AUTOATES V3.0: AUTOMATED ATES MAPPING AT SCALE ACROSS WESTERN CANADA

John Sykes[1*], Declan Knies[1], Pascal Haegeli[1,2], Kristin Anthony-Malone[1], and Grant Statham[3]

[1]*Simon Fraser University Avalanche Research Program, Burnaby, BC, Canada*
[2]*WSL Institute for Snow and Avalanche Research SLF, Davos, Switzerland*
[3]*Parks Canada, Banff, AB, Canada*

ABSTRACT: Avalanche Terrain Exposure Scale (ATES) maps help backcountry travelers understand how exposed a piece of terrain is to avalanche hazard. Making those maps by hand is slow and expensive, so they exist for only a few high-use areas. Producing consistent maps for every avalanche forecast region in western Canada, 400,000 km², is a different problem: it removes the option of local tuning. This paper describes the working autoATES v3.0 model for that production run, and the changes from autoATES v2.0 that the domain forced. A domain-wide Sentinel-2 forest product, built from paired summer and winter scenes at 10 m spatial resolution, supplies a consistent canopy layer from the Coast Range to the Rockies. The classifier uses the area of openings in that canopy (gap area, in hectares), not average canopy percentage, as the forest variable. Treeline is estimated per forecast subregion rather than as a single elevation. Elevation comes from a global 30 m surface model, chosen for consistency over a lidar patchwork. Potential release areas and runout are each produced as a typical scenario and a more inclusive infrequent scenario, so low-angle persistent-weak-layer terrain can be represented without compromising the typical layer. The ATES classifier follows the published ATES v.2 technical model: floors for the bold defaults, then an averaging step for the remaining evidence. Typical and infrequent PRA and runout are the regional calibration surface. Canadian production maps use both scenarios and classes 0–4; the same method can rate a typical scenario alone as classes 1–4, with Simple as the default. Observed avalanches and expert maps answer different questions; neither is ground truth. Connaught Creek, mapped independently of autoATES, is the transfer test: the classifier agrees with the expert consensus at the same level that three mappers agree with each other. Maps edited from an autoATES draft are how the reference set can grow. Finished maps will be hosted by Avalanche Canada and made freely available. A best-practices document is in development. The aim is a defensible, consistent product for western Canada, not a closer fit to any one expert map.

KEYWORDS: Automated avalanche terrain mapping; Avalanche Terrain Exposure Scale (ATES); runout simulation; model validation; trip planning

1. INTRODUCTION

The Avalanche Terrain Exposure Scale (ATES) classifies mountainous terrain based on the degree of exposure to avalanche hazard (Statham et al. 2006; Statham and Campbell 2025). Traditional ATES mapping relies on manual delineation of polygons or trail segments and is both time- and cost-intensive, so high-quality maps exist for only a few high-use areas (Sykes et al. 2024). Automated ATES (autoATES) was developed to make terrain ratings feasible at much larger scales (Larsen et al. 2020; Toft et al. 2024).

The Search and Rescue New Initiatives Fund (SAR-NIF) project “Large-scale avalanche terrain exposure scale mapping for mountainous regions in Canada” aims to produce consistent ATES maps for every avalanche forecast region in western Canada, a domain of roughly 400,000 km². A central problem at that scale is whether the model should change between mountain ranges and snow climates. The default is to keep the same settings everywhere unless local validation showed a difference that is physically plausible, and not an artifact of how the data were collected or of one expert’s mapping style. That test is still open. The configuration described here is the current working model, which uses consistent parameters for all areas.

Experience developing autoATES v2.0 and applying it in regions of the United States and Canada set four targets for this project. (1) Producing a consistent forest-canopy layer, designed around accurately capturing avalanche paths and treeline forest cover. (2) Paired typical and infrequent scenarios for potential release area and runout, calibrated to observed avalanches, so one parameterization does not have to stand in for both everyday events and the larger, less frequent avalanches that still belong on an ATES map. (3) An ATES classifier built to follow the workflow in the ATES v.2 technical model: defaults for the priority criteria, then an averaging step to combine additional terrain characteristics. Regional calibration is intended to be captured in the start-zone and runout

* *Corresponding author address:*
John Sykes (presenting author), Simon Fraser University Avalanche Research Program,
Burnaby, BC, Canada; email: john_sykes@sfu.ca

https://doi.org/10.15788/1790099379

scenarios, so the classifier does not have to be retuned for different regions. (4) Validation design that uses observed avalanches to calibrate release and runout, and expert maps — including maps edited from an autoATES draft — to score the classifier.

