MACHINE LEARNING FOR AUTOMATED AVALANCHE TERRAIN EXPOSURE SCALE CLASSIFICATION – A COMPARISON OF TWO APPROACHES

Kalin Markov[1,5]*, Yves Bühler[2,7], Andreas Huber[6], Momchil Panayotov[3,4,5]

[1] National Institute of Geophysics, Geodesy and Geography, Bulgarian Academy of Science, Sofia, Bulgaria
[2] WSL Institute for Snow and Avalanche Research SLF, Davos Dorf, Switzerland
[3] University of Forestry, Sofia, Bulgaria
[4] Bulgarian Extreme and FreeSkiing Association - BEFSA, Sofia, Bulgaria
[5] Bulgarian Avalanche Association, Sofia, Bulgaria
[6] Austrian Research Centre for Forests (BFW), Department of Natural Hazards, Innsbruck, Tyrol, Austria
[7] Climate Change, Extremes and Natural Hazards in Alpine Regions Research Centre CERC, Davos, Switzerland

ABSTRACT: Backcountry travel in mountainous regions in Bulgaria has increased significantly in recent decades, heightening the need for effective avalanche risk management. The Avalanche Terrain Exposure Scale (ATES), first introduced in 2004 in Canada, classifies mountain terrain based on its potential exposure to avalanche hazards. Whereas early applications relied on expert-based manual mapping, subsequent studies introduced semi-automated and fully automated workflows. The most recent approaches introduce machine learning methods for automated terrain classification according to ATES criteria.

In this study, we compared two machine learning approaches (ATES 1 and ATES 2) for ATES classification for the territory of Pirin National Park, Bulgaria, one of the most popular regions in the country for backcountry skiing. Both approaches used the same training data, a Random Forest model with the same architecture and hyperparameters, and shared input features of slope and percent canopy cover. They differed in the remaining input features used for training procedures and they also used different approaches for determining the potential avalanche release areas (PRA). In addition, in model ATES 1, the runout modelling was performed using AvaFrame/com4FlowPy, while in ATES 2 - from RAMMS::Avalanche runout modelling.

The resulting ATES models were evaluated on a set of 36 separate validation polygons distributed throughout the territory of Pirin National Park. The validation polygons were labelled by three independent experts and located outside the valley where training occurred. Both models performed similarly well with respect to calculated performance metrics. In model ATES 1, the pixel-based accuracy was 87%, with 94% of the polygons classified correctly (when taking the majority classification for each polygon). For model ATES 2, both metrics were very similar, with the pixel-based accuracy being 86% and the polygon-based – 92%. The most classification errors occurred in the challenging terrain class for both models.

The results demonstrate the ability of machine learning models to generalize well across entire mountain ranges and to adapt robustly to varying input feature sets. At the same time, they indicate that mapping the challenging ATES terrain class remains difficult for predictive models and is therefore associated with the most errors.

KEYWORDS: avalanches, ATES, machine learning, com4FlowPy, RAMMS, Bulgaria

1. INTRODUCTION

In the past decades, there has been an increase in backcountry travel in the mountains in Bulgaria. Unfortunately, there have also been accidents, some of which fatal (Panayotov et al., 2021). For this reason, developing solutions for risk management in avalanche terrain is of essential importance. One tool that can be used for aiding route planning in avalanche terrain is the avalanche terrain exposure scale (ATES). ATES is a popular, globally-used scale to classify avalanche terrain based on its potential exposure to avalanches. It is only based on terrain characteristics and does not take into account current meteorological or snowpack conditions. It was created in Canada in 2004 by Statham et al. (2006) and recently its second version has been released (Statham and Campbell, 2025). The second version distinguishes between the following

* Corresponding author address:
Kalin Markov, National Institute of Geophysics, Geodesy and Geography, Bulgarian Academy of Science, Sofia, Bulgaria
tel: +359988917295
email: kalin.markov@gmail.com

terrain classes – non-avalanche, simple, challenging, complex, and extreme avalanche terrain.

