–

[a] Schöttner et al. (2026) uses grain-form-specific power-law fits without a single shared density range in the current implementation. Sub-grain-form codes are resolved to their two-character parent class for Young's-modulus and Poisson's-ratio methods.

2.4 Framework for calculation pathways

Slab properties do not have independent inputs. For example, density is an input to multiple downstream calculations. The choice of one density method propagates into estimates of Young's modulus, Poisson's ratio, shear modulus, and downstream slab properties. The same field observation can therefore lead to many distinct combinations of methods to calculate a target a parameter such as bending stiffness, D_{11} (Fig. 4). We define a "calculation pathway" as the set of upstream methods and parameters used to estimate a target parameter.

https://doi.org/10.5194/egusphere-2026-4761
Preprint. Discussion started: 18 August 2026

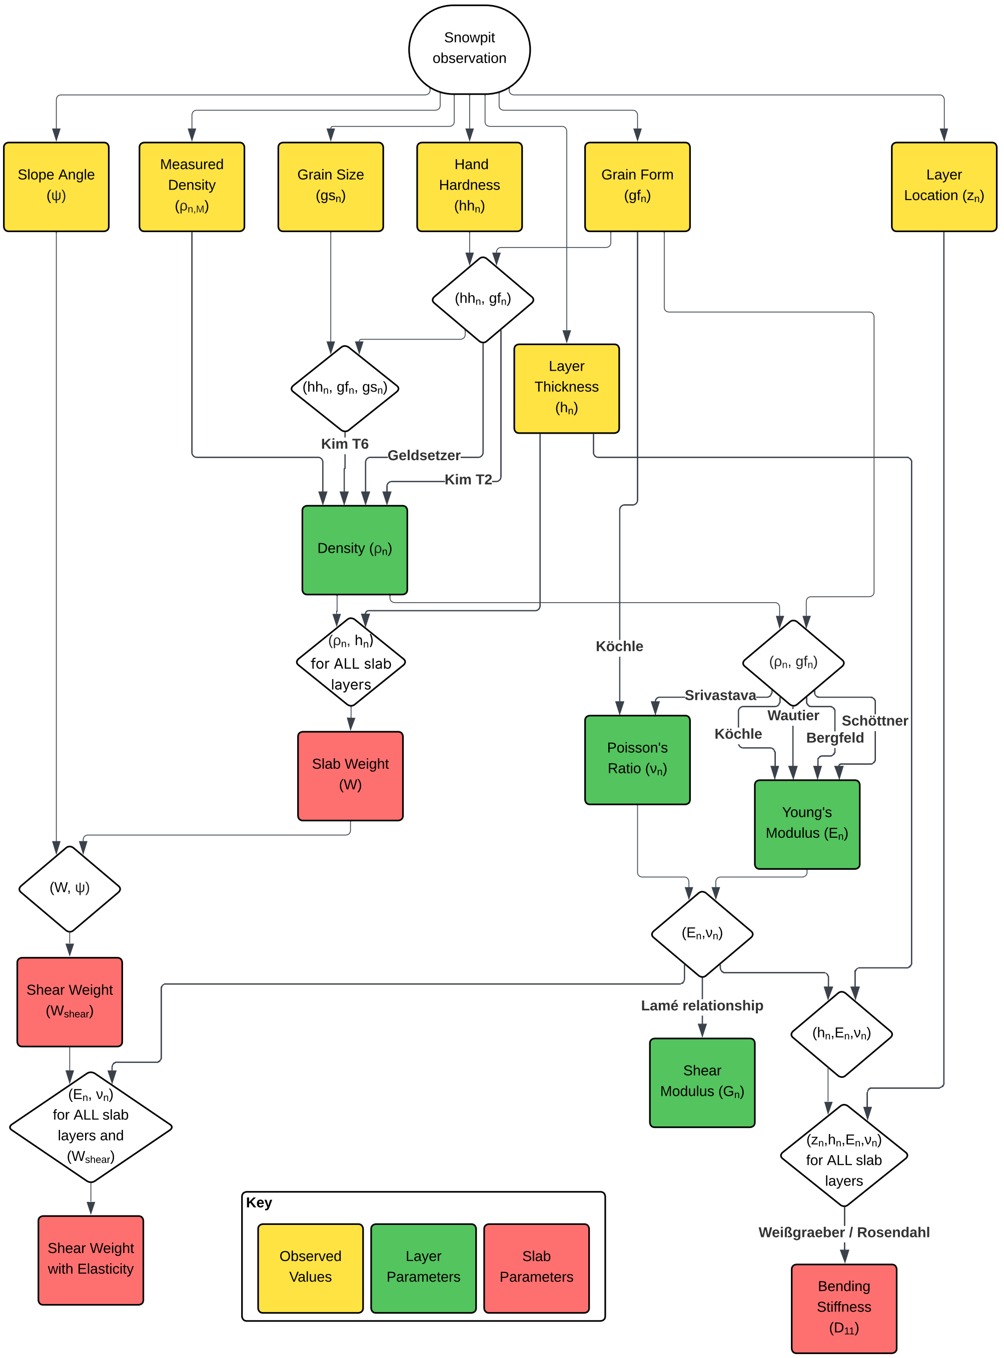

Figure 4. The directed acyclic graph (DAG) used to enumerate and execute the implemented calculation pathways. Yellow rectangles are observed snowpit values; green rectangles are layer-level parameters; red rectangles are slab-level parameters. Diamond shapes are merge nodes that collect the multiple inputs required by a single parameterization; rectangular parameter nodes are resolved by selecting one incoming method per pathway. Labeled arrows represent named parameterization methods; unlabeled arrows represent direct data flow.

https://doi.org/10.5194/egusphere-2026-4761
Preprint. Discussion started: 18 August 2026

2.4.1 Data structures for implementing parameterizations

To implement the calculation pathway described, we developed a specific set of data structures within SnowPyt-MechParams to represent snow pits, layers, and slabs.

The general hierarchy of the data structures is similar to the `snowpit` object in SnowPylot, with the addition of the `Slab` class. The `Pit` class in SnowPyt-MechParams stores identifying information about the snowpit, a list of layers that make up the snowpit, stability test results, and a list of slab objects that are created based on the user defined definitions of the weak layer.

The `Layer` class stores measured properties and calculated properties of the layer. Some calculated properties are derived directly from measured properties to aid with method implementation. For example, the hand hardness of a layer is recorded as a string, while the hand-hardness index (HHI) is a numerical representation of that value and is computed upon initialization of the `Layer` object.

Other calculated properties are populated by the `ExecutionEngine` during pathway execution. These include calculated density, Young's modulus, Poisson's ratio, and shear modulus.

A `Slab` is defined as the set of `Layer` objects in a `Pit` that are above the defined location such as the failure depth of a stability test result. One `Pit` can have multiple `Slab` objects, for example, if multiple stability tests are conducted with failure on different layers.

The `Slab` class stores calculated properties of the slab and information about the layers that constitute the slab. Specifically, the `Slab` object carries the constituent `Layer` objects, slab weight per unit area, W, along with its normal, W_{normal},and shear, W_{shear}, components, and composite stiffness properties A_{11}, B_{11}, D_{11}, and A_{55}. These properties are populated by the `ExecutionEngine` during pathway execution.

Fig. 5 summarizes the `Pit-Layer-Slab` hierarchy in SnowPyt-MechParams and highlights which measured and computed fields feed the pathway engine.

In our analysis, we parse SnowPilot CAAML.xml files into SnowPylot `snowpit` objects, and then use those objects to create SnowPyt-MechParams `Pit` objects.

https://doi.org/10.5194/egusphere-2026-4761
Preprint. Discussion started: 18 August 2026