The objective of this paper is to describe the pragmatic changes from autoATES v2.0 to the working v3.0 model that address those challenges. This paper is written for avalanche-terrain modelers working with autoATES or similar tools, not as a general introduction to ATES mapping. We describe each change at the level needed to understand the design choice; supporting statistical detail is left to the poster and to the companion papers. The model described here is the working version for the production run; PRA, runout, and classification parameters may still change before the forecast-region maps are published for general use. Section 2 covers input datasets and the PRA and runout sub-models. Section 3 describes the classification approach. Section 4 discusses validation. Section 5 outlines the remaining steps to production. Detail that does not fit here is left for the Rogers Pass multi-agency comparison (Sykes et al. 2026) and the runout calibration paper (Knies et al. 2026).

2. ADJUSTMENTS TO INPUT DATASETS AND SUB-MODELS

2.1. Forest cover from Sentinel-2

Production mapping across 400,000 km² requires a forest layer that is consistent from the Coast Range to the Rockies. Canopy cover and gap area describe different properties of the same forest (Bebi et al. 2021). Cover is how continuous the forest is in a neighbourhood. Gap area is the size, in hectares, of a non-forested opening enclosed by that forest. Potential release area uses canopy cover as a terrain property that lowers the likelihood of release. The ATES classifier uses gap area because the rating depends on the size of those openings, not only on the density of the surrounding canopy.

The canopy product is built from Sentinel-2 Level-2A imagery at 10 m spatial resolution, using paired summer (July–September 2024–2025) and winter (February–March 2023–2025) scenes. Sentinel-2 is free and available for mountain ranges worldwide. Summer imagery provided greenness and vegetation indices. Winter separates conifer from leafless deciduous cover and identifies avalanche paths that descend through narrow gullies in otherwise mature canopy that are difficult to detect due to shadows in summer imagery. The operational layer is binary, coniferous forest versus not. The option exists to include a deciduous forest class but this functionality has not been tested thoroughly. We trained the classifier on region-specific samples and on a single merged training set from all zones; the merged set performed better, so the operational product uses one model across the domain. Canopy cover is the fraction of conifer cells in a 3×3 neighbourhood (30 m).

Gap area is derived from the canopy-cover raster. Cells with cover below 30 % are classed as open. Adjacent open cells that touch on an edge or a corner are grouped into a single opening, and each open cell is assigned the area of that opening in hectares, capped at 100 ha. Forested cells are set to zero. Cells outside the mapped domain are treated as forest so that separate openings are not merged across a tile edge. The gap term is only evaluated where mean canopy cover within 300 m exceeds 30 %, so alpine terrain is not read as one unbounded opening.

The same canopy layer is used to estimate regional treeline elevation in each Avalanche Canada forecast subregion based on the 95th percentile of forested pixels. Where a subregion has too little forest to estimate that percentile, the value falls back first to the mountain range and then to a Coast / Columbia / Rockies group median. Range medians are 1,385 m on the Coast, 1,889 m in the Columbia Mountains, and 2,111 m in the Rockies. In the current pipeline the treeline is used only to segment large potential release areas. It is not an ATES input. It may still be useful for snowpack modeling or forecasting work that needs a consistent real-world treeline boundary.

The current mosaic covers the Avalanche Canada and Parks Canada forecast regions (85 Sentinel-2 tiles across the Coast, Columbia, Rockies, and Northwest; Figure 1). This is the first consistent forest dataset for the study area, specifically developed to accurately capture avalanche terrain. The classification model will be published open source so the same method can be applied in other regions.

2.2. Regional scale elevation data

Elevation comes from the ALOS AW3D30 global product, a 30 m surface model. We treat this as a low-resolution baseline for the whole study area rather than stitching in higher-resolution elevation where it exists. That higher-resolution data is not available widely enough to justify a patchwork. In