In the beginning, spatial ATES maps as well as linear route ATES classifications were mainly produced manually (Parks Canada Agency, 2017). This process is extremely time consuming, especially for larger regions, and requires that experts be very familiar with the local mountains, which is difficult to achieve. Automated approaches using GIS pipelines have since been developed – including AutoATES 1.0, AutoATES 2.0 and adaptations of them to different regions (Huber et al., 2023; Larsen et al., 2020; Markov and Panayotov, 2024; Markov et al., 2024; Panayotov et al., 2024; Sykes et al., 2024; Toft et al., 2024; von Avis et al., 2023).

The most recent automated ATES approach uses machine learning for creating the maps (Markov et al., 2025). The approach mirrors that of AutoATES 2.0 and consists of the same PRA identification based on Veitinger et al. (2016) and runout modelling steps using AvaFrame/com4FlowPy (D'Amboise et al., 2022). The final classification stage uses a Random Forest classifier instead of predefined rules and thresholds.

In this paper, we first took the approach from Markov et al. (2025) and expanded the study region, applying the approach to the entire territory of Pirin National Park, Bulgaria. We call this approach ATES 1. Next, we performed a different modelling approach using a different PRA algorithm and then RAMMS::Avalanche runout modelling, substituting a simple-empirical model (com4FlowPy) with a more complex physically-based model (Bühler et al., 2013; Bühler et al., 2018; Bühler et al., 2022; Sykes et al., 2022). This approach generates a different set of input features. A model retraining was performed, after which the newly trained model was again applied to the entire territory of Pirin National Park to generate an ATES map. We call this approach ATES 2. Finally, we extended the validation procedure described in Markov et al. (2025) by creating an additional set of expert-labelled polygons scattered around the territory of the national park, entirely separate from the training data. Both ATES approaches were validated using these polygons.

With this study design, we aimed to address two main research questions:

1) How does the Random Forest AutoATES classifier generalize to areas that have not been used for training, but with similar topographical characteristics?
2) How well is the Random Forest AutoATES able to adapt to different input features fed into it, specifically with changed PRA and runout models?

2. STUDY SITE AND METHODS

2.1 Study area

The study area encompasses the entire Pirin National Park, located in southwestern Bulgaria. The national park occupies the northern part of the mountain range, including its highest peak Vihren (2914 m a.s.l.). The total area of the national park is 40 356 ha. The region consists of multiple prominent glacially-carved valleys, with relative vertical relief up to 1000 m. The lower elevations are predominantly forested, while locations between 2200 m a.s.l and 2500 m a.s.l are often covered by dwarf mountain pine (Pinus Mugo). The mountain summits are usually rocky, often with steep slopes, and sometimes contain cliff faces and rock bands. Avalanche danger is present in most areas and reaching most high peaks, even sometimes huts and shelters, requires travelling through avalanche terrain (Pirin National Park, 2022). The training data from Markov et al. (2025) is located in the Bunderitsa valley, which is one of the major valleys in the area and also the most frequented one due to easy access from Bansko Ski Resort (located underneath Todorka Peak). The validation polygons are spread out in the remaining part of the national park. Figure 1 shows the entire study region, its location within Bulgaria, and also the training and validation data locations.

2.2 Input data and methods

Automated algorithms for ATES map generation require two main input data layers – a digital terrain model and a layer describing the structure of the forest. In this study, we used a 5 m resolution digital terrain model from Pirin National Park (Pirin National Park, 2022), based on digitized maps with a scale of 1:5000. We used a percent canopy cover layer from Copernicus (Copernicus, 2020). in order to describe the forest structure, with a resolution of 10 m. We up-sampled the forest layer using bilinear interpolation so that it has a resolution of 5 m and matches the digital terrain model exactly. Since we used the model chain developed by Markov et al. (2025), the machine learning model requires training data to learn the relationship between input features and ATES classes. The training data was obtained directly from that study.

