ATES Mapping for the Alpine Region

Technical report on method and implementation

Timo Braun

topoˆalpin

Document version 2.3 — 22 October 2026

describes product version 2.0

Contents

1 Introduction . . . 4
2 What version 2 changes . . . 5
2.1 One computation for the Alps instead of a mosaic of countries . . . 5
2.2 Release areas as in the reference . . . 5
2.3 Mobility from the size of the release area . . . 5
2.4 A classification without cell count . . . 6
2.5 Measured against the reference . . . 6
2.6 What it means for the map . . . 6
3 State of the Art . . . 7
3.1 Skitourenguru (RSA) . . . 7
3.2 CAT/ATH (SLF / White Risk) . . . 7
3.3 AutoATES v2 (Toft et al. 2024) . . . 7
3.4 Developments in 2026 . . . 7
3.5 Comparison . . . 8
4 Methodology . . . 8
4.1 Data . . . 8
4.2 PRA — release areas . . . 12
4.3 Wind shelter . . . 14
4.4 Flow-Py runout . . . 14
4.5 From the release mask to release segments . . . 17
4.6 Classification . . . 19
4.7 Alpine calibration . . . 21
5 Implementation . . . 22
5.1 Verification strategy: two checks, two questions . . . 22
5.2 Use of external tools instead of reimplementation . . . 23
5.3 Computation in tiles with a margin . . . 24
5.4 Memory . . . 25
6 Verification against the reference implementations . . . 25
6.1 Two questions, two runs . . . 25
6.2 End-to-end workflow without intermediate-input injection . . . 26
6.3 Documented deviations from the reference configuration . . . 26
6.4 What the comparison does not cover . . . 28
6.5 Effects of tiled hydrological processing . . . 28
7 Borders and the outer edge . . . 29
7.1 Inner borders . . . 29
7.2 The outer edge . . . 29
8 Plausibility check . . . 30
8.1 Model output . . . 30
8.2 Comparison with the per-region runs . . . 31
8.3 Comparison with accident locations . . . 31
9 Discussion . . . 33
9.1 Strengths . . . 33
9.2 Limitations . . . 33
9.3 Outlook . . . 35
10 Coverage . . . 36
11 Licences . . . 37

11.1 Software 37
11.2 Data 38
References 38

1 Introduction

Scope and status of this report. This technical report documents the data, algorithms, parameter choices and implementation checks used to produce the topo^alpin avalanche-terrain layer. The workflow is derived from AutoATES and the Austrian ATI model chain, with documented adaptations. Agreement with the reference implementations is assessed separately from the validity of the terrain classification. **The 5-m Alpine product has not yet undergone independent, regionally representative validation against expert ATES assessments.**

Three terms are used consistently: *verification* — does the implementation compute the specified method? (Section 6); *plausibility check* — are the results consistent with known observations? (Section 8); *validation* — does the model demonstrably meet defined quality requirements for its intended use? (not yet done; Section 9.3).

The report is written by the provider of the product and has been reviewed once externally (October 2026); it is not a peer-reviewed publication. The software, configuration and inputs are not published, so the results cannot currently be reproduced independently.

The Avalanche Terrain Exposure Scale (ATES) classifies terrain by the severity of avalanche exposure and is a central tool for risk assessment in the backcountry (Statham et al.[1]). The original classification was manual, based on 11 qualitative terrain criteria, and depended heavily on the expertise of the mapper.

What ATES is not: a statement about avalanche danger. In the conceptual model of avalanche hazard by Statham et al.[2], terrain is one of several components — alongside avalanche problem, likelihood of triggering and expected size, all of which vary with the snowpack. An ATES map is independent of the day: it states what exposure a place **has**, not what danger **prevails** on a given day. Every statement in this document is to be read in that sense, including the evaluation against accident locations in Section 8. In recent years, automated methods have increasingly complemented manual mapping and made ATES maps possible for large regions.

The updated ATES v2.0 (Statham & Campbell[3]) expands the original three-level scale to five classes by adding Class 0 (Non-avalanche) and Class 4 (Extreme), and introduces more quantitative criteria:

Class	Colour	Description
Value 0 (model)	— (not drawn)	No avalanche exposure detected by this model. **Not** a verified ATES Class 0 (Non-avalanche) assessment — see Section 4.6.3
Simple	Green	Low exposure, simple terrain
Challenging	Blue	Moderate exposure, runout zones or moderate PRA
Complex	Pink	High exposure, significant release and runout zones
Extreme	Violet	Very high exposure, steep PRA, terrain traps

Table 1: ATES classes according to v2.0 (Statham & Campbell[3]). Colours as rendered in the map tiles and the app. Class 0 is optional in the scale; this product carries a model value 0 instead, which is not rendered. On glaciers, a product rule outside ATES v.2 sets a minimum class of 2 (Section 4.6.2).

Various approaches to automation exist (Schmudlach & Köhler[4], Harvey et al.[5][6], Toft et al.[7]). More recent work investigates machine-learning approaches as an alternative (Markov et al.[8]). The approaches differ in method, coverage and availability (see Section 3).

This document describes how topo^alpin — an application for tour planning, navigation and mountain weather in the Alps (web and iOS) — uses ATES maps as a terrain overlay and as the consequence dimension in Tour-Check. The implementation is based on AutoATES v2 with the Alpine calibration (Huber et al. [9]) and computes the whole Alpine region on one 5-m grid, from official elevation models that are laser-scanned almost everywhere. The classification follows an adapted AutoATES workflow; a glacier-related minimum-class rule is retained as an additional product heuristic and is not part of ATES v.2 (Section 4.6.2). All parameters used come from the referenced publications and are documented there; our own adjustments are identified in this document.

2 What version 2 changes

Version 1 of the ATES map was released in spring 2026 and is the one the app has used since. It was an implementation of AutoATES v2 with the Alpine parameters of Huber et al. [9], computed per country or province on the grid of each national elevation model and put together into one map afterwards. Version 2 keeps the model and its classes. What changes is how faithfully the model is implemented, on which data it runs, and how the result is checked. This chapter summarises the differences for a reader who knows version 1; the following chapters describe version 2 on its own.

2.1 One computation for the Alps instead of a mosaic of countries

Version 1 computed fourteen regions independently — Switzerland, Austria, Bavaria, Slovenia, seven French départements, South Tyrol, Trentino and the Aosta Valley — each on its own grid and with its own forest, and cut each result at the national border. Beyond its border a region knew neither the terrain nor the forest, and an avalanche starting on one side of a border ridge was not followed to the other. Version 2 computes the whole Alpine region on one 5-m grid from one elevation model, one forest mosaic and one glacier mosaic, in tiles with a 10-km margin whose seams are checked line by line (Section 4.1, Section 5.3). The map has no seams at national borders, and it covers four regions that version 1 did not: Lombardy, Piedmont, Veneto and Friuli Venezia Giulia, where forest comes from the Copernicus Tree Cover Density because no region-wide LiDAR canopy model is available (Section 4.1.1).

2.2 Release areas as in the reference

In version 1 a release area was every cell whose fuzzy membership exceeded the threshold. Version 2 follows the BFW model chain completely: the wind-shelter index enters in radians with the membership parameters of the model chain, a slope window removes terrain below 28° and above 60°, a ruggedness mask removes about a tenth of the remaining release area, and small and thread-like areas are sieved out (Section 4.2). The release areas are then divided into release segments along hydrological subcatchments, a step that version 1 did not have (Section 4.5).

2.3 Mobility from the size of the release area

Version 1 ran every start cell with the same friction angle of 26° and an energy-height limit of 270 m. In version 2 each run gets its friction angle and velocity limit from the area of its release segment: a large release area runs further than a small one, as in the BFW model chain (Section 4.5.3). The energy-height limit of each run follows from the velocity limit of its segment and ranges from about 9 m for the smallest to about 215 m for the largest release areas; the scalar values of version 1 remain only as an unused fallback (Section 4.4.3).

The runout core itself was rewritten. The core of version 1 propagated in generations and passed a cell's stored flux on again each time the cell was reached; on the test terrains it reached 1.28 to 1.41 times the area of com4FlowPy. The core of version 2 keeps the state of each run separately and agrees with com4FlowPy 2.1 cell for cell on the tested inputs (Section 4.4, Section 6). Forest acts through friction only, as in the model chain; the additional removal of mass (detrainment) of version 1 is switched off.

2.4 A classification without cell count

Version 1 combined three inputs into the base class: slope, runout angle and the number of release areas reaching a cell. The cell count is not part of the Eastern Alpine calibration (Huber et al.[9], Tab. 3) and was removed; class 4 is tested on the unsmoothed slope, as in the reference configuration (Section 4.6). Glaciers are taken from official inventories where they exist instead of from OpenStreetMap everywhere (Section 4.1.2).

2.5 Measured against the reference

Version 1 was described against the publications. Version 2 is measured against the source code of the authors: every stage is compared cell by cell with the BFW model chain and with com4FlowPy, the released workflow is run end to end against ours, and every parameter is read from their configuration and checked (Section 6). Where version 2 departs from the reference, the departure is named and measured (Table 12). The same document describes the product and is generated from its configuration: no number the pipeline sets is typed into the text.

2.6 What it means for the map

The two versions were compared cell by cell on the 212,481 km² that both assess — the fourteen regions of version 1 — with the version 1 rasters that the app uses today, resampled by nearest neighbour onto the Alpine grid.

The amount of terrain with a class stays the same; its grading changes. In both versions about a third of the compared area carries a class of 1 or higher (35.1 % in version 1, 34.0 % in version 2), and 86.6 % of the cells keep their class. Within the terrain of class 1 or higher, however, version 2 shifts weight from *simple* to *challenging* and *complex*: class 1 falls from 43.4 to 32.3 %, class 2 rises from 25.4 to 28.6 % and class 3 from 22.0 to 31.1 %, while class 4 falls slightly, from 9.2 to 8.1 %. The most frequent changes are 2→3 (2.2 % of the compared cells), 1→0 (2.0 %), 1→2 (1.9 %) and 1→3 (1.5 %). Overall 7.7 % of the cells move up and 5.6 % down.

The change grows with the relief. In the Drôme, the Isère and Bavaria 91–92 % of the cells keep their class; in Switzerland 85 %, in Austria 89 %. In the high mountains — Savoie, Hautes-Alpes, Alpes-Maritimes, South Tyrol, Trentino — it is 76–80 %, and in the Aosta Valley 71 %, where 20.7 % of the cells are raised. This is where large release areas and long runouts dominate, and where the size-dependent mobility of version 2 lets large avalanches run further than the single friction angle of version 1 did. The individual changes of Section 2 act together here; their shares have not been separated.

At the borders version 2 sees more. In Austria 7.4 % of the cells in the 10-km band along the border are raised against 5.9 % in the interior, in Bavaria 8.1 % against 1.1 %: the terrain and forest beyond the border, which version 1 did not know, now contribute release areas and runout.

Slovenia is the exception. There version 2 is lower on 9.6 % of the cells and higher on only 1.6 %, mainly 2→1, 2→0 and 3→1. The difference is larger within the Alpine arc, where both versions use

the same Slovenian LiDAR forest (12.5 % lower), than outside it, where version 2 uses Copernicus (7.9 % lower); the forest source is therefore not the cause. The cause has not been isolated.

At accident locations version 2 is not worse. Of the Swiss ski-touring accidents, 87.8 % lie in class 2 or higher at the reported point in version 1 and 89.3 % in version 2, a difference of +1.5 percentage points on the same accidents (95 % interval +0.7 to +2.3); in Austria 81.7 % and 82.0 %, a difference that cannot be distinguished from zero (Section 8). The redistribution between the classes does not lower the share of accident locations assigned class 2 or higher.

3 State of the Art

3.1 Skitourenguru (RSA)

A purely DEM-based method built on the concept of the "Relevant Slope Area" (RSA) (Schmudlach & Köhler [4]): for every terrain point a direction-dependent polygon is delineated that represents the relevant avalanche terrain. The ATES classification follows from weighted geomorphological properties of the RSA (slope, plan curvature, terrain fold). The method needs no simulation model and covers the entire Alpine region, but the weighting factors are not published.

3.2 CAT/ATH (SLF / White Risk)

A physically based method using RAMMS simulations (Harvey et al. [5], updated in CATv2/ATHv2 [6]). Release potential comes from a two-dimensional density estimate over more than 5,000 mapped release areas; runout zones from RAMMS::EXTENDED simulations of size-class 3 avalanches, with release polygons delineated automatically after Bühler et al. [10] and a fracture depth of 0.6 m. The result is CAT maps (Classified Avalanche Terrain) and ATH maps (Avalanche Terrain Hazard, $H = \sqrt{r \times c}$ with release potential r and consequences c). CATv2 operates on a 5-m DEM and is used for tour planning in the SLF White Risk app. The maps are available under CC BY 4.0, the code used to generate them is not publicly available.

3.3 AutoATES v2 (Toft et al. 2024)

An open-source Python framework (GPL-3.0) with three modules (Toft et al. [7]):

1. **PRA (Potential Release Area):** fuzzy logic on slope, wind shelter and forest (Veitinger et al. [11])
2. **Runout:** the Flow-Py energy-line model (D'Amboise et al. [12])
3. **Classification:** a multi-stage lookup table (slope × travel angle × cell count × forest × PRA)

Validated against expert maps in two areas of western Canada: the F1 score rises from v1 to v2 from 64 % to 77 % at Bow Summit and from 40 % to 71 % at Connaught Creek; the difference between the two areas shows how strongly performance depends on the study area. The forest modules of com4FlowPy (Huber et al. [13]) represent the effect of forest on avalanche propagation by raising the friction on forested cells.

3.4 Developments in 2026

autoATES v3.0 (Sykes et al. [14]) scales automated ATES mapping to about 400,000 km² of western Canada. It replaces average canopy cover with the area of forest gaps, introduces two scenarios (typical and infrequent) for release areas and runout, which captures low-angle terrain with persis-

tent weak layers, and outputs classes 0–4. On a transfer test (Connaught Creek) the classifier is reported to agree with expert consensus about as well as three mappers agree with each other.