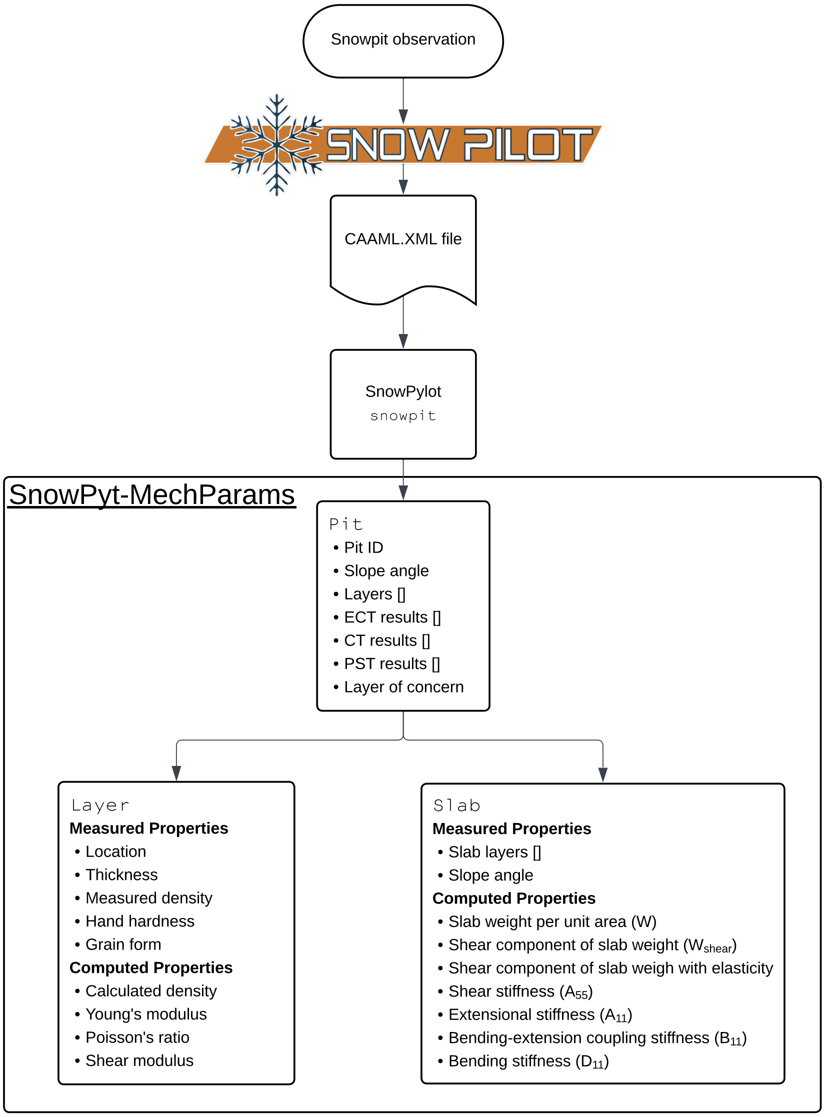

Figure 5. SnowPyt-MechParams data model. Each SnowPylot `snowpit` object is used to construct a SnowPyt-MechParams `Pit` object containing a list of `Layer` objects and a list of `Slab` objects. `Layer` objects carry both measured properties (location, thickness, measured density, hand hardness, and grain form) and computed properties (calculated density, Young's modulus, Poisson's ratio, and shear modulus) that are populated by the pathway `ExecutionEngine`. `Slab` objects store the constituent `Layer` objects and the computed slab-level properties: slab weight per unit area, W, shear weight W_{shear}, extensional stiffness, A_{11}, bending-extension coupling stiffness, B_{11}, bending stiffness, D_{11}, and shear stiffness, A_{55}.

https://doi.org/10.5194/egusphere-2026-4761
Preprint. Discussion started: 18 August 2026

2.4.2 Representing parameterizations as a directed, acyclic graph

A graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ is a collection of nodes (vertices, $\mathcal{V}$) and edges (connections, $\mathcal{E}$) that connect the nodes. In a directed graph, the edges specify a direction that enforces flow from input to output. A cycle is a closed path, or loop, that starts and ends at the same node. An acyclic graph is a graph where there are no cycles and a directed acyclic graph (DAG) is a directed graph that contains no cycles (Thulasiraman and Swamy, 1992).

In our framework, we represent the select set of parameters and methods as a DAG where nodes represent parameters, or collections of parameters, and edges represent the flow and transformation of data. Fig. 4 shows the DAG for the selected set of parameters and methods. The graph is directed because data must flow from snowpit observations to calculated parameters, where downstream parameters require upstream parameters as inputs. The graph is acyclic by nature because no parameter's computation can depend, directly or indirectly, on itself.

There are two types of nodes that represent distinct logical behavior:

- **Parameter Nodes** ($\mathcal{V}_P$): These nodes represent field measurements and mechanical parameters. In Fig. 4, parameter nodes are represented by boxes and the color indicates whether the node represents an observed value, layer parameter, or slab parameter.

- **Merge Nodes** ($\mathcal{V}_M$): These nodes represent the aggregation of multiple inputs required to implement a method represented by an outgoing edge. In Fig. 4, merge nodes are represented by diamonds and the node label indicates the parameters that are aggregated by the node.

Other graph components are:

- **Source Node** (S): The source node, S, is the snowpit observation. It has no incoming edges in the forward DAG, where information flows from observations toward calculated parameters. During backward traversal from a target parameter, the source node, S, becomes a terminal leaf of each enumerated pathway.

- **Edges** ($\mathcal{E}$): All edges are directed and edges are used to represent simple data flow or an implemented parameterization. Parameterizations are named in reference to the author(s) of the paper that proposed the method (section 2.3).

2.4.3 Finding calculation pathways

Two rules define the backward traversal of the graph (Fig. 4). At a parameter node, incoming edges represent alternative methods for calculating the same parameter, so the algorithm selects one incoming edge at a time (OR logic). For example, density can be obtained from a matched measurement or estimated using Geldsetzer, Kim T2, or Kim T6. At a merge node, incoming edges represent jointly required inputs, so the algorithm includes every incoming edge (AND logic). A density method that requires hand hardness and grain form must therefore include both observations.

Because merge nodes combine required inputs, the resulting object is usually not a linear path from source to target. Instead, the resulting "path" from the target back to the source is a branched dependency structure that collects required inputs and

https://doi.org/10.5194/egusphere-2026-4761
Preprint. Discussion started: 18 August 2026

methods. We use the project-specific term "PathTree" for this structure. In the implementation, the corresponding internal Python class is `PathTree`. Each PathTree has the requested target as its root, one branch for each jointly required input, and the snowpit source node at every terminal leaf. If a required input has multiple valid upstream alternatives, the algorithm combines those alternatives with the other required inputs using a Cartesian product.

Algorithm Goal: Given a target node $v \in \mathcal{V}$, find the set of all valid subgraphs $\mathcal{G}' \subseteq \mathcal{G}$ such that:

1. v is the root of $\mathcal{G}'$.
2. All leaves of $\mathcal{G}'$ are the source node S.
3. All logical constraints of $\mathcal{V}_P$ and $\mathcal{V}_M$ are satisfied within $\mathcal{G}'$.

The algorithm employs backward recursive traversal with memoization. Starting from target node t, dependencies are traced recursively toward source node S, building PathTrees that respect the dependency structure of parameter nodes and merge nodes.

For a node v, let $\text{PathTrees}(v)$ denote the set of all PathTrees rooted at v and terminating at S.

Base case:

$$\text{PathTrees}(S) = \{S\}$$

The source node (snowpit observation) is the terminating condition for the backward traversal. Thus, S is the source of the forward DAG, but also the terminating leaf of each path tree rooted at a target node v.

Recursive case for parameter nodes (OR logic):

$$\text{PathTrees}(v) = \bigcup_{\text{for each incoming edge } e \text{ from } u} \text{extend}(\text{PathTrees}(u), e)$$

where u is the predecessor at the start of incoming edge $e = (u, v)$, and extend adds e and node v to each path tree rooted at u. The union represents the selection of one incoming method at a parameter node.

Recursive case for merge nodes (AND logic):

$$\text{PathTrees}(v) = \text{PathTrees}(u_1) \times \text{PathTrees}(u_2) \times \ldots \times \text{PathTrees}(u_n)$$

where $u_1, \ldots, u_n$ are all predecessors of merge node v. Each element of the Cartesian product selects one valid upstream path tree for every required predecessor; the selected trees and their incoming edges are then joined at v.

Fig. 6 shows this logic for the simplest target, density. All four panels are derived from the common DAG in Fig. 4, but each panel represents a distinct PathTree and therefore a distinct calculation pathway. The direct-measurement PathTree contains a single branch from the snowpit observation to measured density and then to density. The Geldsetzer and Kim T2 PathTrees each branch to require both hand hardness and grain form. The Kim T6 PathTree requires those observations plus grain size.

https://doi.org/10.5194/egusphere-2026-4761
Preprint. Discussion started: 18 August 2026