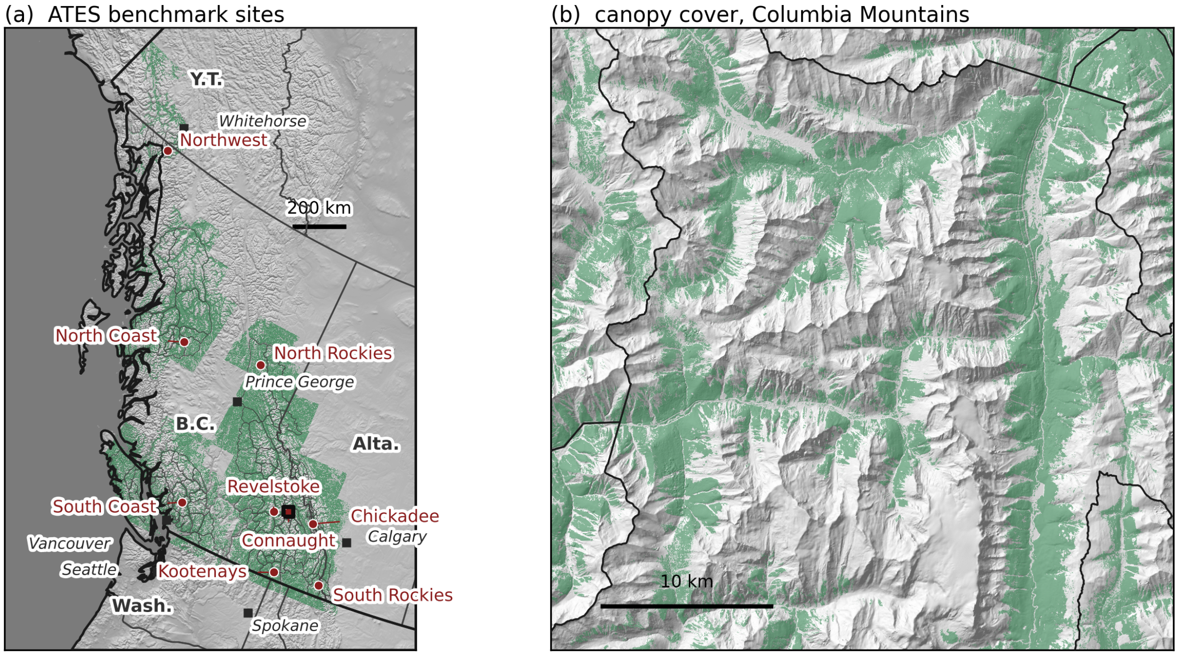

Figure 1: (a) Sentinel-2 conifer canopy across the Avalanche Canada forecast regions, on Natural Earth shaded relief. Red points are the nine ATES benchmark mapping sites. Provincial, state, and national borders and selected cities are shown for orientation. The box is the extent of (b). (b) Canopy cover in the Columbia Mountains. Imagery is at 10 m spatial resolution. Summer scenes are from July–September 2024–2025 and winter scenes from February–March 2023–2025.

forested areas the elevation model surface sits on the canopy top rather than on the ground, so slope, curvature, and wind shelter are biased in those areas. No competing global DEM avoids that, so it is a limit of working at this scale, not a reason to switch products. A comparison with other global surface models, including Copernicus GLO-30 and Canadian MRDEM, is planned before the maps go to production. How much the ratings change with DEM resolution is still an open question. We plan to measure that with 12 m TanDEM-X data, acquired for scientific use through DLR, at the ATES benchmark sites. One consequence of the 30 m surface is that slope is smoothed. Very steep ground that should rate Extreme, and isolated low-angle start zones that should stay Simple or non-avalanche, are both harder to pick out from the elevation data alone.

2.3. Potential release areas: typical and infrequent scenarios

The terrain characteristics of avalanche start zones can vary significantly depending on snowpack conditions and weak layer properties. Under typical snowpack conditions avalanches tend to fail largely between 30–50°, which is how most avalanche terrain models characterize slope distributions (Sykes et al. 2026). However, avalanches can fail on lower angle terrain, especially during periods with active persistent weak layers. Surface hoar is especially notorious for failing on low angle slopes, with slope angle distributions for documented events including terrain below 20° (Atkins et al. 2026). Automated ATES mapping requires capturing the full spectrum of potential release areas in order to accurately characterize terrain exposure.

The working autoATES v3.0 model therefore defines two PRA scenarios rather than averaging them into one layer. The typical scenario is tighter and captures terrain that could produce avalanches under common conditions. The infrequent scenario is more inclusive and captures low-angle slopes as well as the potential for avalanches within forested terrain. The updated autoATES v3.0 PRA step also segments output raster data into discrete polygons that enable characterization of start zone size as well as release likelihood.

PRA parameters were calibrated against 613 start zones mapped by operating ski-area programs at four sites (Lake Louise $n{=}66$, Whitewater $n{=}97$, Blackcomb $n{=}221$, Whistler $n{=}229$). The polygons originate from InfoEx control paths and were expanded to include all in-bounds start zones, including terrain managed by ski cutting and skier traffic.

A grid search evaluated a large set of PRA parameter combinations. Detection rate is the fraction of mapped start-zone cells that the model flags (cell-level recall). Because a model that flags the entire domain would achieve a detection rate of 1, performance is plotted against the ratio of flagged area to mapped start-zone area (Figure 2). The upper envelope of the resulting set of PRA model configurations is the set of parameter combinations that achieve the highest detection rate for any given flagged-area ratio.