The BFW model chain (Avalanche Terrain Identification, ATI[15]) is published as open software and extends AutoATES with a segmentation of the release areas along subcatchments and size-dependent mobility: each release segment gets its own reach angle and velocity limit from its size. Our implementation follows this chain (Section 4.5) and is verified against it (Section 6). Work on automated thalweg identification builds on the same chain (Spannring et al.[16]).

Neither development makes a change of method for this product necessary by itself; they do show that the size-dependent parametrisation and the treatment of class 0 are active questions of the field rather than settled practice.

3.5 Comparison

Criterion	STG[4]	CAT/ATH[6]	AutoATES v2[7]	ML[8]
Entire Alps	Yes	CH only	Yes	No (Pirin)
Runout model	No (RSA)	RAMMS	Flow-Py	No
Release vs. runout	Implicit (RSA)	Yes	Yes	No
Open source	No	No	Yes (GPL-3.0)	Yes
Proprietary tools	No	RAMMS	No	No
Forest integration	Binary	Binary	PFC + friction	Feature
ATES-compatible	Yes	Yes (ATHc)	Yes	Yes

Table 2: Comparison of existing approaches to automation.

4 Methodology

4.1 Data

The model works on 5×5 m, on **one** grid for the whole Alpine region: the Alpine DEM of 14 October 2026 (EPSG:3035, 192,217 × 124,536 cells, 291,635 km² of valid terrain). It is assembled from the official regional DEMs — swissALTI3D (Switzerland and Liechtenstein), BEV ALS-DTM (Austria), BVV DGM1 (Bavaria), IGN LiDAR HD MNT (seven French départements), ARSO DMR (Slovenia) and the DTMs of the Italian regions and provinces (Aosta Valley, Piedmont, Lombardy, South Tyrol, Trentino, Veneto, Friuli Venezia Giulia). Each source is brought onto the common grid once, clipped to its border polygon plus 75 m and blended with distance weights in a 75-m band where two sources overlap; at the nine seams the median elevation difference is 0.3–2.5 m. The resolution matches swissALTI3D, which Harvey et al.[6] use for CATv2, and is well above what AutoATES v2 typically uses (10–30 m).

Why one grid. A region computed on the grid of its own source knows neither the terrain nor the forest beyond its border, and many regions meet in the high mountains (Mont Blanc, Monte Rosa, Ortler, Dolomites). With one DEM, one forest mosaic and one glacier mosaic for the whole region, a national border is not a line at which the computation changes (Section 2); how the computation is nevertheless split into pieces is described in Section 5.3.

The cost is one resampling: each source DEM is converted once from its native grid to EPSG:3035 (bilinear). Where a source has gaps at the border, these were filled from the neighbouring DEM before assembly (Aosta Valley 5.9 km² along the border ridge, and the other Italian regions), but

only where the neighbour agrees with the region's own DEM (aborted if the interquartile range of the elevation differences in the overlap exceeds 15 m, or if more than 10 % of the overlapping cells fall outside the measured range). Remaining holes are water surfaces with documented limits and the Lessinia in Veneto (Section 7).

Source quality differs. The DEMs are not of equal quality. The Lombardy DTM (5 m, 2015) is largely derived from contour lines, with laser scanning only along the rivers; in the Bormio area it is visibly smoother than its neighbours and shows stripes that are present in the original. Slope there is less reliable than elsewhere (Section 9.2.2).

Forest data are derived per country from high-resolution elevation models — mostly from LiDAR campaigns, in Switzerland and Bavaria instead from dense matching of stereo aerial images (WSL VHM, BVV DOM20). From these sources a normalised surface model (nDSM / CHM / VHM, 20 cm–2.5 m) is obtained. Pixels with nDSM ≥ 5 m are classified as tree (Huber et al.[9], Schumacher et al.[17]). The 5 m threshold separates forest from dwarf pine and shrubs in the Alpine treeline zone. The binary tree pixels are aggregated to the DEM grid by block averaging and give the **Percent Forest Cover (PFC)** per 5 m cell (0–100 %).

Buildings and infrastructure also produce high values in the nDSM. An official forest mask is therefore applied per country (pixels outside forest polygons → PFC = 0). The masks come from national forest inventories or land-use mapping and are methodologically independent of the LiDAR signal.

- **Switzerland:** WSL Vegetation Height Model NFI, 2025 edition (1 m, EPSG:2056) [18] — a ready-made nDSM from stereo aerial images of the ADS-80/100 sensor, normalised against swissALTI3D; ODbL with DbCL. Not to be confused with the Vegetation Height Model LiDAR NFI from swissSURFACE3D (0.5 m,[19]), a separate dataset.
- **Austria:** BEV ALS DOM/DGM (1 m, LiDAR, EPSG:3035) — nDSM = DOM − DGM. Buildings and infrastructure are masked with the BFW forest map (Bauerhansl et al.[20], CC BY 4.0). The method follows Huber et al.[9].
- **Italy (South Tyrol):** the province's DOM and DTM (2.5 m, EPSG:25832) from the same survey of 2004–2006 — nDSM = DOM − DTM, area-weighted to the 5-m cell; DOM and DTM share grid and survey. Buildings and infrastructure are masked with the forest type map (CC0, mapped 2001–2008). If a block is missing, the run aborts. More recent surveys of the province cover only parts (Merano/Salorno 2009, Dolomites 2010, settlements 2013, glaciers). Data source: Autonomous Province of Bolzano – South Tyrol, Open Data.
- **Italy (Trentino):** canopy height model (CHM, 1 m) from the 2006–2007 flights of the provincial survey 2006–2008, computed as DSM − DTM by M. Dalponte (FEM) and published on Zenodo (record 14930083; Provincia Autonoma di Trento, CC BY). It comes from the same survey as the Trentino DEM; heights and forest come from one source, and 0.09 km² (0.0015 %) of the province are without measurement. The CHM of the 2014–2018 survey on the same record was considered and rejected: it leaves 349 km² of the province uncovered (Brenta, Pergine, western Val di Fassa, Pale di San Martino), and its nodata value 0 cannot be told apart from open terrain. Buildings and infrastructure are masked with the province's forest map "Bosco" (as of 31 December 2017, 1:10,000, classes 3100 and 3222). CHM and mask predate Vaia (October 2018).
- **Italy (Aosta Valley):** DSM and DTM of the harmonised 2005/2008 survey (2 m, EPSG:23032, CC BY 4.0) of the Regione Autonoma Valle d'Aosta — nDSM = DSM − DTM. Both models come from the same dataset (`DSM0508_002` / `DTM0508_002`), which guarantees pixel-exact co-registration. Buildings and infrastructure are filtered with the forest mask of the Ambiti Inedificabili Art. 33 (legally designated forest).

- **France:** IGN MNH LiDAR HD (Licence Ouverte 2.0) — normalised height model (nDSM) from the LiDAR HD campaign, obtained via the Géoplateforme (`data.geopf.fr`). The seven départements are covered by 50,599 tiles of 1 km², acquired between 14 April 2021 and 10 February 2023, mostly June to August; sensors RIEGL VQ-780 II-S (04, 05, 06, 26, 38) and VQ-1560 II (73, 74). IGN delivers the tile at 0.50 m; here it is requested at 1 m, giving 25 sub-pixels per 5-m cell. Building filter: OCS GE 2.0 (IGN, Licence Ouverte 2.0), forest classes CS2.1.1.1 (broadleaved), CS2.1.1.2 (coniferous) and CS2.1.1.3 (mixed), per département the edition at most one year from the survey (D004 2021, D005 2022, D006 2020, D026 2023, D038 2021 with the official correction, D073 2022, D074 2023); BD Forêt v2 was not used, because its editions (2014–2021) lie up to seven years from the canopy height. Each tile is requested several times with growing pauses; if it is still missing, the run aborts, so that a missing tile can never enter as "no forest" (Section 5.1). DEM: IGN LiDAR HD MNT (0.5 m, the same campaign as the canopy height, EPSG:2154, Licence Ouverte 2.0).
- **Slovenia:** CLSS nDMP (50 cm, EPSG:3794) — normalised digital surface model from the Ciklično Lasersko Skeniranje Slovenije. The nDMP represents the normalised canopy height directly — methodologically equivalent to the WSL VHM (CH) and IGN MNH (FR). Acquired from 21 March 2023 to 9 December 2025 in 16 blocks; the date of each tile is recorded, taken from the tile index of the CLSS viewer (WFS layer `fmp_clss:tileindex_clss_zdruzeno`, as of 7 October 2026). Slovenia is computed only in the Alpine arc: the visible area of the apps plus the runout halo of 3500 m, 9,944 of 23,038 tile entries. Sixteen tiles there come from the 2011 survey and fill gaps of the 2023–2025 campaign; they are recorded with the forest raster. For one tile (433_127, Jelovica) the nDMP is not offered; it is computed from the CLSS point cloud of the same survey. Tiles are fetched one by one at a fixed pace and reduced to their contribution to the 5-m cell; if a contribution is missing at assembly, the run aborts. Building filter: MKGP RABA (`RABA_ID` = 2000), 2026 edition, analogous to the BFW forest map (AT) and OCS GE (FR). The edition is chosen deliberately: its entries range from 28 August 2023 to 30 June 2026 and so cover the CLSS window, whereas the 2023 edition starts two years earlier; between the two editions, 76.2 km² change sides (0.64 % of the forest area, net +22 km²). DEM: ARSO DMR (1 m, EPSG:3794, CC BY 4.0). The surveying authority GURS makes the CLSS data for all of Slovenia freely available for viewing and download (clss.si).
- **Bavaria:** BVV DOM20 (20 cm, EPSG:25832, CC BY 4.0) — image-based surface model from the Bavarian aerial survey (dense matching of aerial images). nDSM = DOM − DGM1 per kilometre tile at 1 m, i.e. the ground from the DGM1 of the same kilometre, not from the 5-m DEM. The tile list comes from the LDBV catalogue; as in Slovenia, the computation runs in two phases — each tile is fetched, checked for length and position and reduced to its 5-m contribution, then everything is assembled, and if a contribution is missing, the run aborts. Buildings and infrastructure are filtered with the combined ATKIS Basis-DLM forest mask (`adv:AX_Wald` + `adv:AX_Gehoelz`, BVV WFS, CC BY 4.0). Combining both layers captures closed main stands as well as edge vegetation and transition zones, which is essential for forest braking of runout in Alpine terrain. Like Slovenia, Bavaria is computed in the Alpine arc: the perimeter of the Alpine Convention plus 3500 m, as far as its DEM reaches.

4.1.1 Copernicus in regions without LiDAR forest

For Lombardy, Piedmont, Veneto and Friuli there is no region-wide canopy height model from a single LiDAR survey. There the Copernicus High Resolution Layer **Tree Cover Density** applies (CLMS, VLCC series, 10 m, EPSG:3035, values 0–100 %) for the reference year 2024[21]. The series starts in 2018; only the later years show the forest damage from Vaia (October 2018) and the bark beetle, which hit Veneto and Friuli in particular.

Limits. Compared with LiDAR forest, Copernicus has three known limits:

No tree height dwarf pine (2 m) and high forest (20 m) cannot be told apart. Markov et al.[8] note that the product partly records dwarf pine as forest and does not detect small gaps — both act on the dangerous side.

10 m instead of 1 m forest edges and clearings blur.

A change of source at the border where LiDAR forest meets Copernicus, the forest effect changes with the source.

How large the difference is can be seen in South Tyrol, where both exist (LiDAR 2004–2006 with the forest type map against Copernicus 2024, 7,399 km², at 25 m): mean canopy cover 29.1 % against 38.5 %, correlation 0.80, the same AutoATES forest class on 68 % of the area. Copernicus records some 450 km² more "forest", concentrated in Überetsch, Vinschgau and around Merano — orchards and vineyards. Conversely, LiDAR records 84 km² of dense forest where Copernicus sees open terrain, concentrated around Welschnofen, Karersee and Eggental, an area hit by the Vaia storm of 2018. The comparison does not separate actual forest change from differences between the two products; both are visible in it.

Rule. LiDAR forest with an official mask where it exists, otherwise Copernicus; never both within one region. In the Alpine forest mosaic every 5-m cell has exactly one source, which is carried in a source map. Within the region of a LiDAR source, Copernicus is never used: a cell without a LiDAR value receives 0 % with a source code of its own, and the area is reported per region. Nothing is filled in. Where two border polygons overlap, the overlap is split along its centre line: each cell belongs to the region whose own area it lies closer to. The LiDAR forests are brought onto the grid of the Alpine DEM by nearest neighbour, so that at the edge of a source nothing is mixed with values beyond its coverage; Copernicus is resampled bilinearly from 10 to 5 m. Copernicus covers everything outside the LiDAR regions, including the fringe of the Alpine DEM beyond the national borders; where it reports a cell as outside the land area (sea), the forest is 0 %.

The forest mosaic of 19 October 2026 passes its hard checks (checksums of all inputs, grid identical to the DEM, no DEM cell without a forest value, LiDAR only within its own region, never Copernicus within a LiDAR region). LiDAR regions without a LiDAR value: 101.8 km² in Haute-Savoie (Lake Geneva, where the IGN raster ends at the shore), 6.3 km² in Bavaria and 0.08 km² in Trentino.

Trees that are not forest. With every tile, Copernicus provides a **Forest Additional Support Layer** (FADSL) that marks trees in agricultural use (from CLC 2018) and in urban context (from IMD 2018 and CLC 2018). In the LiDAR regions the official forest masks remove such areas; for the Copernicus regions we measured the effect, on the 10-m grid of the TCD and with slope from the same regional DEMs as in the Alpine DEM. Of 15,349 km² of canopy cover in the four regions, FADSL marks 75.9 km² (0.49 %). Of this, 4.7 km² lie on terrain steeper than 30°; above 1000 m, in or within 100 m of terrain steeper than 30°, 3.3 km² (0.022 %) remain. The affected area is small under these screening criteria, and FADSL is not applied; its effect on the final classes has not been quantified.