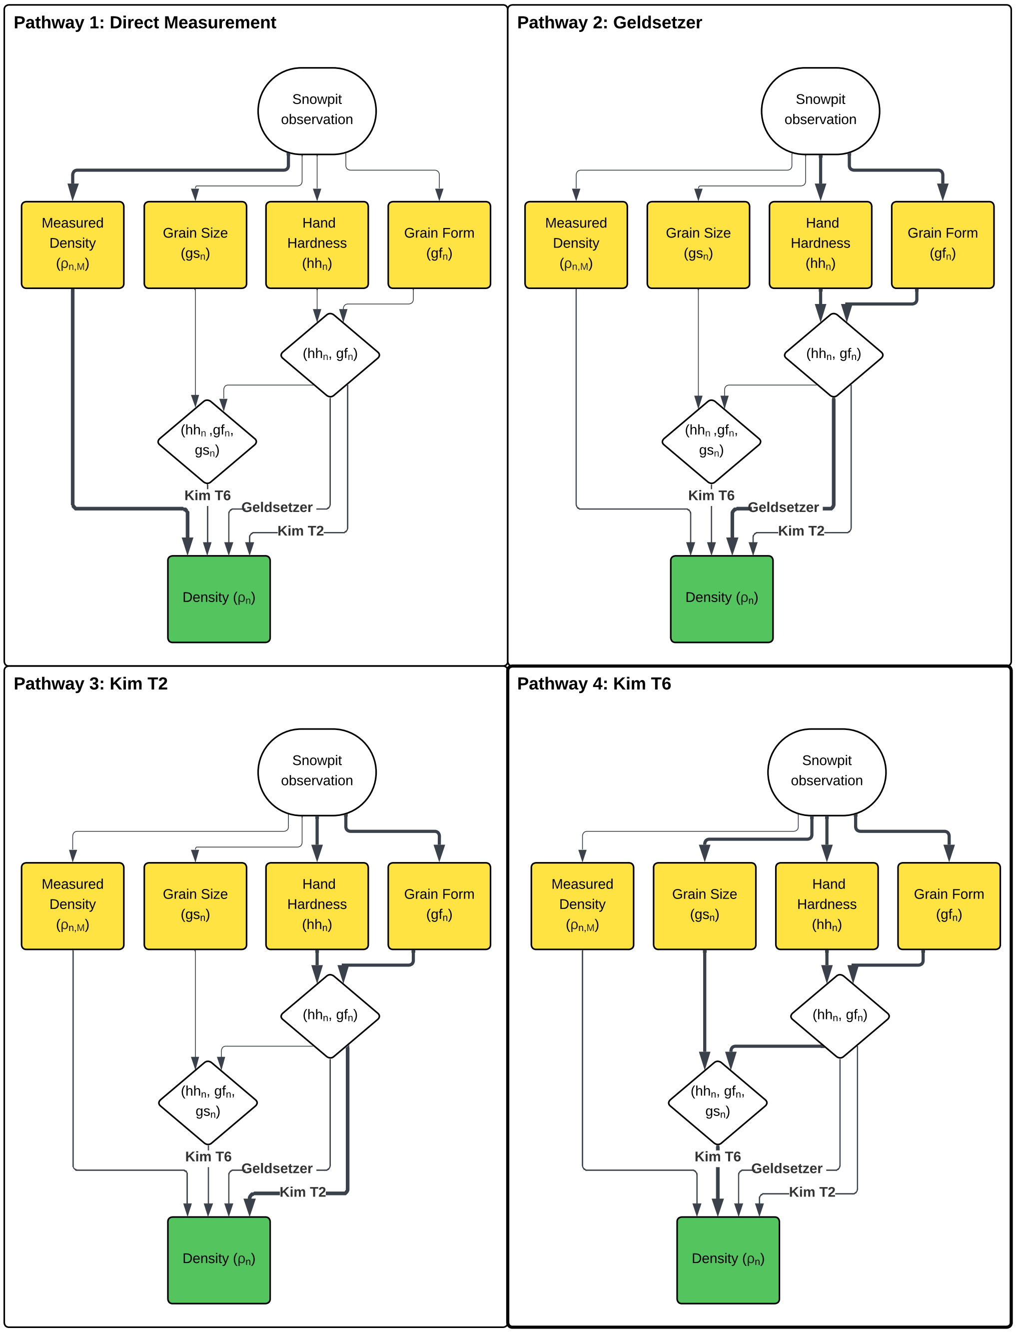

Figure 6. The four unique density pathways (PathTrees) enumerated from the DAG, each drawn as a subgraph of Fig. 4. Pathway 1 (Direct Measurement) uses a matched density observation. Pathway 2 (Geldsetzer) applies the Geldsetzer and Jamieson (2000) regression from hand hardness and grain form via the `(hh, gf)` merge node. Pathway 3 (Kim T2) applies the Kim and Jamieson (2014) Table 2 regression from the same inputs. Pathway 4 (Kim T6) additionally requires grain size, activating the `(hh, gf, gs)` merge node.

For more complex targets, the same rules are applied recursively. One key structural feature is the shared density node (Fig. 4). Density is an input to all Young's modulus parameterizations and to the Srivastava Poisson's ratio parameterization. To keep each pathway internally consistent, SnowPyt-MechParams uses a single density method within a pathway. Following traversal, each internal PathTree is converted to a structured `Parameterization` object that is executable by the `ExecutionEngine`. A fingerprint of the parameterization's input parameters and method choices is used to remove duplicates before execution.

The same pathway structure also preserves the distinction between layer-level and slab-level calculations. Density, Young's modulus, Poisson's ratio, and shear modulus are calculated independently for each layer. Slab weight (Eq. 1) and slab stiffness properties (Eqs. 7 - 10) are calculated for the slab as a whole because they depend on the complete set of layers that comprise the slab. Therefore, a slab-level pathway succeeds if and only if every layer in the slab can complete the required upstream layer-level calculations and the required pit- or slab-level observations, such as slope angle, are also available.

In our implementation, there are four available methods for estimating density, four methods for estimating Young's modulus, and two methods for estimating Poisson's ratio. The shear-weight target requires density, layer thickness, and slope angle, which results in four distinct calculation pathways. The shear-weight-with-elasticity target adds Young's modulus and Poisson's ratio availability for every slab layer. Because the density method is shared upstream of those calculations, this yields $4 \times 4 \times 2 = 32$ distinct slab-input pathways.

2.4.4 Executing parameterizations

The `ExecutionEngine` is a Python class that executes the resulting `Parameterization` on a SnowPyt-MechParams `Slab` object.

For each enumerated pathway, the execution engine first evaluates the required layer-level calculations (e.g. density, Young's modulus, and Poisson's ratio) in dependency order. It then evaluates slab-level calculations using those layer results. Previously computed values are cached by layer, parameter, and method so that shared upstream calculations are reused across pathways without changing the method choices within a pathway.

A pathway fails for a given `Slab` if any required field measurement data are absent, or if any of the input data for a parameterization is outside the defined validity range (Table 3). For example, if the grain form of an input layer is outside the method's supported grain form list, the pathway will fail for that slab.

Calculated parameters are only populated in `Layer` and `Slab` objects as far as the pathway is successful. For example, if density was successfully calculated for some layers in the slab, but not all, the successful `Layer` objects will have populated values for `density_calculated`, but the `Slab` object will not have a populated value for any calculated parameters that rely on density (which happens to be all slab level parameters in the example implementation).

2.5 Constructing `Slab` objects from ECT observations

To create `Slab` objects from the SnowPilot dataset, the SnowPilot CAAML.xml files were parsed into SnowPylot `snowpit` objects which were then used to create SnowPyt-MechParams `Pit` and `Layer` objects. For each ECTP result in a pit, the

https://doi.org/10.5194/egusphere-2026-4761
Preprint. Discussion started: 18 August 2026

weak layer was identified by comparing the recorded failure depth to the `depth_top` and `depth_bottom` of the layers in the pit. A SnowPyt-MechParams `Slab` object was created from the layers above the weak layer (not including the weak layer itself). The `Layer` objects within the `Slab` contain the measured hand hardness, grain form, grain size, and thickness of each layer, and the `Slab` contains the slope angle of the pit. A single Pit object can yield multiple `Slab` objects if it has multiple ECTP results.

This procedure yielded 14,776 slabs from 50,278 snowpits that were used to test the framework at population scale.

3 Results

3.1 Coverage comparisons of target parameters

In our example implementation, the graph yields 4 unique pathways to ascertain shear weight, W_{shear}, and 32 pathways to ascertain shear weight with elasticity properties (Figs. 4 and 6). The difference arises from the combinatorial nature of the calculation pathways. Each density method branches into four Young's modulus parameterizations, yielding 16 combinations that are then combined with two Poisson's ratio methods. The shear-weight pathway with the highest coverage succeeds for 5,470 of 14,776 ECTP slabs (37.0 %), whereas the highest coverage shear-weight-with-elasticity pathway succeeds for 687 slabs (4.6 %). Figure 7 compares these coverage ceilings for the four shear-weight pathways and the ten highest coverage shear-weight-with-elasticity pathways.

The four shear-weight pathways have markedly different practical coverage. The highest coverage pathway uses the Kim T2 density regression and succeeds for 5,470 slabs (37.0 %). It is followed by the Geldsetzer relationship with 4,187 slabs (28.3 %), Kim T6 with 697 slabs (4.7 %), and direct matched density observations with 109 slabs (0.7 %).