On all mapped start zones (panel a), AutoATES v2.0 reaches a detection rate of 0.65 at 1.8 times the mapped start-zone area. The typical parameterization lies farther along the same envelope (0.73–0.74 at 2.3–2.4 times). The infrequent parameterization is displaced well to the right (0.94 at 6.1 times). Panel (b) shows why the second parameterization is required: below 30°, v2.0 detects almost none of the mapped start zones, the typical parameterization detects 0.29, and the infrequent parameterization detects 0.86.

Determining the optimal model calibration is only partly a statistical choice. Further analysis of model parameters against validation data are still being reviewed, but they do not decide how much residual miss is acceptable on the published maps. Capturing every mapped start zone, especially isolated low-angle releases that a 30 m surface tends to flatten, requires conservative parameters and a large extra-area cost. That is an operational decision about residual risk, constrained by the DEM, that impacts end user communication and needs to be weighed in that context.

2.4. Runout: typical and infrequent scenarios

Runout has the same one-layer problem as PRA. A parameterization tuned to typical events misses the furthest-running avalanches. A parameterization tuned to far-running events pushes exposure too far down every path. AutoATES v3.0 therefore uses a typical and an infrequent runout scenario as well.

Current working values come from BC Ministry of Transportation observed-avalanche records (Knies et al. 2026). Typical uses a 30° alpha angle and infrequent uses 18°, a conservative envelope rather than a regional fit, and will be tested again before production. Dynamic alpha angles that scale with PRA size are also under consideration and still need testing in western Canada (Hesselbach et al. 2026).

Against a small sample of large-magnitude avalanches mapped at Rogers Pass in winter 2026, the infrequent scenario covers 98.1 % of the mapped footprint and the typical scenario covers 69.7 % (Sykes et al. 2026). Those figures are a check on reach, not a complete runout calibration.

3. CHANGES IN ATES CLASSIFICATION

The approach to ATES classification has fundamentally changed in this version. AutoATES v2.0 represented slope angle and forest density and combined them with PRA and runout through thresholds that could be retuned based on regional differences (Sykes et al. 2024). The working v3.0 classifier follows the workflow the ATES v.2 technical model describes for a human mapper (Statham and Campbell 2025). Regional differences in avalanche process are meant to be characterized by the typical and infrequent PRA and runout scenarios, which are calibrated to observed release and runout (Section 2), rather than by adjusting classification thresholds for each mountain range. Typical events are those encountered under normal conditions; infrequent events are the larger, less frequent releases that occur during especially unstable snowpacks or very large storms. The same classifier is used from the Coast Range to the Rockies.

3.1. Following the ATES v.2 technical model

The ATES v.2 technical model specifies eight terrain parameters across five classes and two kinds of criteria. Bold defaults are critical attributes: if one is met, the rating is at least that class. Six of them exist: four in the frequency–magnitude row, exposure at Complex, and route options at Challenging. The remaining criteria are weighed by judgment. The classifier follows that structure.

Tier A implements the bold defaults that a cell-based map can reach as floors: if a default is met, the class cannot fall below that rating. Infrequent runout reach sets Simple; typical runout reach sets Challenging. Exposure at Complex is treated as typical overhead hazard: enough runout flow through the cell to indicate start zones above it. Route options at Challenging is not implemented. It is a property of the lines a party might travel, not of a cell. The frequency–magnitude default at Complex (1:1 for Size 3 and larger) is not used as a floor; it needs a release volume these layers do not carry. Extreme has no bold default in the tech-