Product	Resolution	Method	Limitation for ATES
Copernicus TCD	10 m	Sentinel-2	No tree height, cover only
ETH Global Canopy Height[22]	10 m	GEDI + Sentinel-2	RMSE 6.0 m at $h_{\min} = 5$ m
Meta/WRI CHM v2[23]	1 m	Maxar + DINOv2	Modelled ($R^2 = 0.86$)

Table 3: Satellite-based forest products. Copernicus TCD is chosen as an official European product updated every year; the two canopy height models are modelled, not measured.

On the age of the LiDAR data. Forest that has disappeared since the survey — windthrow, bark beetle — would still suggest a protective effect in the model. This mainly concerns South Tyrol and Trentino, whose surveys predate Vaia; the comparison with Copernicus above points to the affected areas in South Tyrol without quantifying them. A check of old LiDAR forest against a recent satellite product would be suitable for this; it is not part of the pipeline.

4.1.2 Glaciers

Before generalisation, glacier cells are assigned a minimum class of 2. This is an additional product rule, not part of ATES v.2 (Section 4.6.2); its input is a glacier mask, whose sources follow.

The source depends on the region, and where an official inventory exists it takes precedence over OpenStreetMap:

Switzerland swissTLM3D (official land cover, object types "Gletscher" and "Schneefeld Toteis").
Austria Austrian Glacier Inventory GI 5, surveys 2021–2023 [24], 20 subregions as GeoJSON, CC BY 4.0.
France OCS GE 2.0 by IGN, class CS1.2.3 "Névés et glaciers", per département the same edition as the forest mask (Section 4.1).
Other regions OpenStreetMap (`natural=glacier`, ODbL 1.0).

As with forest, the glacier mosaic (18 October 2026) uses an official inventory only within its own region; overlaps of border polygons are split along their centre line, and a cell counts as glacier if its centre lies in a glacier polygon. Of 1,682 km² of glacier, 874.6 km² come from swissTLM3D, 283.8 km² from GI 5, 203.9 km² from OCS GE and 319.3 km² from OpenStreetMap — the last in the regions without an official inventory (Italy, Bavaria, Slovenia; Section 10).

The difference is not academic. In the Galtür test area OpenStreetMap carries 15.37 km² of glacier against 11.67 km² in GI 5, i.e. 32 % more; of 673,154 affected cells, 408,415 lie in both sources, 206,475 only in OSM and 58,264 only in GI 5. The direction matches the finding of Hartl et al.[25], who measure a 31 % loss of area in Austria over two decades: OSM often carries outlines that no longer match today's ice margin. Since the glacier rule **raises** the class, OSM would not be conservative here but simply wrong — it would record melted-out debris as glacier.

For the same reason we use GI 5 (surveys 2021–2023) and not the inventory Huber et al.[9] rely on, GI 4 (2015, Buckel & Otto[26]): at this rate of glacier retreat, outlines six to eight years older differ noticeably from today's ice margin.

4.2 PRA — release areas

Release areas are determined with the fuzzy-logic approach of Veitinger et al. [11], modified for AutoATES v2 (Toft et al.[7]), in four steps.

4.2.1 1. Membership functions

Three terrain variables are normalised to $[0, 1]$ with Cauchy functions:

$$\mu(x) = \frac{1}{1 + \left(\frac{x-c}{a}\right)^{2b}}$$

Variable	a	b	c	Unit of x
Slope	11	4	43	degrees
Wind shelter	2	5	2	**radians**
Canopy cover	40	3.5	−15	percent

Table 4: Cauchy parameters of the pipeline. The same values are hard-coded in the model chain.

The wind-shelter index is in radians, not degrees. The Cauchy parameters are calibrated for that; in degrees the membership would drop to exactly zero for more than half of all pixels and act as a veto through the fuzzy AND. With $c = 2$ far outside the attainable range ($\pm\pi/2$), the "bell" is in effect a monotonic ramp passing through 0.5 at $S_x = 0$; this is exactly what Veitinger et al.[11] (Fig. 3) and Toft et al.[7] (Fig. 2) show.

Toft et al.[7] give 3/10/3 for the wind-shelter membership. We use Veitinger's values 2 /5 /2, which the BFW model chain also uses. The choice matters cell by cell: with 3/10/3, 89,996 of 90,000 compared cells differ from the model chain, although the two variants overlap to 95 % over 101 km² (Jaccard) — an overlap measure over an area says nothing about whether the same cell gets the same membership.

4.2.2 2. Fuzzy AND

The three membership values are combined with the operator of Werners[27]:

$$\lambda = 1 - \min\bigl(\mu_s, \mu_w, \mu_f\bigr)$$

$$\mu_{\mathrm{PRA}} = \lambda \cdot \min\bigl(\mu_s, \mu_w, \mu_f\bigr) + (1 - \lambda) \cdot \frac{\mu_s + \mu_w + \mu_f}{3}$$

4.2.3 3. Slope window and ruggedness

Two steps, both of which **remove** area:

Slope window. Areas with less than 28° or more than 60° slope receive membership zero. Markov et al.[8] name this explicitly as a departure from Toft et al.[7], who assign low but non-zero values there, and refer to Veitinger et al.[11] and Bühler et al.[10].

Ruggedness. The Vector Ruggedness Measure (Sappington et al.[28]) is computed on the 5-m grid: slope and aspect with the Horn method on the elevation model, then the dispersion of the surface normals in a fixed 3×3 window, as in Veitinger et al.[11] and the BFW model chain. Every cell with a value of 0.02 or more has 1 subtracted from its PRA score and therefore drops out of the release area — in effect a binary mask. Huber et al.[9] describe this as a variant that "performed well in comparison with the documented PRAs".

The threshold was determined on a 10-m elevation model, and the measure is scale-dependent: on the same terrain the 3×3 window covers 15 m at 5 m grid spacing and 30 m at 10 m. We compute on the native grid, because the reference chain does and the cell comparison depends on it. Measured over 101 km², the mask removes 9.5 % of the PRA on the native 5-m grid and 11.1 % when computed at 10 m; the threshold is therefore not transferred without effect (Section 9.2.1).

4.2.4 4. Threshold and sieve filter

Pixels with $\mu_{\mathrm{PRA}} > 0.3$ count as release area. The value comes from Hesselbach [29] (§ 4.1.3, HSS optimum 0.36, TSS optimum 0.27, 0.30 chosen for the model **with** the ruggedness mask) and is confirmed by Markov et al. [8]. Without the mask it is not supported.

Finally, connected areas below 1000 m² drop out, and a neighbourhood filter removes single cells and threads: two passes over the mask, first the four direct neighbours (kept from 3), then the four diagonals (from 2).

4.3 Wind shelter

The wind-shelter index quantifies how strongly a pixel is enclosed by the surrounding terrain (Winstral et al.[30], Plattner et al.[31]). Positively sheltered terrain accumulates wind-transported snow and favours slab formation.

The computation follows the AutoATES v2 reference (Toft et al.[7]): for every pixel, the terrain angles $\arctan\left(\frac{z_{\mathrm{neighbour}} - z_{\mathrm{centre}}}{\mathrm{distance}}\right)$ to **all** pixels within a search radius of 60 m are computed (omnidirectional, 360°). The output value is the 50th percentile of all angles, i.e. the median.

That is not Veitinger's choice. Veitinger et al. [11] replace the maximum of the index with the **third quartile** ("using the third quantile instead of the MAX"), so that a single large obstacle — a boulder on an open slope — does not determine the whole value, and compute on a 2-m elevation model. The model chain sets `prob = 0.5`; we follow it because the cell comparison depends on the same number. The direction of the difference can be stated: the median lies below the third quartile, so the index comes out lower and marks less terrain as wind-sheltered.

On the radius: the model chain sets `radius = 6` — a **number of pixels**, i.e. 60 m on a 10-m elevation model. We compute at 5 m and therefore carry the length (60 m, 12 px), not the pixel count; Section 4.7.1 quantifies the effect of conflating physical distance with pixel count. That the radius depends on the resolution and not on the pixel count is published practice: for their resolution comparison Sykes et al.[32] state that they adjusted "the windshelter radius input parameter of the PRA model based on the resolution of each DEM" and left everything else unchanged.

The index itself goes back to Winstral et al.[30], in the form modified by Plattner et al.[31]. The 60 m are used because the model chain uses them; we know of no source available to us that establishes them as an optimum.

Veitinger et al. [11] set the omnidirectional mode (`windtol = 180°`) as default — a particular wind direction is a scenario parameter for event-based analyses, not for a static terrain classification.

4.4 Flow-Py runout

Avalanche propagation is simulated with the Flow-Py energy-line model (D'Amboise et al.[12]). Flow-Py models propagation as a dense-core avalanche (no powder avalanches) by breadth-first search on the DEM raster.

Our runout core is a reimplementation, not a call: the core is time-critical — representing 214 million release cells as a Python object graph with one object per event is computationally infeasible. For this reason there is the reference comparison against com4FlowPy 2.1 (Section 6).

4.4.1 Energy line

For every neighbouring cell the remaining energy is computed:

$$z_{\delta,\text{next}} = z_\delta + \left(h_{\text{current}} - h_{\text{neighbour}}\right) - d \cdot \tan(\alpha_{\text{eff}})$$

Propagation stops when $z_\delta \leq 0$ (energy used up) or the flux falls below R_{stop}.

4.4.2 Flux routing

Mass is distributed to neighbouring cells according to the terrain gradient:

$$\tan\beta = \tan\left(\frac{\arctan(dz/d) + \pi/2}{2}\right)$$

$$w_i = p_i \cdot \tan\beta_i^{\text{exp}}$$

$$\text{flux}_i = \frac{w_i}{\sum w} \cdot \text{flux}_{\text{parent}}$$

Here p_i is the directional persistence and exp the routing exponent. Sub-threshold flux is redistributed to the remaining neighbours (mass conservation).

4.4.3 Parameters

Which values act. In production, every run starts from a release segment (Section 4.5), and every segment carries its own reach angle and velocity limit from its size (Section 4.5.3). These segment values override the scalar defaults of the runout for that run, exactly as in com4FlowPy, where a valid value in the variable rasters replaces the scalar for the start cell. The velocity limit is converted to an energy-height limit with $z_{\delta,\text{lim}} = u_{\text{max}}^2/(2g)$. The scalar values remain in the configuration as a fallback for start cells without valid segment values; since only segments start, the fallback does not act in the product.

Parameter	Value	Source
Active in production, per release segment		
α — reach angle	α(size), 23–33.3°	ATI model chain, Section 4.5.3
$z_{\delta,\text{lim}}$ — energy-height limit	$u^2_{\text{max(size)}}/2g$, 9.3–215 m	ditto
exp — spreading exponent	8, fixed	ditto; the size-dependent raster is not read
R_{stop} — flux threshold	0.0003	Huber et al.[9, Tab. 1]; Markov et al.[8, Tab. 2]
$\Delta\alpha^{\text{for}}_{\text{max}}$ — max. friction increment	21°	ATI configuration; Markov et al.[8, Tab. 2]
$\Delta\alpha^{\text{for}}_{\text{min}}$ — min. friction increment	0°	ditto
$v^{\text{for}}_{\text{lim}}$ — velocity limit of the forest effect	45 m/s	ditto
Scalar fallback, not acting in the product		
α	26°	Markov et al.[8, Tab. 2]
$z_{\delta,\text{lim}}$	185 m	ditto; ATI configuration

Table 5: Runout parameters. The ranges follow from the smallest segment (500 m², size 1.57) and the upper bound size 5. All values are taken from the configuration of the pipeline when this document is generated.

The energy height is the classical energy line of runout computation [33]: a block accelerating from height z_δ without friction loss reaches $v = \sqrt{2gz_\delta}$. The velocity limit of a segment is therefore not a second assumption alongside its energy-height limit but the same one in different units.

The scalar values in the literature. The scalar energy-height limit is given differently in three places:

Source	$z_{\delta,\mathrm{lim}}$	Area
Huber et al.[9], Section 2.3 there	270 m	Eastern Alps
ATI configuration of the AutoATES workflow (BFW)	185 m	default of the chain
Markov et al.[8], Tab. 2	185 m	Pirin

Table 6: The scalar energy-height limit in the three sources. With otherwise identical values ($\alpha = 26°$, exp $= 8$, $R_{\mathrm{stop}} = 3 \times 10^{-4}$), Huber alone gives 270 m.

These values belong to a runout with one global parameter set. In the size-dependent chain they are only the fallback, and the question which of them is right does not affect the product: there, the limit of each run follows from the size of its segment. We keep 185 m as the fallback because the ATI configuration sets it.

4.4.4 Forest effect

Forest raises the friction angle α on forested cells; the increment scales linearly with PFC/100 and decreases with the energy height: at $z_\delta = 0$ the full increment acts, above a threshold none. The threshold is configured as a velocity, $v_{\mathrm{lim}}^{\mathrm{for}} = 45\,\mathrm{m/s}$, and com4FlowPy converts it to an energy height as $z_\delta = v^2/\left(\sqrt{2}g\right)$ (`flowClass.py`, lines 113–114), not with the $2g$ of the energy line used elsewhere in the model. We reproduce this conversion exactly, so the forest effect ends at $z_\delta \approx 146$ m. Three quantities are to be kept apart: the configured velocity (45 m/s); the energy height the code derives from it and compares against, which is also the threshold that acts in the product (146 m); and the physical reading of these numbers with the energy line ($2g$), under which 146 m would correspond to 53.5 m/s and 45 m/s to 103 m.

Detrainment is switched off. com4FlowPy has two forest modules — `forestFriction` (friction only) and `forestDetrainment` (friction plus mass removal). `forestFriction` is used, as in the BFW model chain. The detrainment values remain in the configuration so that a comparison against the other module stays reproducible; they act only when it is selected explicitly.

How strongly the forest acts. Huber et al.[13] study the forest-friction module with two parameter sets on generic terrain: with 0–11° and 30 m/s it has "almost no effect on modeled runouts"; with 0–23° and 45 m/s the runout is reduced markedly where the forest structure index is 0.6 or higher. The values of the model chain, which we use (0–21°, 45 m/s Markov et al.[8, Tab. 2]), lie close to the second set.