The shear-weight-with-elasticity target compounds the shear-weight limitations with the validity domains of the Young's modulus and Poisson's ratio methods (Table 3). Two pathways tie for the highest coverage at 687 slabs (4.6 %): Kim T2 density paired with either Wautier or Schöttner for Young's modulus and with Köchle for Poisson's ratio. Two additional pathways using Geldsetzer density, Wautier or Schöttner Young's modulus, and Köchle Poisson's ratio succeed for 684 slabs (4.6 %). Coverage declines when the density method is less widely applicable, when the Young's modulus method has narrower grain form or density support, or when the Poisson's ratio step uses Srivastava, which imposes an additional density threshold.

The highest coverage pathways all use Köchle for Poisson's ratio, indicating that once shear weight and Young's modulus are available, the Poisson's-ratio method is the final bottleneck at slab scale. This is particularly visible for pathways based on Srivastava, which fall outside the eight highest coverage pathways. Figure 8 shows the stepwise loss along the Kim T2–Wautier–Köchle pathway which has the highest coverage in in this implementation.

https://doi.org/10.5194/egusphere-2026-4761
Preprint. Discussion started: 18 August 2026

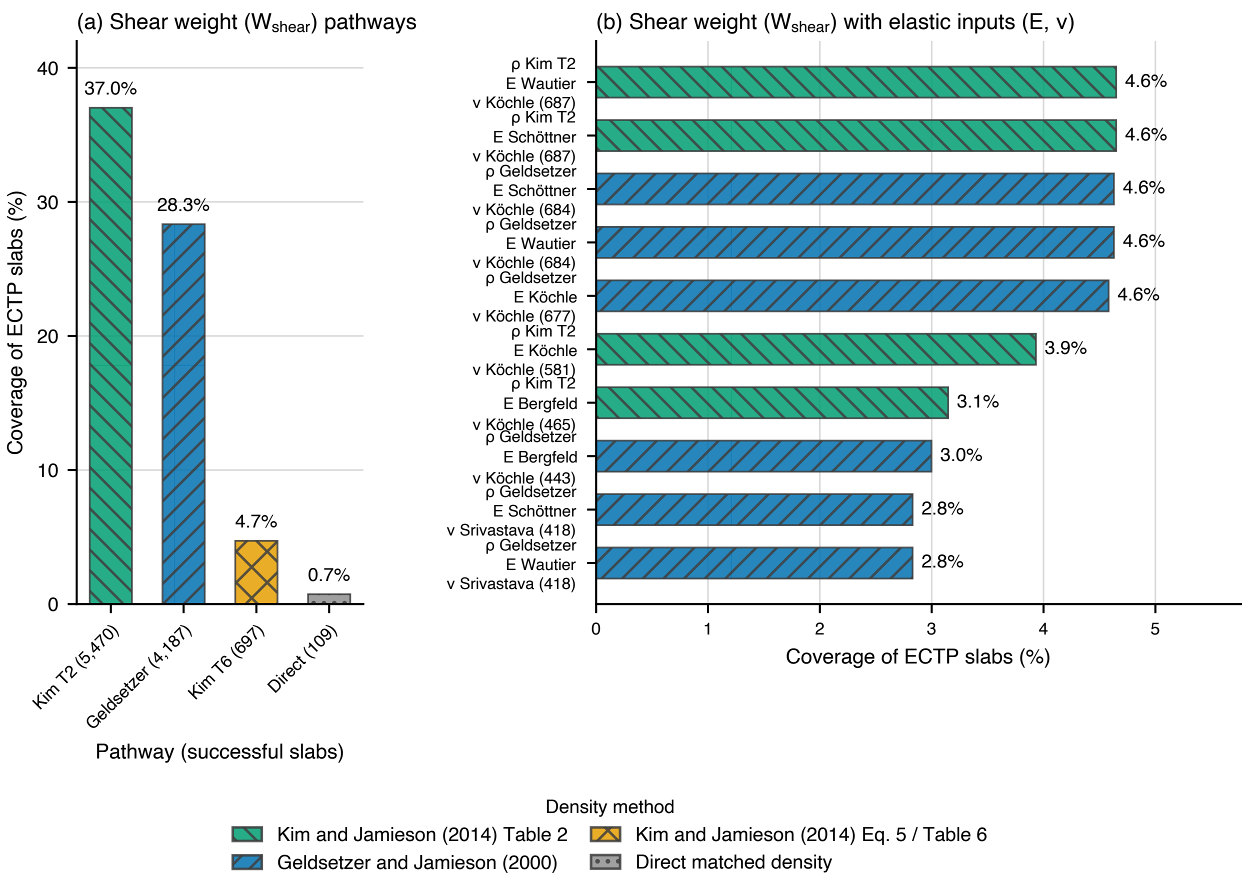

Figure 7. Coverage ceilings for the slab-property pathways evaluated on 14,776 ECTP slabs. Panel (a) shows the four pathways to ascertain shear weight, W_{shear}, from density, layer thickness, and slope angle. Panel (b) shows the ten shear-weight-with-elasticity pathways with the highest coverage out of the 32 implemented calculation pathways, where each pathway also requires valid Young's modulus and Poisson's ratio for every slab layer. Colors indicate the density method used upstream of each pathway.

https://doi.org/10.5194/egusphere-2026-4761
Preprint. Discussion started: 18 August 2026

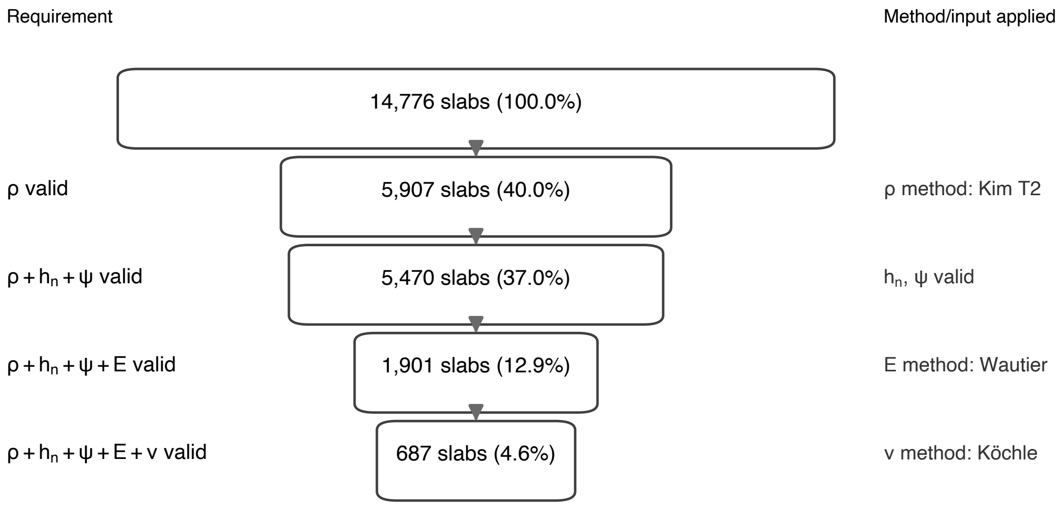

Figure 8. Attrition of slabs along the best-coverage shear-weight-with-elasticity pathway (Kim T2 → Wautier → Köchle). Counts show how many of the 14,776 ECTP slabs retain valid layer-wise inputs after each additional requirement is imposed; annotations identify the method or observed input applied at each funnel layer. The final step is successful only when density, layer thickness, slope angle, Young's modulus, and Poisson's ratio are all available for the slab. Boxes are not to scale.

https://doi.org/10.5194/egusphere-2026-4761
Preprint. Discussion started: 18 August 2026

3.2 Bending stiffness, D_{11}, magnitude and uncertainty comparison

The preceding section established that different pathways reach different subsets of slabs. Also of interest is the range in magnitude and propagated uncertainty of a target parameter across pathways that do succeed on the same slab. To quantify, we consider the bending stiffness, D_{11}, which depends on all upstream layer parameters: density, Young's modulus, and Poisson's ratio (Fig. 4, Eq. 9).

Of the 14,776 ECTP slabs, the 16 highest–coverage pathways share a subset of 317 slabs that produce a finite, positive calculated value for D_{11}. The requirement that all 16 pathways succeed restricts every layer in the common-slab subset to grain forms that resolve to the rounded-grain parent class (RG), because rounded grains are the only grain form for which all four Young's-modulus methods are simultaneously valid (Table 3). This subset of slabs are primarily composed of a single layer (258 of the 317 common slabs); the remaining slabs contain multiple layers that resolve to RG, often through differing rounded-grain subclassification (e.g. wind packed rounded grains, RGwp).