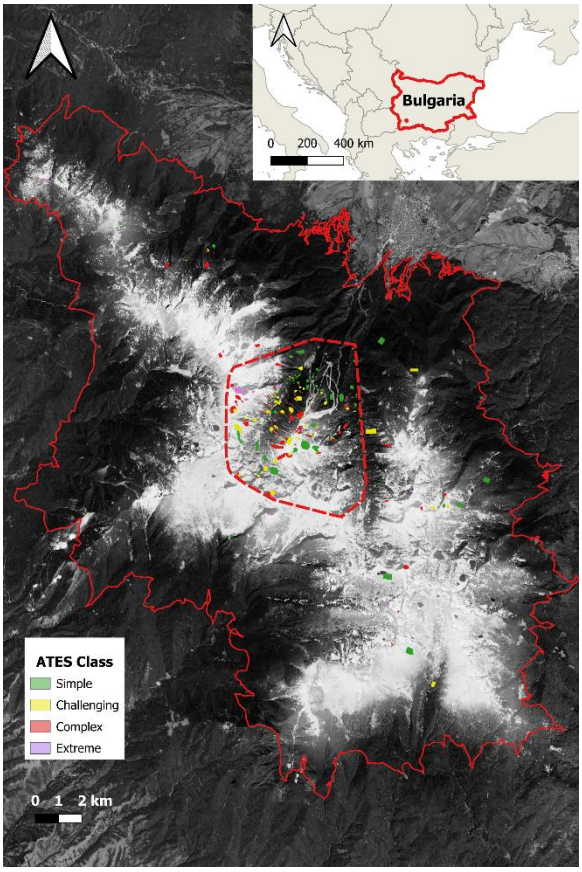


Figure 1: The entire study region and its location within Bulgaria. The dashed lines outline Bunderitsa valley, where all training polygons are located in. The small shapes represent the training and validation polygons.

2.2.1 Approach ATES 1

The first step involved running the algorithm from Markov et al. (2025) exactly as is, except for the fact that we performed its steps with the native 5 m resolution rather than the resampled 10 m resolution used in that paper. We named this approach ATES 1. In that workflow, the PRA calculation and runout modelling mimics that of AutoATES 2.0. The PRA threshold for binarization of the continuous PRAs was chosen to be 0.3, in accordance with Markov et al. (2025). com4FlowPy was run using the forest friction plugin, with parameters being chosen after a small case study for the region (Markov et al., 2025). The input features used for the classification were slope, percent canopy cover, continuous PRA ranging from 0 to 1, and the com4FlowPy outputs of local travel angle, cell counts, maximum modelled avalanche intensity per cell, and the sum of routed flux. These correspond to model RF2 described in Markov et al. (2025). The retraining in this study was required because the whole automated workflow from the previous study was performed using 10 m resolution while we chose to run that workflow with 5 m resolution to preserve more terrain details and match the original input DEM's resolution. We then applied the retrained model to the entire territory of Pirin National Park.

2.2.2 Approach ATES 2

The second part of our study (ATES 2) involved using a different approach to calculate the PRAs and runout modelling using the Swiss Large Scale Hazard Indication Modelling (LSHIM) approach, which uses RAMMS::Avalanche instead of com4FlowPy. The PRA delineation was performed with the method described by Bühler et al. (2018). The algorithm uses an object-based image analysis approach and derives several data layers from the input DTM, such as slope, aspect, curvature, and ruggedness. The algorithm requires a binary forest/non-forest layer as well, which we obtained by taking the input percent canopy cover layer and converting all pixels with values of 0% to non-forest, and all other pixels to the forest class. A series of consecutive processing steps were executed to calculate the PRAs, where a few iterations of further segmentation were performed to delineate individual PRAs, dependent on the required scenario (return periods 10, 30, 100 and 300 years according to the Swiss guidelines). The runout simulations were then performed using these PRAs, the binary forest layer, and the DTM, which were fed into RAMMS::Avalanche. We ran the avalanche simulations using a 30-year return period and a release depth of 0.5 m to estimate more frequent scenarios, as the previous ATES maps for the region and most Swiss avalanche terrain hazard maps aim to simulate typical skier-triggered avalanches around size 3.