Path length for the forest effect. The friction increment is only applied once a run has a three-dimensional path length greater than zero. com4FlowPy 2.1 computes this length with the north–south step twice instead of the east–west and north–south steps (`flowClass.py`, line 218), so that a horizontal east–west step without elevation change counts as zero length. With a skip distance of 0 m, as configured, this matters only for flat east–west steps directly next to a release cell, where the forest friction is skipped. AvaFrame merged a correction on 14 October 2026 (PR #1342, intended for release 2.2, which had not been published by 21 October 2026). We compute the length correctly; this is the only point at which our runout deliberately departs from com4FlowPy 2.1 (Table 12, Section 6).

4.4.5 Size-dependent parameters

α and the velocity limit are **not** global but determined per release area from its size; the equations are in Section 4.5.3. The runout reads two rasters from the mobility stage for this.

4.4.6 Absorption

com4FlowPy has a switch `previewMode` that absorbs incoming flux at a cell already reached. The reference runs with it **off**, and so do we. With absorption the runout angle γ is underestimated by as much as 35°, and classes change only downwards.

4.4.7 Start mask

Runs start from the **segments** (Section 4.5), not from the cleaned release mask: the BFW model chain builds the start mask for com4FlowPy from `praAreaM > 0`, i.e. after the size filter of the segmentation. Remnant areas below 500 m² have no size and therefore no mobility parameters; they do not start.

The difference is not small: starting from the cleaned mask, the remnant areas — 3.5 % of the release cells in the reference area — would run with the global α and raise the class of 1.27 million cells. The comparison against the full workflow covers this handover (Section 6.2).

4.4.8 Output

Runout angle $\gamma = \arctan(\Delta h / \text{ flow distance})$ per pixel, as the maximum over all paths reaching the cell. Cells not reached remain empty. A release cell that no other run reaches carries $\gamma = 0°$ — its own run at distance zero. This is how com4FlowPy writes it (`nodata` only where `cellCounts = 0`), and the classifier reads 0° as reached.

4.5 From the release mask to release segments

Between the release mask (Section 4.2) and the runout lie three steps. They divide the release area into segments and give each segment its own mobility parameters. The runout is then an ensemble of start-cell-based propagation computations with segment-dependent parameters (Section 4.4). FlowPy is a conceptual reach and routing model; its routing flux is not an avalanche mass, and the assumed release depth serves only to parametrise the mobility, not to claim a volume-true dynamic avalanche simulation.

4.5.1 Subcatchments

The mask is broken down into hydrological subcatchments: filling sinks, D8 flow direction [34], accumulation, stream network, stream segments, catchments, and a majority filter to smooth the boundaries. Streams begin at 5 ha of upslope area, and smoothing uses a window 10 px wide. The reference configuration also carries a minimum stream length (`minLength` = 100); it enters no computation step of the reference chain — it is only part of the names of intermediate files — and it does not in ours either.

This chain is standard hydrology and is **called**, not reimplemented: the same WhiteboxTools binaries that the reference chain uses (Section 5.2). What remains our decision is the order and which intermediate raster goes into which step — that accumulation receives the **filled elevation model** and not the flow direction is a choice no tool complains about. Exactly such decisions are what the comparison checks (Section 6).

4.5.2 Segmentation

A **segment** is a connected piece of release area within one subcatchment. The same area across two gullies thus becomes two segments, which become two avalanches and not one large one — that is the whole purpose of the previous step. Segments below 500 m² drop out.

The reference intersects polygons for this; we intersect the two label rasters. The result is the same (measured: no differing cell, identical areas), without a geometry library in the path.

4.5.3 Mobility from size

From the area A_{rel} of a segment and the assumed release depth of 1 m, a continuous **size** is determined, and from it the runout parameters for exactly this segment:

$$\text{size} = 2 + \log_5\left(A_{\text{rel}} \cdot d \cdot 10^{-3}\right)$$

$$\alpha(\text{size}) = 32° - (\text{size} - 2) \cdot 3° \quad u_{\max}(\text{size}) = 20 + (\text{size} - 2) \cdot 15 \quad \exp(\text{size}) = 75 \cdot 0.64^{\text{size}}$$

The size is capped at $\text{size} = 5$, and $u_{\max}$ is bounded below at $u_{\max} \geq 5$ m/s. A larger release area therefore gets a flatter friction angle and a higher velocity limit and runs further — the size dependence that a global α cannot represent.

The runout reads α and $u_{\max}$ as rasters (`variableAlpha`/ `variableUmax` on), but not the spreading exponent: stage 2 produces a raster for it as well, which the runout discards (`variableExponent = False` in the reference configuration), computing with $\exp = 8$.

Which scale is "size"? The anchor sits on the EAWS avalanche size[35]: with a release depth of 1 m, the formula gives exactly size 2 for 1,000 m³, which is the EAWS guide value for the medium avalanche relevant to skiers. Above that, however, the two scales diverge, because the formula rises in steps of five and EAWS in steps of ten:

Size	Formula	EAWS guide value
1	200 m³	100 m³
2	**1,000 m³**	**1,000 m³**
3	5,000 m³	10,000 m³
4	25,000 m³	100,000 m³
5	125,000 m³	> 100,000 m³

Table 7: The size measure of the mobility stage against the EAWS guide values. Equal at the anchor, different in step width.

The upper bound $\text{size} = 5$ thus corresponds to 125,000 m³ and not to EAWS size 5. **"size" is therefore not an EAWS measure** but a scale of its own, anchored to it at size 2 — which is essential for interpreting the mobility parameters and not obvious when reading the formula.

What the scalar fallback values mean in this scheme helps to contextualise the parametrisation: $\alpha = 26°$ corresponds to size 4.00, $z_{\delta,\text{lim}} = 185$ m via $u_{\max} = 60.2$ m/s to size 4.68 and $\exp = 8$ to size 5.01. The three fixed values that appear in the literature as **one** parameter set therefore do not lie on the same avalanche size — one reason why the size-dependent chain replaces them per segment.

4.6 Classification

The ATES class results from two independently computed input classes, which are combined by their maximum and then mapped to ATES 0–4 together with canopy cover and release status through a lookup table (Toft et al.[7]). The implementation follows `autoATESClassifier` line by line (Section 6).

Slope is **truncated towards zero** to whole degrees before it is tested (`slope.astype("int16")` in the reference), and the thresholds are tested with "greater than". On the untruncated slope this gives intervals that differ from the threshold values by one degree; the table gives both. The runout angle is not truncated and is tested with "greater than or equal". Class 4 is tested against a smoothed slope; the window is 1 pixel, so the smoothing has no effect — this matches the reference configuration, and we follow it.

Input	Class 1	Class 2	Class 3	Class 4
Slope: threshold	$> 28°$[†]	$> 28°$	$> 39°$	$> 45°$
Slope: effective interval	—[†]	$29° \leq s < 40°$	$40° \leq s < 46°$	$s \geq 46°$
Runout angle	$\gamma \geq 0°$	$\gamma \geq 26°$	$\gamma \geq 36°$	—

Table 8: Classification thresholds and the intervals in which they act on the untruncated slope s. [†] Class 1 cannot be reached via slope, because SAT01 = SAT12 = 28°. The intervals are covered by a test through the production path, with values directly below, on and above every boundary.

Overhead exposure and cell count are not part of the classification. Huber et al.[9] (Tab. 3) do not use them, and we follow that. This sacrifices the ability to detect channelled terrain — gullies in which several avalanche paths converge. That is a real loss, but the same one the Alpine calibration accepts.

4.6.1 Forest and release area

The base class is combined with the forest code (open/sparse/mixed/dense from canopy cover, limits 10 /25 /65 %) and the release status through a table of 40 entries. Dense forest lowers the class, release raises it.

Two peculiarities of the reference that we adopt: **in the runout, forest does not count** — where a modelled avalanche arrives, canopy cover is set to "open", because the forest demonstrably did not stop it. And **class 4 in forest becomes class 3**.

4.6.2 Glaciers

An additional product rule, not ATES v.2. Before generalisation, glacier cells are assigned a minimum class of 2. The rule comes from the original ATES (Statham et al. [1]), in which all glaciated terrain defaulted to at least Class 2, and Huber et al. [9] introduce it as a binary glacier layer "to disallow simple terrain on glaciers"; in the released BFW model chain the corresponding code is switched off. ATES v.2 has removed glaciation as an input parameter (Statham & Campbell [3], Section 3.3), because the rule also raised flat and low-angled glaciers without avalanche terrain to Class 2; crevasses remain only as a terrain-trap consideration.

We keep the rule as a product heuristic for a touring map, on which glacier travel with crevasses and ice breaks is part of the terrain the user has to judge, and we name it as such. It raises the class of low-angled glacier terrain that ATES v.2 would rate Class 1 or 0; the class shown on a glacier is therefore not a pure snow-avalanche rating. The effect on steep glaciers is small, because they are reached by runout anyway — in the reference area the rule changes 10 of 520,262 glacier cells; for the Alpine product it has not been quantified.

The generalisation runs after the rule and can still remove a glacier patch smaller than 1000 m² that is surrounded by value 0. In the Alpine product this leaves 4,071 of 67.3 million glacier cells (0.10 km², 0.006 %) in value 0 or class 1; all other glacier cells carry class 2 or higher.

4.6.3 Model value 0 — no modelled exposure

ATES v.2 (Statham & Campbell[3]) carries Class 0 (Non-avalanche) as an **optional** fifth class: "the use of Class 0 is optional due to the reliability needed to make this assessment; otherwise, Class 1 includes Class-0 terrain", and its determination "requires high confidence in the assessment and can have little to no uncertainty".

The automated methods consistently do without it. Toft et al.[7] give liability as the reason — "there is significant liability associated with deeming an area safe from avalanche hazard" — and too little validation data. Markov et al.[8] explicitly assign Class-0 terrain to Class 1. Huber et al.[9] (Tab. 3) have no 0/1 boundary.

What our value 0 means. The product carries a value 0 for terrain in which **this model detects no avalanche exposure**: slope below 29° (Table 8) that no modelled runout reaches. It is derived from the same terrain model as classes 1–4 and carries the same limitations. **It is not a verified ATES Class-0 assessment**, which would require high confidence and little uncertainty; throughout this document it is called "value 0", not "Class 0" or "Non-avalanche".

Why the value is kept. Whole countries are mapped including lowlands, not Alpine study areas. In the Alpine product 68.7 % of the assessed area carries value 0, in Switzerland 59.6 %. Of the terrain below 600 m, 95.1 % (Switzerland: 96.2 %) carries value 0, and 58.7 % of the value-0 area lies below 600 m (Switzerland: 42.3 %). Merging it into class 1 would draw all of this terrain green as "Simple".

Display. The map does not draw value 0; it is transparent. This is a deliberate product decision: on a touring map of the Alps, a colour over every valley floor and the whole foreland would carry no information for route planning and would unsettle rather than inform. Transparent therefore means "no avalanche exposure detected by this model", not "safe"; the map legend has to say so.

Not assessed. The product stores value 0 and "not assessed" separately: nodata marks only cells for which no slope can be computed — outside the DEM, in DEM gaps and on open water. The reference implementation writes its class 0 as nodata, so that terrain without exposure and terrain without data share one value; every assessed cell receives the same class as in the reference, only the encoding differs (Table 12). autoATES v3.0 (Sykes et al.[14]) separates the two in the same way.

In the app, the two are not distinguished: the map leaves both transparent, and the terrain service returns no class for either, so that neither sets a lower bound in Tour-Check. We checked whether this matters on the routes the app offers. Of the 13,436 winter routes of the app (Switzerland, Austria, France, Bavaria), sampled every 5 m — 14.9 million points — **none** lies on a cell that was not assessed. Within the Alpine region, cells not assessed are water surfaces and the Lessinia in Veneto (Section 7). For Italy and Slovenia the app offers no routes; there the check has not been made.

4.6.4 Generalisation

Connected areas below 1000 m² are removed and assigned to the surrounding class — Toft et al.[7]: "removing clusters smaller than a specified area and incorporating it to the surrounding class". The filter comes from AutoATES v1.0; Sykes et al.[32] determined 30,000 m² as optimal for the resolutions tested there, the Alpine calibration uses 1000 m².

Because the surroundings are often value 0, the filter cannot reclassify a small steep spot but only remove it. A protection against this was tested and rejected: in the reference area it recovers 104,079 cells, spread over 16,185 patches with a median of 75 m², all isolated. The result would be isolated points in empty terrain. Moreover, the PRA stage already sieves the same patches out of the release areas (1000 m²), so no avalanche is computed there at all.

4.7 Alpine calibration

We adopt the BFW parametrisation, published by Huber et al.[9] (Tab. 1 and 3) for the Eastern Alps and continued by Markov et al.[8] (Tab. 2) — the latter for the Pirin Mountains, with the same authors and Eastern Alpine starting values but adapted to the terrain there. The column "Eastern Alps" below therefore denotes **our** values and their respective source, not a single Eastern Alpine publication. The Canadian comparison values come from Toft et al.[7] (Tab. 1 and 2), who refer to Sykes et al.[32] for them.

Parameter	**Canada**[7]	**Eastern Alps**	**Reason for the difference**
SAT — slope	18/28/39°	28/39/45°	slab avalanches in the Alps from about 28°
AAT — runout angle	24/33°	26/36°	calibrated against Tyrolean avalanches
PRA threshold	0.15	0.3	fewer false alarms in the Alps
ISL_SIZE	30,000 m²	1000 m²	finer resolution
Canopy cover (TREE)	20/55/75 %	10/25/65 %	Alpine treelines, denser stands
α	25°	per segment, 32° at size 2	size-dependent mobility (Section 4.5.3)
$z_{\delta,\mathrm{lim}}$	270 m	per segment, from $u_{\max}$	size-dependent mobility; scalar fallback 185 m (Section 4.4.3)

Table 9: Eastern Alpine parametrisation against the Canadian recommendations. Our values from the configuration of the pipeline.

On the SAT row: Toft's Table 1 has three limits (SAT12/23/34 = 18/28/39°) and **no** 0/1 limit; a value of 15 found in the source code is a default that the publication does not support. Likewise for AAT: Toft's table gives AAT12 = 24° and AAT23 = 33°, no AAT1. The model chain contains an `AAT1 = 18` whose line in the classifier is commented out and has no effect.