To isolate pathway effects, we normalized each pathway's D_{11} value by the median D_{11} across all 16 successful pathways for the same slab (Fig. 9). The paired ratios show consistent grouping by Young's-modulus method within the common-slab subset: Wautier pathways fall above the median, as do Köchle pathways to a lesser extent. Schöttner and Bergfeld pathways fall below the median, with Bergfeld nearer to it.

https://doi.org/10.5194/egusphere-2026-4761
Preprint. Discussion started: 18 August 2026

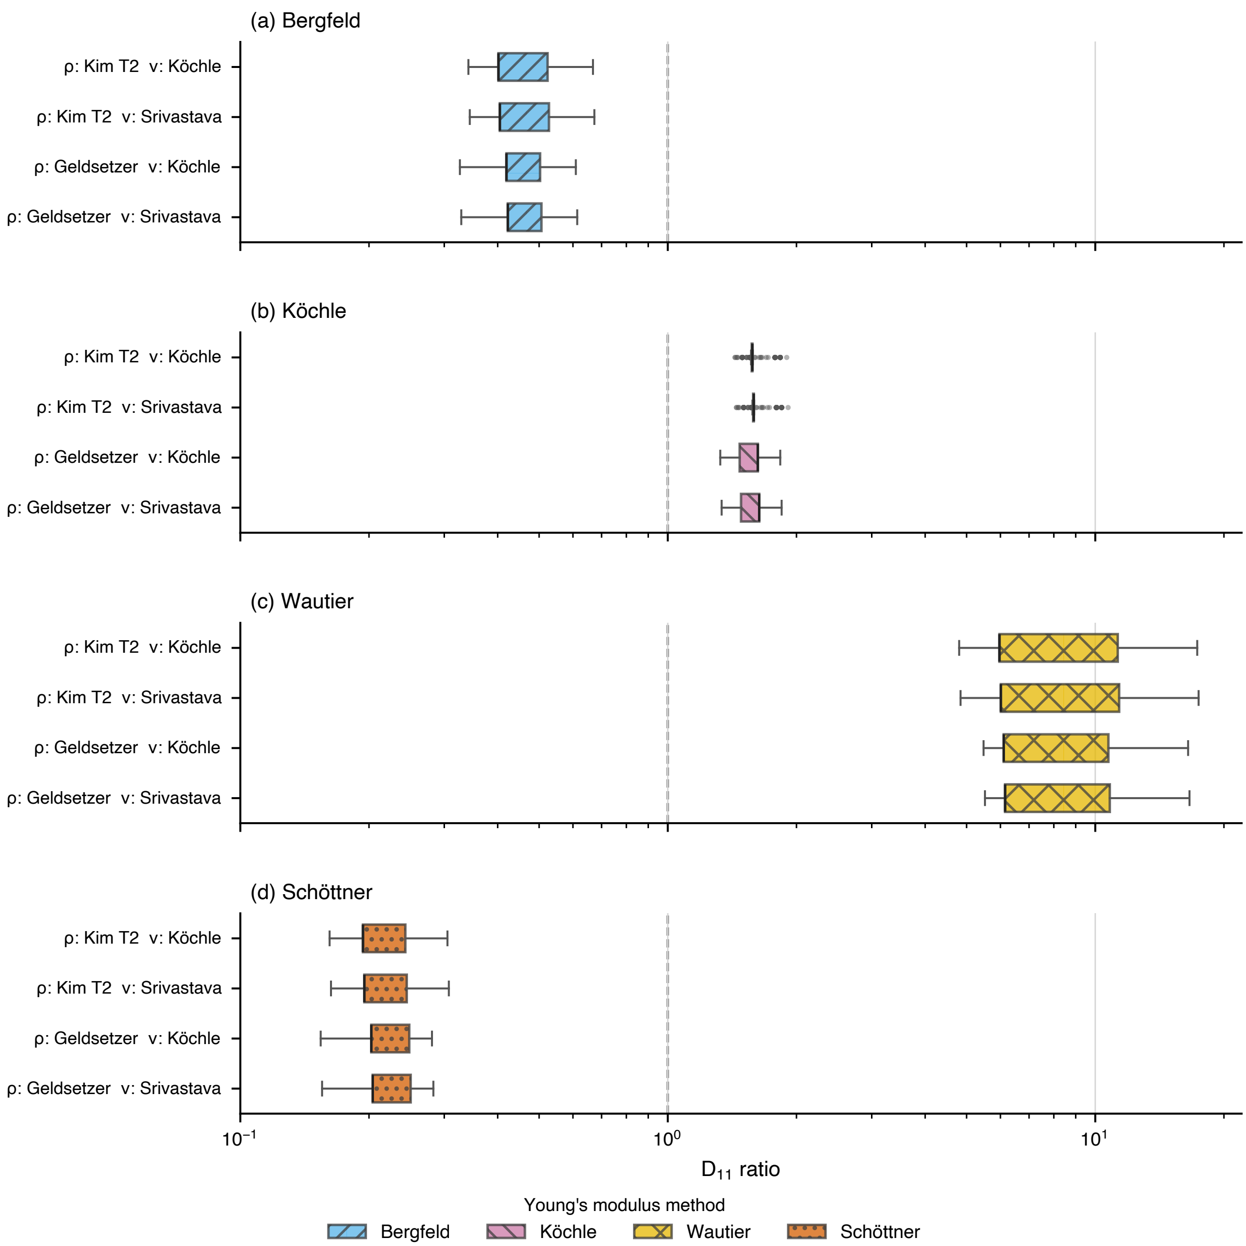

Figure 9. Box plots of paired pathway effects for D_{11} across the 317 slabs where all 16 highest-coverage pathways succeed, grouped and colored by the method used to calculate Young's modulus: (a) Bergfeld et al. (2023), (b) Köchle and Schneebeli (2014), (c) Wautier et al. (2015), and (d) Schöttner et al. (2026b). Each value is the pathway-specific D_{11} divided by the median D_{11} across all 16 successful pathways for the same slab. The dashed vertical line marks parity with that median.

https://doi.org/10.5194/egusphere-2026-4761
Preprint. Discussion started: 18 August 2026

For each slab, a spread ratio is calculated as the maximum over the minimum D_{11} values across the 16 valid pathways. For three unique slabs selected at the 30$^{\text{th}}$, 60$^{\text{th}}$, and 90$^{\text{th}}$ percentiles of the spread-ratio distribution (Fig. 10), the propagated measurement uncertainty (represented by error bars) is smaller than the horizontal separation between Young's-modulus groups, and also consistent within groups that share the same Young's-modulus method. Across the 317 common slabs, the median relative input-measurement uncertainty is 78.4 % (IQR: 55.8–84.8 %), while the median within-slab spread ratio (max/min D_{11}) is 32.

https://doi.org/10.5194/egusphere-2026-4761
Preprint. Discussion started: 18 August 2026