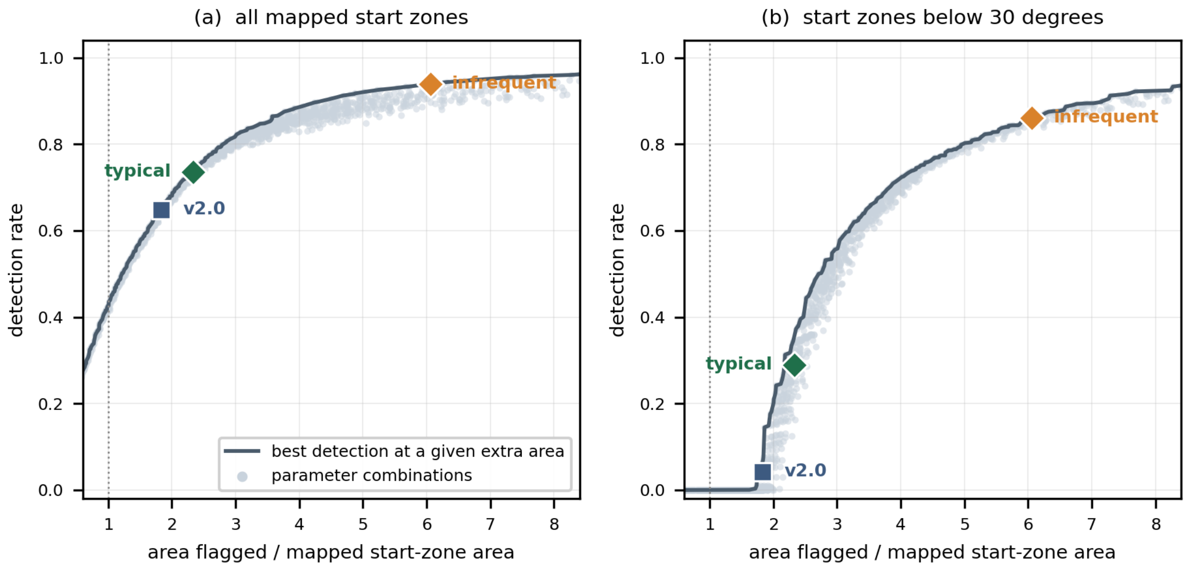

Figure 2: Potential release area detection rate against extra area flagged, at four ski-area validation sites. Grey points are parameter combinations. (a) All mapped start zones. (b) Mapped start zones below 30°, where the infrequent scenario earns its additional area. Published autoATES v2.0 sits on the envelope of what these data allow; the v3.0 typical and infrequent scenarios span that point by design.

nical model. The classifier still applies a floor of slope greater than 45° and open canopy, because that is the defining attribute of Extreme terrain.

Tier B assigns ordinal scores on PRA, runout, neighbourhood slope, and forest-opening size. On PRA and runout, Simple and Challenging read the infrequent envelope; Complex reads the typical layer, because a Complex rating is a statement about how often the terrain is exposed. Neighbourhood slope and forest-opening size do not depend on the scenarios. A score can stay silent when that layer has nothing to say. Silence is required: a modeled release of zero on wind-scoured rock means avalanches are not likely to release there, not that the terrain is Simple. Scores that spoke are averaged. The class is the higher of that average and the floor, so a floor cannot be averaged down.

Class 0 is a third, optional step, used only when an infrequent envelope is supplied. It is the only dedicated definition of non-avalanche terrain. After the floors, the rule tests how far a cell is from infrequent start zones and runout, the slope angle (the cell and a 300 m neighbourhood), and forest cover. Those readings are averaged, and class 0 is assigned only where that definition is clearly met. Floors can only raise a class, so they never create class 0. This step can lower a cell to 0, including over a floor. The bar is set high so that rating terrain non-avalanche is conservative. Cells the infrequent runout actually reaches almost never pass, because they are not far from runout. Using two scenarios and including a non-avalanche class are the Canadian production setup, not a requirement of the v3.0 ATES classifier. Typical PRA and runout are enough to assign Challenging, Complex, and Extreme. If no infrequent scenario is supplied, those scores stay silent and the study area defaults to Simple, producing a 1–4 map.

The method speaks to five of the eight technical-model parameters. A parameter is claimed only where a layer genuinely corresponds to it. Slope angle and forest density, start-zone size and density, runout, frequency–magnitude (through the two scenarios), and exposure are implemented. Slope shape was tested as curvature, and terrain traps as fall exposure and forest interaction with runout. Those inputs did not definitively improve classification accuracy, so they were not added, in order to keep the model simple. Further tests are planned to finalize the input rasters and the Tier B criteria. Route options remains the one unimplemented bold default. A zone-wide automated rating therefore answers a different question from a route-specific ATES rating, even when the class names match. Route-scale use is left to later trip-planning tools (Section 5).