The Cauchy parameters are identical to the model chain and are hard-coded there (Section 4.2.1). The flux threshold $R_{\mathrm{stop}} = 0.0003$ follows Huber et al.[9, Tab. 1] instead of the Flow-Py default (3×10^{-3}).

4.7.1 On resolution

The Eastern Alpine parametrisation was determined on a **10-m** elevation model (Huber et al.[9]; Markov et al.[8] explicitly name 10 m as the input resolution of the PRA model). We compute at **5 m** because our elevation and forest data support it — Markov's choice of 10 m is justified by elevation models from digitised military maps and a forest product with 10 m grid spacing, neither of which is our situation.

Slope. Slope is computed once from the 5-m elevation model with the Horn method (3×3 window, edges computed), without smoothing; the same slope enters the release areas, the classification and the ruggedness. That the PRA has been used at 2 m (Veitinger et al.[11]) and at about 10 m (Toft et al.[7]) does not show that the whole model chain with this parametrisation is suitable at 5 m; that has not been tested.

Three parameters are defined in pixels or cells and therefore mean something different at a different grid spacing. We carry them as physical quantities and convert them to pixels:

Quantity	at 10 m	This implementation	Basis
Wind-shelter search radius	6 px = 60 m	60 m	Sykes et al.[32]
Stream threshold	500 cells = 5 ha	5 ha	area threshold
Smoothing window	5 px = 50 m	50 m	window length

Table 10: Quantities defined in pixels, converted to 5 m.

Sykes et al.[32] are explicit about the search radius: they adapt it to the resolution of each elevation model and keep all other parameters fixed. It is the only one they treat this way.

What cannot be converted: the slope thresholds, the PRA threshold and the Cauchy parameters are thresholds on quantities that themselves depend on resolution — finer resolution can preserve steeper local terrain features. Sykes et al.[32] determine these values per area and elevation model by a grid search against expert maps; they find that the slope parameters have the largest effect of all. Without a reference map this cannot be reproduced, and we keep the values unchanged. See Section 9.2.1.

5 Implementation

The model is a Python pipeline and processes regions of any size. Compute-intensive steps are compiled with Numba and run in parallel by chunk or strip.

Relation to the reference. The implementation aims to preserve the reference behaviour; known departures and the scope of the equivalence tests are documented in Section 6. The runout core keeps its state per event instead of per raster position, including the directional persistence summed over all parent cells; apart from the corrected path length of the forest effect (Section 4.4), it agrees with com4FlowPy 2.1 on the tested inputs.

5.1 Verification strategy: two checks, two questions

The comparison in Section 6 checks the **computation**: it runs every stage in memory on small extracts and compares every cell with the reference output, and so demonstrates agreement with the reference implementations for the tested inputs and configurations. What it does not touch is the **production path** — writing COGs, cache, stamps, manifest, generalisation. That is where it is decided whether a correct computation also becomes a file.

The distinction matters because a broken intermediate raster can produce a plausible-looking result: a wind-shelter raster interrupted while being written carries a header and overviews but nodata at full resolution, its parameter fingerprint still matches, and the chain continues with an empty release mask — and since the classification takes $\max(\text{slope class}, \text{runout class})$, the result is a pure slope map. A fingerprint says **with which parameters** something was computed, not **whether** anything is in the raster. Two checks cover this:

Assertion per stage every stage states what it must deliver — grid, type, nodata, value range, and whether a **dense** raster carries values at all. The pipeline checks against this right after building, so an invalid intermediate product cannot propagate into subsequent stages. Sparse rasters (release areas, segments, runout) are exempt: there, extensive nodata is normal.

Acceptance test before every production run the full production path on 320 km² of Kandersteg–Blümlisalp with an empty cache, about half an hour. The area is chosen so that

everything the chain must distinguish occurs — forested valley sides, treeline, rock, glaciers and runout paths down to the valley. Every product is checked against its assertion, plus what must occur on **this** terrain (release cells, reached cells, more than one class) and the completeness of the manifest. An empty stage is an error here, not a terrain finding.

Two further checks secure the intermediate states from which a run resumes. The **stamp check** recomputes per stage the fingerprint that the current code would expect and holds it against the stamp of the file — read-only, so that a stale input is found before a run discards and recomputes everything that depends on it. The **integrity check** detects files left behind by an abort: a tiled GeoTIFF that was never closed properly is larger than the sum of what its directory describes, even when its dimensions and progress file look correct.

For forest a principle of its own applies: **a missing source aborts, it does not become 0 % forest.** 0 % is a valid value, and a raster with silent failures looks like a raster with clearings. Where tiles are fetched one by one — France, Bavaria, Slovenia —, the pipeline retries failed requests with growing pauses and aborts afterwards; in Trentino it aborts if more than 0.1 % of the province remains without canopy height. Austria reads the nDSM from 50-km tiles of BEV; there too the run aborts if the tile search gets no answer or a strip cannot be read. In addition, all forest rasters were reviewed for silent failures (blocks without forest against the forest mask, straight edges between forest and 0) before the forest mosaic was built.

Checks of the delivered product. The Alpine product passes through a further chain of checks, each of which can stop it:

Forest mosaic the hard checks of Section 4.1.1.

Tiles every tile records code, inputs with checksums and duration; the edge check over all cut lines must be clean before the tiles are assembled (Section 5.3).

Copies the product is identified by its checksum; every copy — the archive next to its tiles and the copy that the terrain service reads — is compared against it, and a copy counts as complete only after its checksum has been recomputed where it is stored. Nothing is overwritten; a new version is a new directory.

Map tiles generated from the same raster, z10–z14, by nearest neighbour with exactly the colours of the app, because Tour-Check reads the class back from the pixel colour; a test builds real tiles and fails on any colour outside the five allowed (four classes and transparent). The tile tree is identified by a checksum over all files, recomputed on the server.

Delivery it refuses a product built from code that is not under version control, one whose checksum no longer matches its manifest, and one without a clean edge check.

5.2 Use of external tools instead of reimplementation

The hydrology — filling sinks, flow direction, accumulation, stream network, subcatchments — runs in WhiteboxTools, the same binary that the model chain calls. A separate check of a reimplemented algorithm is therefore not needed; how the tools are called — arguments, preprocessing and the handover between steps — can still differ and is covered by the end-to-end comparison (Section 6.2).

Only two things are reimplemented, and both out of necessity. The **runout core** is time-critical: representing 214 million release cells as an object graph with one object per event is computationally infeasible. And the **majority filter** that smooths the area boundaries has to run in strips — in one piece it needs 26 GB for Switzerland and 73 GB for Austria.

One property of the reference matters for the majority filter: it breaks ties by the iteration order of a hash table, which is random per process. Two runs of the same binary on the same raster therefore differ from each other (measured: 33 cells). Our version is deterministic and takes the smallest value

in a tie. Differences therefore occur only where there is no majority — in the reference area 9 of 28,090,000 cells.

5.3 Computation in tiles with a margin

The Alpine grid has 23.9 billion cells. One run over all of it would take days in one piece, and an interruption or a code change would invalidate everything. It is therefore computed in **tiles**, each of which sees the same complete inputs:

Grid fixed cores of 100 km × 100 km on the grid of the Alpine DEM, 45 of which contain terrain. Each core is computed with a **margin** of 10 km on all sides, cut from the same Alpine DEM, forest mosaic and glacier mosaic, and only the core is kept.

Why 10 km the runout halo of the pipeline is 3500 m beyond that, the subcatchments (Section 4.5.1) act across valleys and enter the result through segments and mobility.

Assembly the 45 cores are placed side by side without overlap and written as one raster. Every cell comes from exactly one core.

Two checks belong to this, and they answer different questions.

Is the map seamless? At each line where two cores meet, the mosaic takes one tile on one side and the other tile on the other. The **edge check** compares both full tiles in a band of 1 km on either side of every such line: a cell at the line itself where the two disagree would be a visible jump. Result for the Alpine run: **73 of 73 cut lines without a jump and without a gap.** Deviations next to a line occur at 8 of the 73 edges, the nearest 30 m from the line; they do not enter the mosaic, because there the respective core applies. Assembly refuses to run unless this check is present for exactly the cores used and every edge is clean; the result is recorded in the manifest.

Is every tile the same as one run over the whole Alps? For this, the **ring check** compares neighbouring tiles in the ring between 3500 m and 10 km from the outer edge of one of them, where both have complete surroundings. The design assumption was that the two agree there cell for cell. **They do not, completely:** in 31 of 146 comparisons, 23,368 of 4.86 billion compared cells (4.8 per million) differ. The differences are attributed to tile-dependent subcatchment delineation, as in Section 6.5; this has not been isolated by a separate test. The subcatchments are formed within each tile, and a catchment reaching across the margin is cut differently in two tiles. A run over the whole Alps in one piece would form them once; with the memory demand of the hydrology (33 bytes per cell, 790 GB) that is not feasible. The ring differences are therefore of the same nature as the tiling within a run (Section 6.5). Their effect on the classes of a single run over the whole Alps has not been measured.

Forty-two tiles were computed with the code of the release and carry its parameter fingerprint. Three were computed shortly before, with two earlier code states. The differences to the release lie only outside the model: in the construction of the forest mosaic and in the catalogue and border handling of the regional DEMs, none of which is read by the tile computation. Every tile records the checksums of its inputs. The three tiles used an earlier forest mosaic; within the extent of each of the three tiles, the earlier and the current forest mosaic are identical cell by cell (576 million cells each, none different), and DEM and glacier mosaic are the same files. The code identifiers and the change set can be provided on request. That the product can be assembled again byte for byte from the stored tiles shows only that the assembly is repeatable, not by itself that the tiles are equivalent; the equivalence rests on the code and input comparison above.

5.4 Memory

Every step is computed in pieces because the raster sizes require it: a tile of 120 km × 120 km has 576 million cells, the whole Alpine grid 23.9 billion.

WhiteboxTools loads every raster completely, measured at 33 bytes per cell — 19 GB for one tile, 790 GB for the Alps. The hydrology therefore runs tiled within each tile as well, with a halo (2048-px pieces, 1600-px halo); areas across piece boundaries are merged via their outlet, not via their number. Island removal runs in strips.

This affects the result and is therefore measured (Section 6.5). The tiling cannot be bit-identical: flow accumulation needs the full upslope area of a cell, and that is unbounded.

Three kinds of settings have to be told apart:

Execution settings number of workers, chunk size of the runout, strip size, cache sizes. They are derived from the cores and the memory of the machine and are **not** part of the parameter fingerprint, because they do not change the result: two machines with different values must deliver the same raster, and tests check this. For the runout this holds by construction — every release cell runs exactly once, and the partial results are merged by maximum.

Decomposition settings within a run that can change the result the tile and halo size of the hydrology (2048 px, 1600 px) and the runout halo (3500 m). They are part of the parameter fingerprint; their effect is measured in Section 6.5.

The tile grid of the Alpine product core and margin of each tile (100 km, 10 km). They are recorded per tile as pixel windows in the provenance of the product; their effect is described in Section 5.3.

6 Verification against the reference implementations

This chapter is about **verification**: does the implementation compute the specified method? Agreement with a reference implementation checks the implementation, not whether its terrain classification is right; that is a separate question (Section 8). The pipeline is a reimplementation, and whether it reproduces the model is measured against the authors' source code, not against a description of it.

Two references, both **called**, not copied:

BFW model chain Avalanche Terrain Identification (ATI [15]), repository `OpenNHM/AvalancheTerrainIdentification`, commit `1326bce` of 1 September 2026 (after release v1.1), EUPL-1.2, with its configuration for the AutoATES workflow — for release areas, hydrology, segmentation, mobility and classification.

com4FlowPy from AvaFrame, release 2.1 (tag `2.1`, commit `9c3cefb8`), for the runout and the stage-by-stage comparison. The end-to-end workflow of ATI requires an interface of AvaFrame that release 2.1 does not have; ATI names a development branch that no longer exists. It was run with the AvaFrame main branch at commit `40a15f53` of 14 October 2026, which also contains the corrected path length (Section 4.4).

On 21 and 22 October 2026, on the code of the release, all checks passed: 153 stage tests and 27 end-to-end tests against ATI, 239 tests against com4FlowPy, and the acceptance run on the production path (Section 5.1).

6.1 Two questions, two runs

The comparison separates two things that are easily confused.

Does our code compute what theirs computes? For this both sides receive the same parameter set. Both pipelines process identical terrain inputs at every stage, and the comparison requires exact agreement, `(ours != theirs).sum() == 0`, plus the same type and shape — except for the documented tie-breaking difference in the majority filter, which in the reference area affected nine area-identifier cells and no final terrain class (Section 9.2.9). Label rasters (area and segment identifiers) are compared as a **partition** — the same cells occupied, the same number of pieces, a bijective mapping. That is no relaxation: if a boundary lay differently, the mapping would no longer be unique.

Test area	cells compared	differing
4 areas of 5 km × 5 km, stages 1.1–1.4, 2, 4	96,000,000	**0**
Reference area 700 km², same stages	169,000,000	**9**
Runout against com4FlowPy, 4 areas	4,000,000	**0**

Table 11: Cell comparison against the reference implementations.

The nine cells in the reference area all lie in windows where the majority filter finds no majority and the choice is arbitrary (Section 9.2.9). On the four test areas the runout angle agrees to within 0.00°, with the same set of reached cells.

Are our parameters theirs? The first run does not answer this: it feeds our values to both sides, so a wrong value would arrive at both and both would compute the same wrong thing. The second run reverses this and uses the released configuration on their side. Every difference must then be explained by a named parameter: the stage is recomputed with **their** value and must match their result completely.

In addition, every parameter is checked individually against their configuration file — read, not copied, because a copied expectation would cement the same error in thinking.

6.2 End-to-end workflow without intermediate-input injection

Do we hand over the same things? The two runs do not answer this question either. They call the reference stage by stage and give each stage **our** inputs — our start mask to the runout, our runout angle to the classification. What is established is "same input, same computation", not "same handover". Settings that the comparison sets for both sides — such as absorption (`previewMode`) — and the handover between stages — such as the start mask and the angle in release cells — are not covered by it.