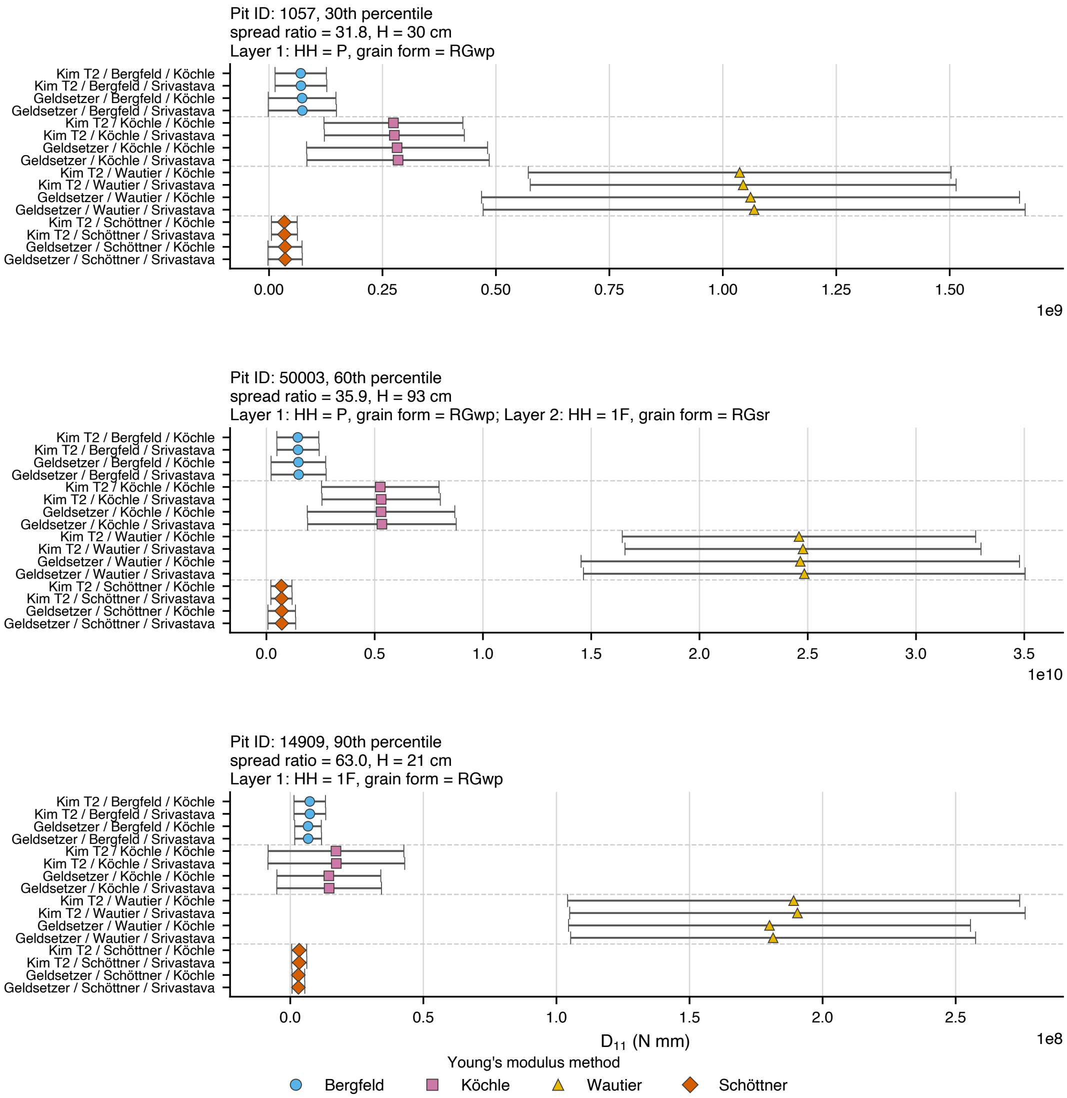

Figure 10. Bending stiffness, D_{11}, and propagated measurement uncertainty for three ECTP slabs selected at the 30th, 60th, and 90th percentiles of the spread-ratio distribution. The spread ratio is the maximum divided by the minimum D_{11} across the 16 pathways for each slab. Panel titles give the pit identifier, percentile, total slab thickness (H), spread ratio, and details about each layer including hand hardness (HH) and grain form. Points are colored by Young's-modulus method; error bars show ± 1 standard deviation of propagated input-measurement uncertainty. Each panel uses an independent linear x axis so pathway separation and uncertainty remain visible despite large differences in absolute D_{11} magnitude among slabs.

https://doi.org/10.5194/egusphere-2026-4761
Preprint. Discussion started: 18 August 2026

4 Discussion

4.1 Reasoning for coverage results

To ascertain shear weight, W_{shear}, layer thickness and density must be known or calculated for each layer in the slab, and the pit must have an observed slope angle. A valid layer in SnowPilot must contain layer thickness information, so layer thickness is always available. Shear-weight coverage is therefore determined by density availability across all slab layers and by the availability of slope angle at the pit level.

A given density parameterization method can only be considered successful on a slab if the method can successfully be executed on every layer within the slab. If any layer within the slab is missing data required to execute the method, or if a measured value falls outside the limits under which the method can successfully be applied, the method is considered unsuccessful for that slab. For example, the Geldsetzer method does not include a parameterization for melt-freeze crust (Fig. 3). If any layers within a slab have this primary grain form, then the Geldsetzer method is unsuccessful for that slab. Similarly, if any layers within the slab are missing recorded grain form, then the Geldsetzer, Kim T2, and Kim T6 methods will all fail.

The very low coverage directly observed layer density (2.8% of layers) reflects the sparsity of density measurements that align exactly with snow-profile layers after applying the density-profile matching rule. The drop in coverage from Kim T2 to Kim T6 shows the cost of requiring grain size in addition to hand hardness and grain form.

Taken together, these results quantify a practical input-assembly problem rather than a complete model-skill comparison. Shear weight has higher coverage from common snowpit observations because its slab property requirement is limited to density, layer thickness, and slope angle. Shear weight with elasticity results in a narrower subset of slabs due to the Young's modulus and Poisson's ratio constraints requiring validity for each layer in the slab. The graph representation makes this loss explicit and identifies where it occurs in the chain of calculation pathways.

4.2 Reasoning for magnitude variability in slab bending stiffness

The systematic difference in bending stiffness magnitude (Figs. 9 and 10) is largely attributed to the underlying Young's moduli parameterizations. The Wautier formulation uses a power-law exponent of 2.34 (Eq. 13), whereas Bergfeld uses 4.4 (Eq. 14) and Schöttner uses 4.6–5.1 depending on grain form (with the exception of SH). Köchle uses exponential regressions (Eq. 11 and 12). These different empirical relationships yield different Young's moduli, and D_{11} scales directly with E (Eq. 9). Plotting Young's modulus as a function of density for the different parameterizations (Fig. 11) shows the same ordering as D_{11} (Figs. 9 and 10). The 30$^{\text{th}}$ percentile example slab (pit ID: 1057) has a spread ratio of 31.8 whereas the 90$^{\text{th}}$ percentile example slab (pit ID: 14909) has a spread ratio of 63 (Fig. 10). The main difference between these two slabs is the hand hardness. The harder slab leads to higher density according to the density parameterizations (Sec. 2.3), which results in a smaller spread in elastic modulii (Fig. 11), and therefore a smaller spread in bending stiffness, D_{11}.

https://doi.org/10.5194/egusphere-2026-4761
Preprint. Discussion started: 18 August 2026

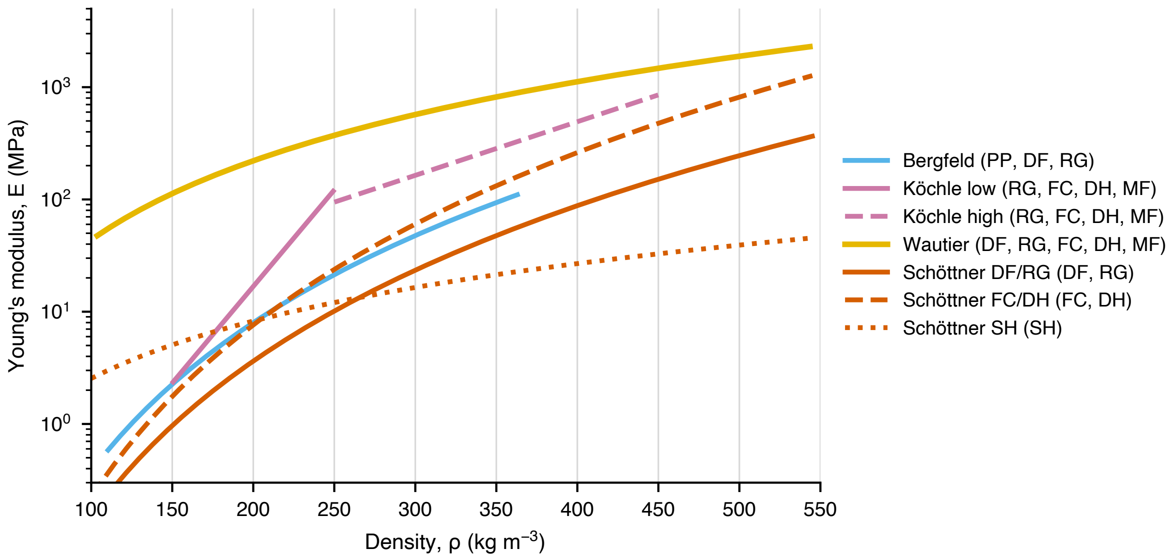

Figure 11. Implemented effective-isotropic Young's-modulus parameterizations as functions of density on a logarithmic vertical scale. Curves are plotted only within each method's implemented validity range. The lower Wautier exponent produces larger E values than the steeper Bergfeld and Schöttner relationships over much of the low-to-moderate density range, supporting the ordering of the D_{11} pathway estimates in Fig. 10.

4.3 Limitations and assumptions

This analysis has several important limitations. First, SnowPilot is a community-sourced database that spans a wide range of observer expertise and completeness; nominal field values are recorded, but spatial variability within a pit is not. Second, our matching of density-profile observations to snow-profile layers is intentionally strict: only exact depth-thickness matches are retained, which likely underestimates the amount of usable direct density information but keeps the analysis reproducible. Third, the slab definition used here is tied to ECTP failure layers and therefore provides a convenient population-scale set of slab stratigraphies, not a complete description of every slab-avalanche problem in the database.

At the mechanical level, the framework assumes effective isotropic, linearly elastic layer properties and perfect bonding between slab layers in the laminate-theory calculations. Real snow is anisotropic and rate dependent, particularly for faceted snow and depth hoar. Coverage is also limited by the published parameterizations themselves: melt-freeze crusts, machine-made snow, surface hoar, and ice formations are poorly represented or absent in several of the implemented Young's modulus and Poisson's ratio methods (Table 3).