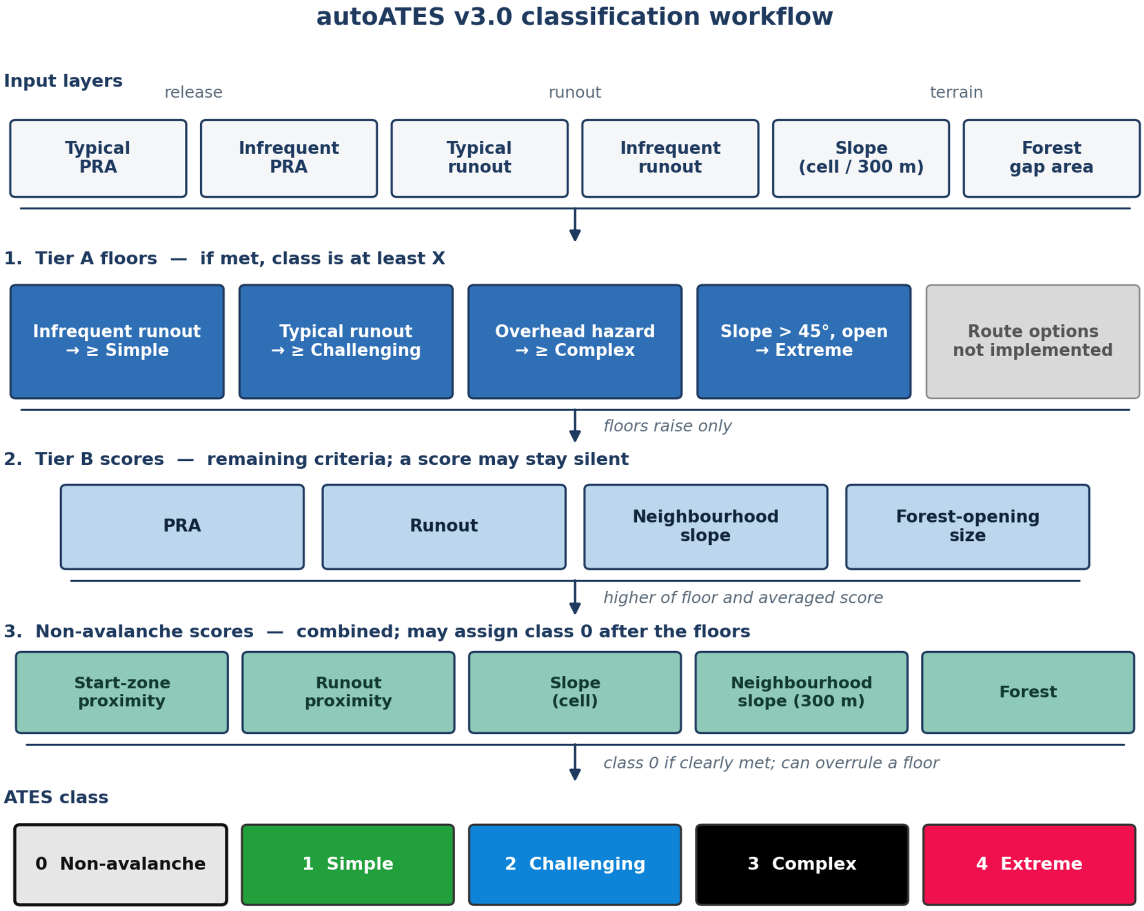

Figure 3: Classification workflow. Input layers feed three stages: Tier A floors (if met, the class is at least X), Tier B ordinal scores on PRA, runout, neighbourhood slope, and forest-opening size (a score may stay silent; scores that spoke are averaged), and an optional non-avalanche step whose five readings — start-zone proximity, runout proximity, local slope, neighbourhood slope, and forest — are combined after the floors. The class is the higher of the floor and the averaged score, unless that step assigns class 0. Class 0 is used only when an infrequent envelope is supplied. Route options is the unimplemented bold default.

4. MODEL VALIDATION

4.1. *Validation data for each step*

Each stage of the autoATES v3.0 workflow is checked against a different reference, because each stage answers a different question. Potential release areas and runout have to be parameterized against observed avalanches. The best records available are start zones mapped by operating ski-area programs, avalanche observations from the BC Ministry of Transportation, and a small set of events mapped at Rogers Pass. The ski-area and highway records represent decades of consistent observation of the same ground, so they are a substantial sample of the start zones and runout paths those operations actually manage. They are not a sample of all avalanche terrain. They miss events that were never recorded, and they are biased by the character of each operation: in-bounds ski terrain, and paths that reach a highway. They also show where avalanches did run, not where one could not.

Expert ATES maps cover terrain completely, including where nothing ran, but they encode a judgment about exposure that no observation can confirm. The two answer different questions. A validation design has to say which question it is asking.

4.2. *Expert maps as a reference*

PRA is scored against the ski-area start zones, always pairing detection rate with extra area (Section 2.3). Runout is calibrated to the Ministry of Transportation record and checked against mapped Rogers Pass events (Section 2.4; Knies et al. 2026; Sykes et al. 2026).

ATES maps are judgments. Each expert has tendencies and a technique, and a cell-based model is more precise than a person will usually draw by hand. Comparing an automated map with an independently drawn one, as at Connaught Creek and Bow Summit, therefore mixes two things: places the mapper thought the model was wrong, and places they simply did not digitize at the same resolution because the work is tedious. The benchmark maps in this project were commissioned the other way around. Mappers were given an autoATES product as a starting point, so that what they changed is, as far as possible, where they thought the model was wrong.