The third run therefore calls the released workflow (`runAutoAtesModelChain.py`) **as a whole**, with nothing but the elevation model and forest, and our pipeline through its production path. The comparison is made at every interface: raw and cleaned release mask, segment areas, the three mobility rasters, the runout angle in **every** cell and the class before and after island removal. On three terrains (narrow valley, high mountains, forested valley floor) everything is cell-identical except `SAT01`; if that is set equal, the finished map is too. At scale — 702 km², all tile seams of the hydrology — release, segments, start mask and mobility are cell-identical. The reference does not compute the runout itself on this area in reasonable time; it is established through the cell comparison with com4FlowPy for the same inputs.

6.3 Documented deviations from the reference configuration

Seven departures from the released configuration, each deliberate and documented: four parameters (`SAT01` and the three resolution-related quantities), the glacier rule, which is not a parameter, the

encoding of value 0 (Section 4.6.3), which changes no class, and the path length of the forest effect, which follows a correction of com4FlowPy (Section 4.4). The three resolution-related quantities are given in **pixels** because that is exactly where they diverge — physically they are equal, and that is the purpose of the correction:

Quantity	Reference	This implementation	Rationale
`SAT01`	0	28	This implementation keeps a model value 0 below Class 1; neither reference does. ATES v.2 lists Class 0 as optional; Toft et al. (2024) omit it for liability reasons, and Markov et al. (2025) merge it into Class 1. Because entire countries are mapped, including lowlands, the value is kept but not rendered. It means no avalanche exposure detected by this model, not a verified ATES Class-0 assessment.
Wind-shelter search radius	6 px at 10 m = 60 m	12 px at 5 m = 60 m	The reference setting radius = 6 was determined on a 10 m DEM and therefore represents 60 m. Transferred as a pixel count, it would shrink to 30 m at 5 m; the physical length is preserved instead.
Stream threshold	500 px at 10 m = 5 ha	2,000 px at 5 m = 5 ha	The reference threshold of 500 cells represents 5 ha of upslope contributing area at 10 m. Retained as a cell count, it would drop to 1.25 ha at 5 m and yield a drainage network four times as dense.
Smoothing window of the subcatchments	5 px at 10 m = 50 m	10 px at 5 m = 50 m	The same conversion preserves the width of the majority-filter window.
Glacier rule	commented out in the released code	enabled (minimum ATES 2)	Huber et al. (2023) describe it explicitly: a binary glacier layer that prevents glaciers from being classified as simple terrain. In ATI's released code the block sits inside a string literal and has no effect. ATES v.2 has removed glaciation as an input. This implementation keeps the rule as an additional product heuristic, applied before generalisation.
Value 0 in the output	written as nodata (−9999)	value 0; nodata only where not assessed	In the reference output, its class 0 and missing data share one value, so terrain without exposure and terrain without data look the same. Here value 0 is stored as 0 and nodata marks cells that were not assessed (no slope, hence no DEM). Every assessed cell receives the same class as in the reference; autoATES v3.0 separates the two the same way.
3D path length (forest)	$\sqrt{(2\ dy^2 + dz^2)}$ in com4FlowPy 2.1	$\sqrt{(dx^2 + dy^2 + dz^2)}$	The 3D distance from the release cell, which decides where forest friction may act, uses the row offset twice in com4FlowPy 2.1. AvaFrame merged the correction on 14 October 2026 (PR #1342, intended for 2.2, not yet released on 21 October 2026); this implementation follows the correction. With a skip distance of 0 m the effect is limited to flat east–west steps directly below a release cell.

Table 12: Deliberate departures from the released configuration of the model chain, taken from the configuration of the pipeline.

Three of them arise from the same cause — quantities defined in pixels or cells whose physical meaning shifts with the grid spacing (Section 4.7.1). Their combined effect in the reference area: 3.85 % of the cells change class, with a net shift of 0.44 % of cells to the next lower class. Corrected individually it would have been more than twice as much, because the stream threshold acts opposite to the search radius.

6.4 What the comparison does not cover

The tiled hydrology is not part of it. The comparison runs the untiled path; it says nothing about the tiling. That is measured separately (Section 6.5).

Stage 4 is not comparable in the second run. It receives the runout angle from our chain, while the preceding stages of the reference run with their values — a mixture of two parameter sets for which there can be no explanation. The first run covers the stage exactly.

The comparison does not see the glacier rule, because the test runs receive no glacier raster. It is named in Section 4.6.2.

The path length of the forest effect is compared against com4FlowPy 2.1, which still has the error corrected in AvaFrame PR #1342. The runout test areas pass nevertheless: in them the correction changes no cell (the case is a flat east–west step directly next to a release cell in forest). The departure is therefore documented (Table 12) but not covered by a cell comparison; that will be possible once AvaFrame publishes the correction.

6.5 Effects of tiled hydrological processing

The stage that forms the subcatchments runs tiled because in one piece it would need 101 GB of memory (Section 5.2). It cannot be bit-identical to the untiled run. The question is therefore not **whether** it differs but **how far** — and the measure for that is not α but the ATES class: it is the product by which someone decides where to ski.

Measured over twelve random terrain pieces of 8192^2 cells (41 km side), tile 2048, halo 1600 px. Share of segment cells whose α differs by more than the threshold:

Area	Relief	Segment cells	$> 0.1°$	$> 0.5°$	$> 3°$
11	3778 m	25.3 M	4.32 %	2.54 %	0.359 %
12	3048 m	17.2 M	0.40 %	0.24 %	0.037 %
9	1549 m	4.95 M	0.19 %	0.14 %	0.014 %
Median over eleven			0.019 %	0.003 %	0

Table 13: Difference in friction angle between tiled and untiled hydrology. Four of the eleven usable areas are bit-identical.

The magnitude of the difference increases with relief: bit-identical up to 1879 m, differing from 1503 m, and the worst area, at 3778 m, is the highest. That is precisely the terrain ATES is about — which is why the median is of little use here. The distribution has a long tail, and the high mountains lie on it.

Almost none of it reaches the class. The same area 11, computed once with and once without tiling through stages 1.4, 2, 3 and 4:

Area 11, 3778 m relief	Share
α differing: 1,952,451 of 25.3 M segment cells	7.71 %
class differing: 6,284 of 54.7 M terrain cells	**0.0115 %**
of which higher 2,989, lower 3,295	1.10 : 1

Table 14: From tiling through to the class, worst area of the sample.

Stage 4 damps by a factor of 311. The reason is the same one by which stage 1.4 amplifies a disturbance from the hydrology roughly 710-fold, only in the other direction: the classification is a

step function over the runout angle. A change in α flips a class only if it pushes γ across AAT_2 or AAT_3; in the middle of a band even a whole degree changes nothing.

No strong net direction in this area. The most frequent changes occur in pairs and almost equally in both directions: 2→1 in 2,219 cells against 1→2 in 1,808, 2→3 in 1,094 against 3→2 in 1,021. No strong net directional shift was observed in this aggregate; local systematic effects remain possible, since up- and down-classifications in different places can cancel in the balance.

This is an empirical result for one area, not a bound. The class effect was measured in the area with the largest observed α differences; an area with smaller α differences can still show more class changes if more cells lie near a class boundary, and the class of a cell depends on all runs that reach it, not only on the α at its own position. The reported class differences apply to the tested area and do not establish an upper bound elsewhere.

For comparison, the scalar path — the fixed α of 26° instead of the size-dependent mobility, the obvious shortcut around the hydrology: it differs by more than 0.5° on 85 % of the cells and changes 0.83 % of the classes. In this area the tiling is thus two orders of magnitude closer to the untiled chain than the fallback it would replace.

7 Borders and the outer edge

7.1 Inner borders

A map composed of per-region runs has two effects at every border that no compositing removes: a region sees neither terrain nor forest beyond its border, and an avalanche releasing on one side and running out on the other is modelled on neither side completely. With one DEM, one forest mosaic and one glacier mosaic for the whole region (Section 4.1) and tiles with a margin (Section 5.3) these inner borders do not arise. A national border is not a line at which the computation changes; what remains is the change of **data source** at the border — a different LiDAR campaign, a different forest product (LiDAR against Copernicus), a different glacier inventory. These changes are part of the input and are named per source in Section 4.1.

How much the change of method moves the result was measured against the per-region runs of Switzerland (10 October 2026) and Austria (11 October 2026), cell by cell on the Alpine grid, separately for the interior (more than 10 km from the border) and the border band (Section 8.2).

7.2 The outer edge

The Alpine DEM has an outer edge, and beyond it there is no terrain in the model. No coarser elevation data are added beyond it: inside, the data are checked; outside there would only be coarser ones. Instead, it was checked where the edge matters at all: along the edge of the **valid** DEM cells (which also catches holes inside), in a strip of 3500 m inwards, per 5-km section — steep terrain over 30°, relief, and whether the terrain **rises** towards the edge, in which case avalanches from outside are possible.

Holes inside 46 holes of at least 1 km², almost all water or plain (Lake Geneva, Lake Garda, the Venetian lagoon and Po delta, the Venetian plain) or flat exclaves. Relevant is only the Lessinia in Veneto (31.7 km², steep valleys), which the regional DTM does not cover. It carries nodata, i.e. *not assessed*, and remains distinguishable from value 0 (Section 4.6.3).

Outer edge 527 sections with steep terrain, 112 of them rising towards the edge, concentrated in the Ligurian Alps (where the edge mostly runs on the ridge), at the edge of the Apennines, in

the Jura, at Mont Ventoux and along the southern edge of Slovenia. A rise towards the edge on a ridge means that avalanches from the other side run out on the other side.

The outer edge therefore lies almost everywhere in plains, at lakes or on ridges. Within 3500 m of it, the classification can miss runout from terrain beyond the edge; this is a limit of the statement, not a computational error, and individual places (Mont Ventoux, Ligurian Alps) can be added later as small additional regions from official sources.

8 Plausibility check

This chapter checks whether the results are consistent with known observations: the class shares per region, the agreement with the per-region runs and the location of avalanche accidents. It is not a validation of the classes (Section 8.3.1).

8.1 Model output

The Alpine product of 21 October 2026 covers 291,640 km² of assessed terrain (the valid area of the Alpine DEM: 291,635 km²). It was computed in 45 tiles of between 5 minutes and 6.3 hours each, 51 hours in total; 42 tiles on one machine with 12 cores and 24 GB, three on a second one. The duration follows the share of steep terrain in a tile, because the runout dominates.

Region	Assessed km²	Class ≥ 1 km²	Value 0 %	1 %	2 %	3 %	4 %
Austria	83,944	24,077	71.3	35.4	27.5	29.7	7.5
France (7 départements)	42,326	15,509	63.4	30.4	31.0	30.7	7.9
Switzerland, Liechtenstein	41,448	16,767	59.5	23.6	32.3	34.7	9.4
Piedmont	25,400	7,524	70.4	38.7	23.8	29.9	7.5
Lombardy	23,845	5,997	74.8	42.3	24.4	26.4	6.9
Veneto	16,669	2,951	82.3	51.8	18.4	19.0	10.8
Bavaria (Alpine arc)	12,439	1,885	84.8	43.6	19.7	28.2	8.5
Slovenia (Alpine arc)	8,126	2,701	66.8	68.2	13.6	13.4	4.8
Friuli Venezia Giulia	7,877	2,484	68.5	51.3	18.3	20.5	9.9
South Tyrol	7,398	4,495	39.2	26.4	31.7	34.2	7.7
Trentino	6,206	3,518	43.3	33.7	24.3	34.7	7.2
Aosta Valley	3,261	2,478	24.0	10.5	31.7	44.1	13.7
Fringe beyond the regions	12,700	970	92.4	93.9	4.7	1.1	0.2
Alpine region	**291,640**	**91,356**	**68.7**	**34.6**	**27.3**	**30.0**	**8.1**

Table 15: Class shares of the Alpine product. Classes 1–4 relative to the terrain with a class ≥ 1, value 0 relative to the assessed area. A cell is assigned to a region by the source map of the forest mosaic (LiDAR regions) or by the regional border (Copernicus regions); the fringe is the part of the Alpine DEM beyond these regions.

The order follows the relief: the Aosta Valley, South Tyrol and Switzerland have the largest shares of classes 3 and 4 and the smallest of value 0; Bavaria and Veneto, whose areas extend far into the foothills and plain, the smallest.

8.2 Comparison with the per-region runs

Switzerland and Austria were also computed per region on their national DEMs, in the national grid, with their own forest only (10 and 11 October 2026, with the same parameters). Between those runs and the Alpine product lie the change of method (Alpine grid, forest and terrain beyond the border, tiles) and one change of code, the path length of the forest effect (Section 4.4). The per-region runs write value 0 as nodata; for the comparison, their nodata within the country counts as value 0.

Area	km²	same class	Alpine higher	Alpine lower
Switzerland, interior	28,680	96.87 %	1.69 %	1.44 %
Switzerland, border band 10 km	12,608	96.87 %	1.69 %	1.45 %
Austria, interior	64,810	98.96 %	0.54 %	0.51 %
Austria, border band 10 km	19,134	98.79 %	0.64 %	0.57 %

Table 16: Alpine product against the per-region runs of 10/11 October 2026, cell by cell on the Alpine grid (per-region raster resampled by nearest neighbour). Interior: more than 10 km from the national border.

The differences are symmetric: about as many cells move up as down, and the most frequent changes are pairs between neighbouring classes (in Switzerland 2→3 on 0.45 % and 3→2 on 0.37 % of the cells, 1→0 and 0→1 on 0.36 and 0.33 %). No strong net directional shift is observed in this aggregate; this is consistent with resampling of the DEM and nearest- neighbour comparison across two grids as causes, but does not prove them, and local systematic effects remain possible. **The border band behaves like the interior.** Switzerland, whose DEM is resampled from EPSG:2056 for the Alpine grid, differs more than Austria, whose DEM is delivered in EPSG:3035 — consistent with resampling as the main source, though the shares have not been separated.