Source-study domains, implementation validity checks, and SnowPilot grain form mappings are related but not identical. Some implementation choices map source studies onto parent grain-form classes, while others enforce a narrower subset of the source domain. For example, the Köchle Young's modulus implementation accepts melt freeze grain forms even though

https://doi.org/10.5194/egusphere-2026-4761
Preprint. Discussion started: 18 August 2026

the source fits were demonstrated with rounded grains, faceted crystals, and depth hoar samples and melt forms were noted as distinct. The Wautier implementation includes melt forms and excludes precipitation particles as a conservative mapping from a broader metamorphism based source domain. These domain choices affect pathway coverage and should be interpreted alongside the numerical uncertainty estimates rather than as purely observational data loss.

4.4 Extensions and future work

The framework is designed to be extensible. Adding a new parameterization requires implementation of the calculation, registration in the dispatcher, and insertion of the corresponding node or edge in the graph. The graph traversal and execution engine then enumerate and compare the new method automatically against the existing pathways. This design means that newly published parameterizations can be integrated with limited changes to the surrounding infrastructure.

This extensibility can be applied to more elaborate formulations of snow's mechanical behavior not yet implemented. Future work could incorporate different constitutive relationships (Gaume et al., 2018; Verplanck and Adams, 2024; Vallero et al., 2025) and weak layer failure criteria (Reiweger et al., 2015; Schöttner et al., 2026a), extending the framework beyond elastic, isotropic slab properties.

Adding new parameterizations may extend the number of calculation pathways for a target parameter. One area of future work is therefore defining a preferred calcuation pathway among valid options the graph enumerates. A higher coverage pathway may be attractive operationally, but a lower coverage pathway may still be more defensible if it better reflects the underlying mechanics or performs better against independent measurements. The framework developed here is therefore best understood as an organizing structure for comparison, not as a ranking system.

Instruments and methods to directly observe mechanical properties could circumvent the framework and minimize uncertainty of snow's mechanical properties and generate validation data for determining accurate calculation pathways. Continuing to develop instruments such as the blade hardness gauge (Borstad and McClung, 2011), snow micropenetrometer (Schneebeli and Johnson, 1998), and SnowSoundPen (Pietro et al., 2024) are all worthy pursuits to this aim.

Finally, snow cover modeling has progressed in recent years and is being adopted as a forecasting tool in locations such as Switzerland, the Western U.S., and Canada. Physics-based snow cover models attempt to create simulated snow stratigraphy across space and time, estimating layer properties such as thickness, density, and grain form. Since avalanche release stems from the mechanical failure of snow, knowledge of the snowpack's mechanical properties across space and time, with associated uncertainty, would be a boon to automated forecasting. The same graph framework could be driven by snow cover model output rather than by hand dug snowpits, similar to Viallon-Galinier et al. (2022), as a step toward automated, mechanically-based avalanche forecasting.

5 Summary and conclusions

Mechanical models of snow stability rely on mechanical properties of snow that are seldom directly measured in snowpits. To ascertain these mechancial properties, published parameterizations are chained together with methods that have have not

https://doi.org/10.5194/egusphere-2026-4761
Preprint. Discussion started: 18 August 2026

previously been systematically enumerated and compared. We develeoped SnowPyt-MechParams, an extensible, graph-based framework that enumerates calculation pathways from snowpit observations to slab-level mechanical properties and propagates measurement uncertainty through each pathway. As an example application of the framework, we developed SnowPylot to parse 10 years of data from the SnowPilot database into structured Python objects and implemented the shear component of slab weight per unit area with and without elasticity as target parameters. We applied the directed, acyclic graph to 14,776 slabs constructed from propagating Extended Column Tests. This example highlights a tradeoff that the mechanically richer target parameter with elasticity comes at the cost of lower coverage of valid calculation pathways. Furthermore, different published parameterizations reaching the same target parameter can disagree substantially in magnitude. Together these results illustrate how the framework can be used to identify bottlenecks and quantified uncertainty in parameters used to implement avalanche release models. Key conclusions are:

– SnowPyt-MechParams provides a reproducible, extensible, graph-based framework for enumerating and executing valid calculation pathways from snowpit observations to mechanical properties of snow.

– SnowPylot converts CAAML.xml SnowPilot export files into structured Python objects, enabling efficient analysis of a large, international snowpit database.

– In the example implementation of 14,776 ECTP slabs, the four shear-weight pathways reach up to 37.0 % coverage, whereas the 32 shear-weight-with-elasticity pathways reach only 4.6 % in the highest coverage case, reflecting the coverage cost of longer parameterization chains.

– Propagated measurement uncertainty compounds through chained calculations. In this example, slab bending stiffness, D_{11}, carries a median relative uncertainty of 78 % (IQR: 56–85 %).

– Because multiple pathways can reach the same target, estimated values vary substantially and this range grows as pathways lengthen. In this example, for slab bending stiffness, D_{11}, the median ratio of maximum to minimum estimate for an individual slab was 32, with ratios up to 63.

– These example results demonstrate that the practical applicability of mechanically richer slab property calculations is constrained not only by model formulation, but by the limited observational coverage and validity domains of currently available slab-input parameterizations. The framework is designed to reveal such constraints for any target parameter, not only the ones demonstrated here.

Code availability. SnowPyt-MechParams is openly available at github.com/connellymk/SnowPyt-MechParams. SnowPylot is openly available at github.com/connellymk/snowpylot.

https://doi.org/10.5194/egusphere-2026-4761
Preprint. Discussion started: 18 August 2026

Data availability. The SnowPilot database is publicly accessible at snowpilot.org. The CAAML export files used in this analysis cover the period 1 September 2014 to 31 August 2025.

Author contributions. MKC did the data curation, formal analysis, and algorithm development. MKC and SVV wrote the manuscript and made the figures with reviewing and editing by PR and VA. All authors were involved in conceptualization. SV and PR were the primary project administrators. Funding was acquired by SVV.

Competing interests. The authors declare no competing interest

Acknowledgements. Thanks to Jim Earl for helping to work with SnowPilot. Thanks to Bastian Bergfeld, Yannik Werner, and Florian Rheinschmidt for a helpful discussions.

Financial support. This worked was funded by Montana State University's Norm Asbjornson College of Engineering via Verplanck's faculty startup funding.

https://doi.org/10.5194/egusphere-2026-4761
Preprint. Discussion started: 18 August 2026

References

Bergfeld, B., Van Herwijnen, A., Bobillier, G., Rosendahl, P. L., Weißgraeber, P., Adam, V., Dual, J., and Schweizer, J.: Temporal Evolution of Crack Propagation Characteristics in a Weak Snowpack Layer: Conditions of Crack Arrest and Sustained Propagation, Natural Hazards and Earth System Sciences, 23, 293–315, https://doi.org/10.5194/nhess-23-293-2023, 2023.

Birkeland, K.: Personal communication, 2025.

Borstad, C. P. and McClung, D. M.: Thin-Blade Penetration Resistance and Snow Strength, Journal of Glaciology, 57, 325–336, https://doi.org/10.3189/002214311796405924, 2011.

CAA Technical Committee: Observation Guidelines and Recording Standards for Weather, Snowpack and Avalanches, Canadian Avalanche Association (CAA), Revelstoke, seventh edn., 2024.

Chabot, D., Kahrl, M., Consulting, K., Birkeland, K., and Anker, C.: A "NEW SCHOOL" TOOL FOR COLLECTING, GRAPHING, AND DATABASING SNOWPIT AND AVALANCHE OCCURRENCE DATA WITH A PDA, in: International Snow Science Workshop, 2004.

Chabot, D., Kahrl, M., and Earl, J.: SNOWPILOT: ONLINE, UPDATED AND FREE, in: International Snow Science Workshop, Breckenridge, Colorado, USA, 2016.

Conger, S. M. and McClung, D. M.: Comparison of Density Cutters for Snow Profile Observations, Journal of Glaciology, 55, 163–169, https://doi.org/10.3189/002214309788609038, 2009.

Fierz, C., Armstrong, R., Durand, Y., Etchevers, P., Greene, E., McClung, D., Nishimura, K., Satyawali, P., and Sokratov, S.: The International Classification for Seasonal Snow on the Ground, UNESCO, IHP–VII, Technical Documents in Hydrology, No 83; IACS contribution No 1, p. 80, 2009.

Gaume, J., Gast, T., Teran, J., van Herwijnen, A., and Jiang, C.: Dynamic Anticrack Propagation in Snow, Nature Communications, 9, https://doi.org/10.1038/s41467-018-05181-w, 2018.

Geldsetzer, T. and Jamieson, B.: Estimating Dry Snow Density from Grain Form and Hand Hardness, International Snow Science Workshop, pp. 121–127, 2000.