How much those maps still resemble the provided automated version is then partly mapper technique, not a ranking of which references are most useful. We measure that resemblance with weighted kappa: chance-corrected ordinal agreement, in which a one-class disagreement counts less than a three-class disagreement. Five of them were edited in place by a professional ATES mapper at Avalanche Canada, as they would to produce a publishable map. They remain close to the version they started from (weighted kappa of about 0.95 to 0.98; Figure 4b). We refer to this group of benchmark maps as ‘corrected’, due to the high degree of similarity to the automated version. That is not a write-off, and it is not independent judgment from a blank page either. It is operational validation: how much of the draft survived a publishable-map edit. Three other mappers also started from a model product and redrew more freely; those maps sit farther from the provided automated map and we refer to them as ‘redrawn’. This is likely due to differences in mapping practices and potentially due to local familiarity with the terrain. Figure 4b shows how much each benchmark site varies from the original automated map that was provided to the mapper. On the five corrected maps the seed agrees more closely with the expert layer than the working classifier does (Figure 4b vs 4a). That is expected: those maps are the reference point, after a publishable-map edit. At Connaught Creek, which was drawn with no seed, the working classifier is closer to the consensus than autoATES v2.10 was (0.73 vs 0.69).

Two independently drawn maps, Connaught Creek and Bow Summit, come from autoATES v2.0 validation and were held out of fitting. They are the only sites with several mappers on the same ground. How much those experts agree with each other is the noise floor: it is the part of any model-versus-expert score that is ordinary disagreement between qualified mappers, not model error. That floor is not a single number. It is 0.72 at Connaught Creek and 0.89 at Bow Summit, and the grey band on Figure 4a is that envelope of disagreement. Only Connaught Creek has a scored v3.0 map so far. The working classifier agrees with that consensus at 0.73, inside the envelope at that site: the model matches the consensus about as well as the three experts match each other. Bow Summit is the planned second independent test once the typical and infrequent stack is built for that site.

That limit does not mean every new area needs a three-expert map. Those maps are time- and cost-intensive, and they are rarely available (Sykes et al. 2026). Their job here is to measure the noise floor so a model-versus-expert score can be read. The path to validation data in new locations is the one used for the commissioned maps: start from the automated product and have a local expert edit where the model is wrong. That is operational validation, and it is how the reference set can grow.

5. NEXT STEPS

Production maps will cover the Avalanche Canada forecast regions at the global 30 m DEM, and will be hosted by Avalanche Canada for free download and third-party integration. The working pipeline has been run end-to-end on domains of several thousand square kilometres (North Rockies, 5,246 km², in 43 minutes), which is enough to show that the 400,000 km² target is computationally feasible. A best-practices document is being developed with an expert panel to guide application of the maps and to state their limitations, including the surface-model bias over forest and the appropriate use of 30 m products. The working classifier still expects both scenarios. An explicit typical-only mode, emitting classes 1–4 with Simple as the study-area default, is planned so the method can be used without adopting the Canadian two-scenario setup. Future work includes an Avalanche Canada trip-planning tool, which is also where route-scale interaction with the maps can return. This research demonstrates a path to consistent ATES mapping across western Canada. Differences among automated approaches are real (Sykes et al. 2026); this classifier is built so regional differences in start-zone and runout modelling enter through the typical and infrequent scenarios, rather than by re-fitting the classification model in each range. The aim is a defensible, flexible product, not a closer fit to any one expert map.