8.3 Comparison with accident locations

The product is compared against avalanche accidents with people involved. The comparison measures one thing: the share of selected accident coordinates that lie, under a fixed query, in terrain the model assigns class 2 or higher. This is a useful inventory check, but neither the sensitivity of the map nor a validation of the individual classes; an accident in correctly rated simple terrain does not by itself show that the rating is wrong.

Region	Source	Licence	Accidents	Period
CH	SLF EnviDat[36]	CC BY 4.0	2,840	1970–2025
CH	LAWIS Public API	attribution	6	2019–2026
AT	LAWIS Public API	attribution	2,480	2000–2026

Table 17: Data sources. Switzerland: ski touring only; Austria: ski touring and unknown activity, since LAWIS has no reliable activity before
2018. Both with at least one person caught. Swiss LAWIS accidents that the

SLF already records (same day, within 1 km) are not counted twice. SLF file of 15 December 2025; LAWIS retrieved 18 May 2026, its public interface has not delivered data since.

The class is read at the reported coordinate and, separately, as the highest class of the cells whose centre lies within 25, 50 and 100 m of it. As a baseline, the same query is made on a pure slope

map from the same elevation model: slope of at least 29° (the effective lower bound of slope class 2, Table 8).

	Switzerland (2,846)				Austria (2,480)			
Query	**Point**	**25 m**	**50 m**	**100 m**	**Point**	**25 m**	**50 m**	**100 m**
Version 2	**89.3**	96.4	98.1	99.0	**82.0**	92.5	95.8	96.5
Version 1 (Section 2)	87.8	96.4	98.2	99.1	81.7	93.1	95.6	96.5
Slope ≥ 29° only	82.1	96.1	98.3	99.1	72.0	91.7	97.9	99.6

Table 18: Share of accident coordinates (%) with class 2 or higher at the point or within the given distance. Classes at the point, version 2: Switzerland value 0 / 1 / 2 / 3 / 4 = 9.1 / 1.7 / 25.9 / 51.7 / 11.7 %, Austria 14.8 / 3.2 / 32.0 / 36.7 / 13.3 %. No accident lies in a cell not assessed.

Three observations.

At the reported coordinate, the model includes more accident locations than the slope baseline; neighbourhood queries provide little additional evidence of model skill. At the reported coordinate, version 2 reaches class 2 or higher for 89.3 and 82.0 % of the accidents, a pure slope map for 82.1 and 72.0 %. Within 50 m, the slope map reaches as much as the model or more (98.3 and 97.9 %). The neighbourhood queries therefore say little about the model: in steep terrain, almost every coordinate has steep terrain nearby. The neighbourhood query shows that most reported accident coordinates are close to terrain assigned class 2 or higher. It does not establish whether discrepancies at the point result from positional uncertainty, from the meaning of the coordinate (release zone, burial site or deposit) or from model omissions.

Where the coordinate is uncertain, the share falls. The SLF coordinate is a point in the release zone, with its accuracy given in metres (0, 100, 250, 500, … m; empty if unknown). For the 1,761 Swiss accidents with an accuracy of 0 m, 91 % lie in class 2 or higher at the point; for the 181 without an accuracy entry, 65 % — and there the slope map reaches only 49 %, so many of these coordinates do not lie on steep terrain at all. In Austria, accidents recorded as ski touring reach 90 %, those with unknown activity, which include off-piste skiing near ski areas, 77 %. The share also rises over time (Austria: 2000s 69.5 %, 2010s 81.0 %, 2020s 88.5 %), consistent with more precise reporting, though this has not been separated from other changes.

Version 2 against version 1, on the same accidents. In Switzerland, 94 accidents lie in class 2 or higher only in version 2 and 51 only in version 1: a difference of +1.5 percentage points (95 % interval +0.7 to +2.3; exact two-sided McNemar test on the discordant pairs, $p < 0.001$). In Austria, 70 accidents are captured only by version 2 and 63 only by version 1: +0.3 points (−0.6 to +1.2; $p = 0.60$), i.e. not distinguishable. The intervals come from 2,000 bootstrap resamples of the accidents, each accident carrying its pair of results (percentile interval). Accidents that are close in space or time are treated as independent; no clustering is accounted for. The intervals describe the uncertainty of this inventory comparison, not the accuracy of the product across the Alps.

Of the accidents below class 2 at the point, 258 (Switzerland) and 367 (Austria) lie in value 0 and 47 and 80 in class 1. That 9.1 and 14.8 % of the accident coordinates carry value 0 is a strong argument against reading this value as safe terrain (Section 4.6.3), even if part of it is due to coordinate uncertainty.

A higher share is not a better model. The metric does not penalise overclassification: expanding the area assigned class 2 or higher can increase the reported share without improving classification accuracy, as the slope map in the neighbourhood queries shows.

8.3.1 Limits of the metric

The metric only checks whether the accident location lies in terrain assigned class 2 or higher — it does **not** check whether the assigned class correctly reflects the exposure. An accident in a steep gully classified as Challenging instead of Complex counts as detected, although the class is too low.

An ordinally correct validation — a confusion matrix over the classes — requires a reference: expert ATES assessments with a documented definition, as Sykes et al.[32] and Markov et al.[8] use them. The SLF maps CAT and ATH[6] are not such a reference: CAT distinguishes other terrain aspects than a four-class ATES map, ATH combines release potential and consequences, and only the classified ATHc is oriented towards ATES. A comparison with them is a comparison of methods, not a measurement of accuracy, and shared inputs and related model assumptions can make the errors of both maps dependent. See Outlook.

In addition, accident locations are not a cross-section of the terrain. They lie where people travel, and that is not evenly distributed. A high detection rate therefore shows that terrain where accidents occurred is largely assigned class 2 or higher — not that the model classifies terrain correctly.

9 Discussion

9.1 Strengths

- **Transparent methodology:** all parameters, formulas and thresholds used are documented in this document, each with its origin — a publication, the reference configuration, or a decision of our own, which is named as such.
- **The entire Alpine region on one grid:** one DEM, one forest mosaic and one glacier mosaic for 291,640 km², computed using an implementation derived from the published AutoATES/ATI methods and verified against specified reference implementations, with explicit release areas and runout simulation, without seams at national borders (Section 5.3).
- **Versioned and traceable:** every number the pipeline sets comes from its configuration into this document; the computation is compared cell by cell with the reference implementations; the product is assembled byte-identically from stored tiles and identified by checksums on every copy. Independent reproduction would additionally require access to the software, configuration, inputs and test data, which are not published.
- **Grid spacing:** 5 m, finer than AutoATES v2 typically works (10–30 m) and the same nominal grid spacing as SLF CATv2. This is a technical property, not a demonstrated gain in classification quality.
- **Eastern Alpine parametrisation:** the thresholds of the BFW model chain (Huber et al.[9]), not the Canadian defaults — adopted, not recalibrated to our resolution (Section 9.2.1).

9.2 Limitations

9.2.1 The parametrisation is calibrated at 10 m, we compute at 5 m

This is the most consequential limitation. The Eastern Alpine parametrisation was determined on a 10-m elevation model; we have converted three quantities defined in pixels (Section 4.7.1), but not the **thresholds** — and these lie on quantities that themselves depend on resolution. Finer resolution can preserve steeper local terrain features; the net effect on threshold exceedance and final classes has not been quantified for this product.

Sykes et al.[32] find that the slope parameters have the largest effect of all and determine them per area and elevation model by a grid search against expert maps. Without such maps this cannot be reproduced. The parameters are therefore adopted without resolution-specific recalibration.

9.2.2 The elevation models are not of equal quality

The model is only as good as the DEM in each place. Most sources are airborne laser scans at 1–2.5 m; the Lombardy DTM (5 m, 2015) is largely interpolated from contour lines, with laser scanning only along the rivers. There, slopes are smoothed, small terrain features are missing, and in the Bormio area stripes from the source are visible. Lombardy is computed and delivered with this DTM; the effect on the classes has not been measured. A replacement by a laser-scanned DEM (the national survey of Italy, where available) is planned and would require recomputing only the affected tiles (Section 5.3). The Aosta Valley DEM is coarser than swissALTI3D, which is visible as a change of detail along the Monte Rosa ridge.

9.2.3 Tiles are not identical to one run in their margin

The ring check (Section 5.3) shows that neighbouring tiles disagree on 4.8 cells per million in their common margin, because the subcatchments are formed per tile. The product contains no seam from this, but it is not proven to be identical to a single run over the whole Alps; the closest measurement of this effect is the one for the tiled hydrology in one area (Section 6.5), which is not a bound for other areas.

9.2.4 The outer edge

Beyond the outer edge of the Alpine DEM there is no terrain in the model. Within 3500 m of it, runout from outside is missing; where the edge does not run on a ridge or in a plain (Ligurian Alps, Mont Ventoux, parts of the Jura and southern Slovenia), classes there can be too low (Section 7).

9.2.5 Wind shelter as a veto

The Cauchy membership of the wind shelter enters through the fuzzy AND, in which the minimum dominates. A low wind-shelter value can therefore suppress release detection — even on steep terrain. Affected are convex ridge backs and windward sides, where the slab probability is indeed lower, but not zero.

This is inherited from the reference design rather than introduced here, and the parameters (2 /5 /2, S_x in radians, Section 4.2.1) are calibrated for exactly this computation. How much release area the veto suppresses on steep terrain has not been quantified.

9.2.6 Ruggedness: a decision against Toft

Toft et al.[7, § 4.1.4] removed ruggedness from AutoATES v2 — at 5–30 m grid spacing the Vector Ruggedness Measure measures catchment rather than slope roughness. The Eastern Alpine calibration, by contrast, reintroduces it as a **binary mask**, and Huber et al.[9] report that this variant "performed well in comparison with the documented PRAs".

We follow the Eastern Alpine variant and apply the mask — it removes about a tenth of the release area (Section 4.2.3).

9.2.7 Terrain traps and remote triggering

Terrain traps are not explicitly classified in this implementation. Their representation from elevation data requires additional definitions and evaluation beyond the adopted workflow; ATES v.2

(Statham & Campbell [3]) treats them as a qualitative feature, and Toft et al. [7, § 4.1.2] report an improvement of 0.08–0.14 % for a runout-angle substitute.

Remote triggering (Harvey et al. [5]) is not represented explicitly: the model identifies release areas and their runout, not the possibility of triggering a release area from a point outside it. Whether and to what extent this contributes to discrepancies at reported accident coordinates has not been quantified. For the Swiss accidents this question is constrained by what the data record: the SLF coordinate is a point in the release zone (Section 8).

9.2.8 No separate cumulative overhead-exposure criterion

Runout modelling represents whether terrain can be reached from release areas above. However, the classification does not include a separate criterion for cumulative overhead exposure or the convergence of multiple avalanche paths, because the Eastern Alpine calibration does not include one. Consequently, terrain affected by several paths may not receive a higher class for that reason alone (Section 4.6).

9.2.9 Ties in the majority filter

Smoothing the area boundaries is arbitrary in a tie, and the reference does not decide it reproducibly (Section 6.1). Our version is deterministic; the difference in the reference area is 9 of 28,090,000 cells and affects only area identifiers, not classes.

9.3 Outlook

Validation. The product is verified against the reference implementations and checked for plausibility, but not validated. A validation that supports a statement such as “validated ATES classification for the Alps” would need several independent reference areas with different terrain and data conditions, expert ATES assessments with a documented definition, spatially separated calibration and test areas, results per class including under-classification, and a controlled comparison of 5-m and 10-m processing. It would also need targeted checks of old forest data, small release areas, the outer edge and problematic DEM regions, and of how terrain uncertainty and missing data propagate into Tour-Check.

Calibration at our resolution. The thresholds come from a 10-m elevation model; a grid search against expert maps, as carried out by Sykes et al. [32] and Markov et al. [8], would set them for 5 m. Without such maps the parametrisation is adopted, not calibrated.

Comparison against CAT/ATH. For Switzerland, the SLF maps (Harvey et al. [6], CC BY 4.0) allow a cell-by-cell comparison of methods, with a documented mapping of their categories to ATES classes. It does not replace an expert reference (Section 8.3.1).

Comparison with Skitourenguru. The terrain indicator of Skitourenguru on the route points of the BSRPD dataset [37] is an independent, Alps-wide terrain judgement; the comparison is to be repeated on the Alpine product.

Accident data. The Austrian accidents were retrieved on 18 May 2026; the evaluation is to be repeated with the current LAWIS data, and extended to Italy, France and Slovenia where accident data with coordinates are available.

Better elevation data for Lombardy (Section 9.2.2).

Remote triggering. Adding the Weibull kernel (Harvey et al. [5]) — requires an Alpine reference dataset.

Machine learning as an alternative to the rule set. Markov et al.[8] replace the rule-based classifier with random forests and provide an alternative classification approach; transfer to new regions, however, still requires suitable training data and independent evaluation — their models are trained and tested on a labelled dataset from the Pirin Mountains. In addition, the class assignment no longer follows from named thresholds but from a model. For a product whose ratings must be traceable, this is a trade-off rather than an unequivocal improvement.

Validation across regions. Further Alpine regions with avalanche cadastre data.

10 Coverage

Region	DEM source	Forest source	Glaciers
Switzerland, Liechtenstein	swissALTI3D (2 m)	WSL VHM NFI 2025 (1 m)	swissTLM3D
Austria	BEV ALS-DTM (1 m)	BEV ALS nDSM (1 m) + BFW forest map	GI 5
French Alps (7 dép.)	IGN LiDAR HD MNT (0.5 m)	IGN MNH LiDAR HD (1 m) + OCS GE 2.0	OCS GE 2.0
Bavaria (Alpine arc)	BVV DGM1 (1 m)	BVV DOM20 + ATKIS Wald/Gehölz	OSM
Slovenia (Alpine arc)	ARSO DMR (1 m)	CLSS nDMP (50 cm) + MKGP RABA 2026	OSM
South Tyrol	Prov. DTM (2.5 m)	Prov. DOM − DTM (2.5 m) + forest type map	OSM
Trentino	PAT DTM 2006–2008	PAT CHM 2006–2007 (1 m) + Bosco 2017	OSM
Aosta Valley	VdA DTM (2 m)	VdA DSM − DTM (2 m) + Art. 33 forest	OSM
Lombardy	Regione Lombardia DTM (5 m)	Copernicus TCD 2024 (10 m)	OSM
Piedmont	ICE 2009–2011 DTM (5 m)	Copernicus TCD 2024 (10 m)	OSM
Veneto	Regione Veneto DTM LiDAR (5 m)	Copernicus TCD 2024 (10 m)	OSM
Friuli Venezia Giulia	Protezione Civile DTM (1 m)	Copernicus TCD 2024 (10 m)	OSM

Table 19: Sources per region in the Alpine product of 21 October 2026. All regions are computed. On the Copernicus forest see Section 4.1.1; data sources and licences are documented in Section 4.1; on the glacier source see Section 4.1.2.