Gerling, B., Löwe, H., and van Herwijnen, A.: Measuring the elastic modulus of snow, Geophysical Research Letters, 44, 11–088, 2017.

Greene, E., Birkeland, K., Elder, K., McCammon, I., Staples, M., Sharaf, D., Trautman, S., and Wagner, W.: Snow, Weather, and Avalanches: Observation Guidelines for Avalanche Programs in the United States, American Avalanche Association, Bozeman, 4 edn., 2022.

Heierli, J., Gumbsch, P., and Zaiser, M.: Anticrack nucleation as triggering mechanism for snow slab avalanches, Science, 321, 240–243, https://doi.org/10.1126/science.1153948, 2008.

Höller, P. and Fromm, R.: Quantification of the Hand Hardness Test, Annals of Glaciology, 51, 39–44, https://doi.org/10.3189/172756410791386454, 2010.

Kim, D. and Jamieson, B.: ESTIMATING THE DENSITY OF DRY SNOW LAYERS FROM HARDNESS, AND HARDNESS FROM DENSITY, in: International Snow Science Workshop, Banff, 2014.

King, J., Howell, S., Brady, M., Toose, P., Derksen, C., Haas, C., and Beckers, J.: Local-Scale Variability of Snow Density on Arctic Sea Ice, The Cryosphere, 14, 4323–4339, https://doi.org/10.5194/tc-14-4323-2020, 2020.

Köchle, B. and Schneebeli, M.: Three-Dimensional Microstructure and Numerical Calculation of Elastic Properties of Alpine Snow with a Focus on Weak Layers, Journal of Glaciology, 60, 705–713, https://doi.org/10.3189/2014JoG13J220, 2014.

LaChapelle, E. R.: The Fundamental Processes in Conventional Avalanche Forecasting, Journal of Glaciology, 26, 75–84, https://doi.org/10.3189/s0022143000010601, 1980.

Lebigot, E. O.: Uncertainties: a Python package for calculations with uncertainties, https://pythonhosted.org/uncertainties/, available at https://pythonhosted.org/uncertainties/, 2010.

McCammon, I.: SLOPE MEASUREMENT FOR HUMANS: INCLINOMETER ERROR AND RISK COMMUNICATION, in: International Snow Science Workshop, Bend, Oregon, 2023.

Morin, S., Horton, S., Techel, F., Bavay, M., Coléou, C., Fierz, C., Gobiet, A., Hagenmuller, P., Lafaysse, M., Ližar, M., Mitterer, C., Monti, F., Müller, K., Olefs, M., Snook, J. S., Van Herwijnen, A., and Vionnet, V.: Application of Physical Snowpack Models in Support of Operational Avalanche Hazard Forecasting: A Status Report on Current Implementations and Prospects for the Future, Cold Regions Science and Technology, 170, 102 910, https://doi.org/10.1016/j.coldregions.2019.102910, 2020.

Nanthikesan, S. and Sunder, S. S.: Anisotropic Elasticity of Polycrystalline Ice Ih, Cold Regions Science and Technology, 22, 149–169, https://doi.org/10.1016/0165-232X(94)90026-4, 1994.

Norwegian Water Resources and Energy Directorate: Felthåndbok for Snø Og Skredobservasjoner, Norwegian Water Resources and Energy Directorate, Oslo, Norway, 3rd edn., 2022.

Pietro, T. D., Kaltenborn, J., Groninger, S., Mitterer, C., and Schneebeli, M.: ACOUSTIC SNOW STRATIGRAPHY MEASUREMENTS - FROM THE SOUND OF BREAKING SNOW TOWARDS ENHANCED SNOWPACK STABILITY ASSESSMENT, in: International Snow Science Workshop, pp. 1065–1072, Tromso, Norway, 2024.

Proksch, M., Rutter, N., Fierz, C., and Schneebeli, M.: Intercomparison of Snow Density Measurements: Bias, Precision, and Vertical Resolution, The Cryosphere, 10, 371–384, https://doi.org/10.5194/tc-10-371-2016, 2016.

Reiweger, I., Gaume, J., and Schweizer, J.: A New Mixed-Mode Failure Criterion for Weak Snowpack Layers, Geophysical Research Letters, 42, 1427–1432, https://doi.org/10.1002/2014GL062780, 2015.

Roch, A.: Les declenchements d'avalanches, in: International Symposium on scientific aspects of snow and ice avalanches, pp. 182–95, Davos, 1965.

Rosendahl, P. L. and Weißgraeber, P.: Modeling Snow Slab Avalanches Caused by Weak-Layer Failure–Part 1: Slabs on Compliant and Collapsible Weak Layers, The Cryosphere, pp. 115–130, 2020a.

Rosendahl, P. L. and Weißgraeber, P.: Modeling Snow Slab Avalanches Caused by Weak-Layer Failure–Part 2: Coupled Mixed-Mode Criterion Forskier-Triggered Anticracks, The Cryosphere, pp. 131–145, 2020b.

Schneebeli, M. and Johnson, J. B.: A constant-speed penetrometer for high-resolution snow stratigraphy, Annals of Glaciology, 26, 107–111, https://doi.org/10.3189/1998aog26-1-107-111, 1998.

Schöttner, J., Piecuch, M., Rosendahl, P. L., Walet, M., Adam, V., Weißgraeber, P., Schweizer, J., and Van Herwijnen, A.: Testing the Strength of Buried Surface Hoar Weak Layers Under Combined Compression and Shear Loading, Geophysical Research Letters, 53, e2025GL120 870, https://doi.org/10.1029/2025GL120870, 2026a.

Schöttner, J., Zeller-Plumhoff, B., Hagenmuller, P., Weißgraeber, P., Rosendahl, P. L., Löwe, H., Schweizer, J., and Van Herwijnen, A.: The Influence of Snow Microstructure on the Compressive Mechanical Properties of Weak Snowpack Layers, Acta Materialia, 302, 121 657, https://doi.org/10.1016/j.actamat.2025.121657, 2026b.

Schweizer, J., Reuter, B., Herwijnen, A. V., and Gaume, J.: Avalanche Release 101, in: International Snow Science Workshop, pp. 1–11, Breckenridge, Colorado, USA, 2016.

Simenhois, R. and Birkeland, K. W.: The Extended Column Test: A Field Test for Fracture Initiation and Propagation, in: International Snow Science Workshop, Telluride, CO, USA, 2006.

Siron, M., Trottet, B., and Gaume, J.: A theoretical framework for dynamic anticrack and supershear propagation in snow slab avalanches, Journal of the Mechanics and Physics of Solids, 181, 105 428, https://doi.org/10.1016/j.jmps.2023.105428, 2023.

Srivastava, P. K., Chandel, C., Mahajan, P., and Pankaj, P.: Prediction of Anisotropic Elastic Properties of Snow from Its Microstructure, Cold Regions Science and Technology, 125, 85–100, https://doi.org/10.1016/j.coldregions.2016.02.002, 2016.

Thulasiraman, K. and Swamy, M. N. S.: Acyclic Directed Graphs, in: Graphs: Theory and Algorithms, chap. 5.7, p. 118, John Wiley and Sons, 1992.

Toft, H. B., Verplanck, S. V., and Landrø, M.: How Hard Do Avalanche Practitioners Tap during Snow Stability Tests?, Nat. Hazards Earth Syst. Sci., 24, 2757–2772, https://doi.org/10.5194/nhess-24-2757-2024, 2024.

Vallero, G., Barbero, M., Barpi, F., Borri-Brunetto, M., and De Biagi, V.: An Elasto-Visco-Plastic Constitutive Model for Snow: Theory and Finite Element Implementation, Computer Methods in Applied Mechanics and Engineering, 433, 117 465, https://doi.org/10.1016/j.cma.2024.117465, 2025.

Verplanck, S. V. and Adams, E. E.: Dynamic Models for Impact-Initiated Stress Waves through Snow Columns, Journal of Glaciology, 70, e59, https://doi.org/10.1017/jog.2024.26, 2024.

Viallon-Galinier, L., Hagenmuller, P., Reuter, B., and Eckert, N.: Modelling Snowpack Stability from Simulated Snow Stratigraphy: Summary and Implementation Examples, Cold Regions Science and Technology, 201, 103 596, https://doi.org/10.1016/j.coldregions.2022.103596, 2022.

Wautier, A., Geindreau, C., and Flin, F.: Linking Snow Microstructure to Its Macroscopic Elastic Stiffness Tensor: A Numerical Homogenization Method and Its Application to 3-D Images from X-ray Tomography, Geophysical Research Letters, 42, 8031–8041, https://doi.org/10.1002/2015gl065227, 2015.

Weißgraeber, P. and Rosendahl, P. L.: A closed-form model for layered snow slabs, The Cryosphere, 17, 1475–1496, https://doi.org/10.5194/tc-17-1475-2023, 2023.