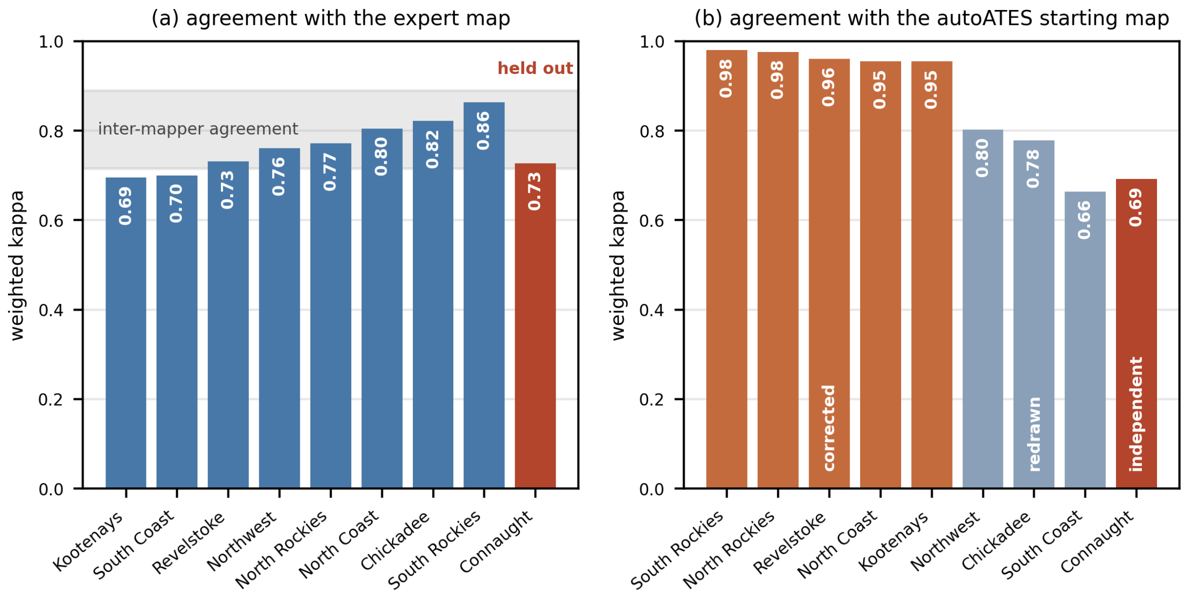

Figure 4: (a) Classifier agreement with each expert map, as weighted kappa. Connaught Creek was held out of fitting. The grey band is the envelope of inter-mapper agreement: 0.72 at Connaught Creek and 0.89 at Bow Summit. (b) Agreement between each expert map and the autoATES version that mapper started from. Five maps were edited in place toward a publishable product and remain close to that starting map; three were redrawn more freely; Connaught Creek was drawn independently, with no starting map.

ACKNOWLEDGEMENTS

We thank Avalanche Canada, the expert mappers who supplied benchmark maps, and the forecasters, ski-area programs and transportation mitigation programs that provided validation data. Financial support for this study was provided by Public Safety Canada through the Search and Rescue New Initiatives Fund (SAR-NIF) grant “Large-scale avalanche terrain exposure scale mapping for mountainous regions in Canada”.

REFERENCES

Atkins, R., Sykes, J., Rigby, N., and Haegeli, P. (2026). A closer look at slope angles of accidentally triggered persistent and deep persistent slab avalanches. In: Proceedings of the International Snow Science Workshop. Whistler, Canada.

Bebi, P., Bast, A., Helzel, K., Schmucki, G., Brozova, N., and Bühler, Y. (2021). Avalanche protection forest: from process knowledge to interactive maps. In: Protective forests as ecosystem-based solution for disaster risk reduction (Eco-DRR). Ed. by Teich, M., Accastello, C., Perzl, F., and Kleemayr, K. London: IntechOpen. https://doi.org/10.5772/intechopen.99514.

Hesselbach, C., Spannring, P., Huber, A., and Fischer, J.-T. (2026). OpenNHM/AvalancheTerrainIdentification: Avalanche Terrain Identification. Zenodo. https://doi.org/10.5281/zenodo.22049056.

Knies, D., Sykes, J., and Haegeli, P. (2026). An automated method for regional avalanche alpha-angle characterization based in western Canada. In: Proceedings of the International Snow Science Workshop. Whistler, Canada.

Larsen, H. T., Hendrikx, J., Slåtten, M. S., and Engeset, R. V. (2020). Developing nationwide avalanche terrain maps for Norway. In: Nat. Hazards 103, pp. 2829–2847. https://doi.org/10.1007/s11069-020-04104-7.

Statham, G. and Campbell, C. (2025). The Avalanche Terrain Exposure Scale (ATES) v.2. In: Nat. Hazards Earth Syst. Sci. 25, pp. 1113–1137. https://doi.org/10.5194/nhess-25-1113-2025.

Statham, G., McMahon, B., and Tomm, I. (2006). The avalanche terrain exposure scale. In: Proceedings of the International Snow Science Workshop. Telluride, CO, pp. 491–497.

Sykes, J., Harvey, S., Spannring, P., Hesselbach, C., Markov, K., Haegeli, P., and Statham, G. (2026). Characterizing differences among automated avalanche terrain mapping approaches: a multi-institution comparison at Rogers Pass, Canada. In: Proceedings of the International Snow Science Workshop. Whistler, Canada.

Sykes, J., Toft, H., Haegeli, P., and Statham, G. (2024). Automated Avalanche Terrain Exposure Scale (ATES) mapping – local validation and optimization in western Canada. In: Nat. Hazards Earth Syst. Sci. 24, pp. 947–971. https://doi.org/10.5194/nhess-24-947-2024.

Toft, H. B., Sykes, J., Schauer, A., Hendrikx, J., and Hetland, A. (2024). AutoATES v2.0: Automated Avalanche Terrain Exposure Scale mapping. In: Nat. Hazards Earth Syst. Sci. 24, pp. 1779–1793. https://doi.org/10.5194/nhess-24-1779-2024.