All regions of the table are part of one product; the Alpine DEM also contains a fringe beyond them, computed with Copernicus forest (Table 15). Not covered are the Lessinia in Veneto (a gap in the regional DTM) and the German Alpine foreland west of Bavaria.

The glacier column shows OSM where no official inventory is used. For Italy this is a known weakness, not the result of a check: Italy maintains its own inventories (GLIMS/RGI contributions), which we have not evaluated. As long as OSM applies, the same effect as in Austria is to be expected — outlines that are too large because they are too old (Section 4.1.2).

11 Licences

Software and data licences are separate questions and are listed separately. The product is a derived dataset. Most data licences below require the attribution of the source and of the licence on every use of the product. Two sources are under the Open Database Licence (ODbL): the Vegetation Height Model of WSL and OpenStreetMap. The ODbL distinguishes between a database, a derivative database and a produced work, and a publicly used derivative database can carry share-alike and access obligations (ODbL §§ 4.3–4.6) that go beyond attribution. How the processing chain and the products derived from these sources are to be classified under the ODbL has not yet been settled; this rights clearance is an open release item for the data product and separate from the scientific status of this report.

11.1 Software

The implementation is the provider's own code and is not published or licensed to third parties. External software is **called**, not copied: WhiteboxTools (MIT) in production; com4FlowPy/AvaFrame and the ATI model chain (both EUPL-1.2) only for verification. The methods of AutoATES v2 (GPL-3.0) are reimplemented from the publications; no code is copied.

11.2 Data

Use	Source (holder)	Licence
Elevation	swissALTI3D (swisstopo)	swisstopo terms of use (OGD)
	ALS-DTM (BEV)	CC BY 4.0
	DGM1 (Bavarian surveying administration)	CC BY 4.0
	LiDAR HD MNT (IGN)	Licence Ouverte 2.0
	DMR1 (GURS / ARSO)	CC BY 4.0
	DTMs of Aosta Valley, Piedmont, Lombardy, Trentino	CC BY 4.0
	DTM 2.5 m (Autonomous Province of Bolzano)	CC0 1.0
	DTMs of Veneto and Friuli Venezia Giulia	IODL 2.0
Forest	Vegetation Height Model NFI (WSL)	ODbL with DbCL
	ALS DOM (BEV); forest map (BFW)	CC BY 4.0
	DOM20; ATKIS Basis-DLM (Bavaria)	CC BY 4.0
	MNH LiDAR HD; OCS GE 2.0 (IGN)	Licence Ouverte 2.0
	DOM, DTM, forest type map (Bolzano)	CC0 1.0
	CHM 2006–2007 (PAT, FEM; Zenodo); forest map Bosco	CC BY
	DSM/DTM 0508 (Aosta Valley)	CC BY 4.0
	Forest mask Art. 33 (Aosta Valley)	CC BY 4.0
	CLSS nDMP and point cloud (GURS)	CC BY 4.0
	RABA land use (MKGP)	no named licence; no restrictions on use
	Tree Cover Density 2024 (Copernicus Land Monitoring Service)	Copernicus data policy
Glaciers	swissTLM3D (swisstopo)	swisstopo terms of use (OGD)
	Austrian Glacier Inventory GI 5	CC BY 4.0
	OCS GE 2.0 (IGN)	Licence Ouverte 2.0
	OpenStreetMap contributors	ODbL 1.0
Evaluation	Avalanche accidents in Switzerland (SLF)	CC BY 4.0
	LAWIS (Austrian avalanche warning services)	attribution

Table 20: Data licences per source. Aosta Valley: the licence the region states for its geodata. GURS: its general terms for public geodata. RABA: no licence is named by the ministry or in the INSPIRE record; the national open-data catalogue lists the dataset with “no restrictions on use” and as available for access and reuse. Its default licence (CC BY 4.0) does not apply, because the data are only linked there; the source is attributed nevertheless. None of this replaces a legal review of the published data product.

References

[1] G. Statham, B. McMahon, and I. Tomm, ‘The Avalanche Terrain Exposure Scale’, in *International Snow Science Workshop, Telluride*, 2006.

[2] G. Statham *et al.*, ‘A conceptual model of avalanche hazard’, *Natural Hazards*, vol. 90, no. 2, pp. 663–691, 2018, doi: 10.1007/s11069-017-3070-5.

[3] G. Statham and C. Campbell, 'The Avalanche Terrain Exposure Scale (ATES) v.2', *Natural Hazards and Earth System Sciences*, vol. 25, pp. 1113–1137, 2025, doi: 10.5194/nhess-25-1113-2025.

[4] G. Schmudlach and J. Köhler, 'Method for an automatized avalanche terrain classification', in *Proceedings, International Snow Science Workshop, Breckenridge, Colorado*, 2016.

[5] S. Harvey, G. Schmudlach, Y. Bühler, L. Dürr, A. Stoffel, and M. Christen, 'Avalanche terrain maps for backcountry skiing in Switzerland', in *Proceedings, International Snow Science Workshop, Innsbruck, Austria*, 2018, pp. 1625–1631.

[6] S. Harvey, M. Christen, Y. Bühler, C. Hänni, N. Boos, and B. Bernegger, 'Refined Swiss avalanche terrain mapping CATv2 / ATHv2', in *Proceedings, International Snow Science Workshop, Tromsø, Norway*, 2024. [Online]. Available: https://arc.lib.montana.edu/snow-science/item/3363

[7] H. B. Toft, J. Sykes, A. Schauer, J. Hendrikx, and A. Hetland, 'AutoATES v2.0: Automated Avalanche Terrain Exposure Scale mapping', *Natural Hazards and Earth System Sciences*, vol. 24, pp. 1779–1793, 2024, doi: 10.5194/nhess-24-1779-2024.

[8] K. Markov *et al.*, 'Machine learning for automated avalanche terrain exposure scale (ATES) classification', *Natural Hazards and Earth System Sciences*, vol. 25, pp. 4375–4403, 2025, doi: 10.5194/nhess-25-4375-2025.

[9] A. Huber *et al.*, 'AutoATES Austria — testing and application of an automated model-chain for avalanche terrain classification in the Austrian Alps', in *Proceedings, International Snow Science Workshop, Bend, Oregon*, 2023, pp. 929–936. [Online]. Available: https://arc.lib.montana.edu/snow-science/item/2989

[10] Y. Bühler, D. von Rickenbach, A. Stoffel, S. Margreth, L. Stoffel, and M. Christen, 'Automated snow avalanche release area delineation – validation of existing algorithms and proposition of a new object-based approach for large-scale hazard indication mapping', *Natural Hazards and Earth System Sciences*, vol. 18, no. 12, pp. 3235–3251, 2018, doi: 10.5194/nhess-18-3235-2018.

[11] J. Veitinger, R. S. Purves, and B. Sovilla, 'Potential slab avalanche release area identification from estimated winter terrain: a multi-scale, fuzzy logic approach', *Natural Hazards and Earth System Sciences*, vol. 16, pp. 2211–2225, 2016, doi: 10.5194/nhess-16-2211-2016.

[12] C. J. L. D'Amboise *et al.*, 'Flow-Py v1.0: a customizable, open-source simulation tool to estimate runout and intensity of gravitational mass flows', *Geoscientific Model Development*, vol. 15, pp. 2423–2439, 2022, doi: 10.5194/gmd-15-2423-2022.

[13] A. Huber *et al.*, 'Regional-scale avalanche modeling with com4FlowPy — potential and limitations for considering avalanche-forest interaction along the avalanche track', in *Proceedings, International Snow Science Workshop, Tromsø, Norway*, 2024. [Online]. Available: https://arc.lib.montana.edu/snow-science/objects/ISSW2024_P4.1.pdf

[14] J. Sykes, D. Knies, P. Haegeli, K. Anthony-Malone, and G. Statham, 'autoATES v3.0: Automated ATES mapping at scale across western Canada', in *Proceedings, International Snow Science Workshop, Whistler, British Columbia*, 2026. [Online]. Available: https://sfuarp.ca/publications/2026_issw_sykes_autoates/

[15] C. Hesselbach, P. Spannring, A. Huber, and J.-T. Fischer, 'OpenNHM/AvalancheTerrainIdentification: Avalanche Terrain Identification'. Zenodo, 2026. doi: 10.5281/zenodo.22049056.

[16] P. Spannring, C. Hesselbach, A. Huber, and J.-T. Fischer, 'Towards a model chain for an automated thalweg identification'. 2026. doi: 10.5194/egusphere-egu26-10580.

[17] J. Schumacher, H. Toft Larsen, J. P. McLean, M. Hauglin, R. Astrup, and J. Breidenbach, 'The utility of forest attribute maps for automated Avalanche Terrain Exposure Scale (ATES) modelling', *Scandinavian Journal of Forest Research*, vol. 37, no. 4, pp. 264–275, 2022, doi: 10.1080/02827581.2022.2096921.

[18] C. Ginzler, *Vegetation Height Model NFI.* (2021). Swiss Federal Institute for Forest, Snow, Landscape Research WSL. doi: 10.16904/1000001.1.

[19] M. Marty and C. Ginzler, *Vegetation Height Model LiDAR NFI.* (2025). Swiss Federal Institute for Forest, Snow, Landscape Research WSL. doi: 10.16904/envidat.664.

[20] C. Bauerhansl, T. Koukal, and K. Schadauer, 'Waldkarte Österreich — Flächendeckende Waldtypenklassifikation für Österreich', *BFW-Praxisinformation*, vol. 14, pp. 12–14, 2007.

[21] Copernicus Land Monitoring Service, 'CLMS VLCC Tree Cover Density (TCD) Europe 10 m yearly V1, reference year 2024'. [Online]. Available: https://land.copernicus.eu/en/products/high-resolution-layer-tree-cover-density

[22] N. Lang, W. Jetz, K. Schindler, and J. D. Wegner, 'A high-resolution canopy height model of the Earth', *Nature Ecology & Evolution*, vol. 7, pp. 1778–1789, 2023, doi: 10.1038/s41559-023-02206-6.

[23] Meta AI and World Resources Institute, 'High Resolution Canopy Height Maps v2'. 2025.

[24] A. Hartig *et al.*, 'The Austrian Glacier Inventory GI 5, 2021–2023, as vector files'. PANGAEA, 2026. doi: 10.1594/PANGAEA.991106.

[25] L. Hartl *et al.*, 'On thin glacial ice: New Austrian Glacier Inventory shows accelerating glacier shrinkage and 31% area loss within two decades', *The Cryosphere*, vol. 20, no. 9, pp. 5157–5180, 2026, doi: 10.5194/tc-20-5157-2026.

[26] J. Buckel and J.-C. Otto, 'The Austrian Glacier Inventory GI 4 (2015) in ArcGis (shapefile) format'. 2018. doi: 10.1594/PANGAEA.887415.

[27] B. Werners, 'Aggregation models in mathematical programming', *Mathematical Models for Decision Support*, vol. 48. in NATO ASI Series F, vol. 48. Springer, 1988.

[28] J. M. Sappington, K. M. Longshore, and D. B. Thompson, 'Quantifying landscape ruggedness for animal habitat analysis: a case study using bighorn sheep in the Mojave Desert', *Journal of Wildlife Management*, vol. 71, no. 5, pp. 1419–1426, 2007, doi: 10.2193/2005-723.

[29] C. Hesselbach, 'Adaptation and application of an automated avalanche terrain classification in Austria', 2023. [Online]. Available: https://permalink.obvsg.at/bok/AC16964320

[30] A. Winstral, K. Elder, and R. E. Davis, 'Spatial snow modeling of wind-redistributed snow using terrain-based parameters', *Journal of Hydrometeorology*, vol. 3, pp. 524–538, 2002, doi: 10.1175/1525-7541(2002)003<0524:SSMOWR>2.0.CO;2.

[31] C. Plattner, L. N. Braun, and A. Brenning, 'Spatial variability of snow accumulation on Vernagtferner, Austrian Alps, in winter 2003/2004', *Zeitschrift für Gletscherkunde und Glazialgeologie*, vol. 39, pp. 43–57, 2006.

[32] J. Sykes, H. Toft, P. Haegeli, and G. Statham, 'Automated Avalanche Terrain Exposure Scale (ATES) mapping – local validation and optimization in western Canada', *Natural Hazards and Earth System Sciences*, vol. 24, no. 3, pp. 947–971, 2024, doi: 10.5194/nhess-24-947-2024.

[33] H. J. Körner, 'Reichweite und Geschwindigkeit von Bergstürzen und Fließschneelawinen', *Rock Mechanics*, vol. 8, no. 4, pp. 225–256, 1976, doi: 10.1007/BF01259363.

[34] J. F. O'Callaghan and D. M. Mark, 'The extraction of drainage networks from digital elevation data', *Computer Vision, Graphics, and Image Processing*, vol. 28, no. 3, pp. 323–344, 1984, doi: 10.1016/S0734-189X(84)80011-0.

[35] European Avalanche Warning Services, 'Avalanche Size'. 2026.

[36] WSL Institute for Snow and Avalanche Research SLF, 'Avalanche accidents in Switzerland since 1970/71'. 2023. doi: 10.16904/envidat.411.

[37] C. Furrer, 'Artificial intelligence in ski touring: prediction of caution- and foot-sections for ski touring routes with a machine learning approach', 2024.