The procedure described resulted in a different set of input features compared to ATES 1. The only shared input features with approach ATES 1 were slope and percent canopy cover. The next input feature was the PRAs, which were now binary (1 or 0) instead of continuous, as approach ATES 2 produces this type of PRA. The remaining input features were simulation outputs from RAMMS: cell counts, max flow depth, max pressure, and max velocity. After obtaining all these input features, we retrained the random forest model from Markov et al. (2025), using the exact same model architecture, hyperparameters, and training data, just using the different input features. The purpose in keeping everything else constant was to assess the impact of only the different input features and test how the machine learning model can adapt to them. We finally applied the trained model to the entire territory of Pirin National Park.

2.2.3 Validation

We validated both ATES 1 and ATES 2 on a separate set of 36 validation polygons. The polygons were delineated to be small and consist of confined and homogenous terrain forms based on aerial imagery and inclination maps. They were then given to three independent local experts to classify based on their expert opinion and years of experience in the area. Then, the most commonly classified ATES class (the mode) from all three experts was assigned to each polygon. This step was necessary, as there was not always complete agreement for all polygons between the experts.

The pixel-based statistics calculated on the validation polygon pixels included accuracy, precision, recall, F1 score, and a precision-recall curve, computed as in Rainio et al., 2024, using the Python programming language (Python Software Foundation, 2024). These metrics were calculated for both ATES approaches by taking the pixels in each validation polygon and comparing their ATES class with the corresponding pixels in the maps from ATES 1 and ATES 2. A group-based accuracy and confusion matrix was also calculated, where for each polygon, the majority predicted ATES class from ATES 1 and ATES 2 was calculated (so from all pixels in the polygon, what was the most commonly predicted one) and compared with the actual label of the polygon.

3. RESULTS

Both approaches ATES 1 and ATES 2 produced very similar ATES maps for the entire Pirin Mountain range (Figure 2). The bottom part of Figure 2 shows a snippet of the ATES classification for a steep, forested region underneath Hleven Peak (southeastern part of the national park) as well as a photo of the same region, highlighting that both approaches capture challenging forested terrain well, but with some differences in the ATES classifications of forest openings.

The side-by-side comparison of the precision, recall, and F1 score values for all pixels of the validation polygons for each ATES class for both approaches shows very good performance for most ATES classes (Table 1). The F1 score, which is a measure of the harmonic mean between precision and recall, shows that the challenging class was the most difficult to delineate in both cases.

Table 1: A side-by-side comparison of the precision (P), recall (R), and F1 score values for all pixels of the validation polygons for each ATES class for both approaches

	ATES 1			**ATES 2**		
ATES Class	P	R	F1	P	R	F1
Simple	0.946	0.937	0.942	0.901	0.939	0.920
Challenging	0.755	0.680	0.715	0.736	0.670	0.701
Complex	0.812	0.903	0.855	0.854	0.884	0.869
Extreme	0.985	0.953	0.969	0.984	0.903	0.941

Figure 3 shows the precision-recall curve per ATES class for both approaches. This metric also demonstrates that the challenging class is the most difficult to classify, while the simple and extreme classes are most successfully classified. Approach ATES 1 is more confident than ATES 2 in the classification for the challenging terrain class.

Table 2 shows the confusion matrices for the polygon classification for both approaches. In model ATES 1, the overall per-pixel accuracy was 87%, with 94% of the polygons classified correctly (when taking the majority classification for each polygon). For model ATES 2, both metrics were very similar, with the pixel-based accuracy being 86% and the polygon-based – 92%. Again, there was some disagreement involving the challenging terrain class for both models.

4. DISCUSSION

The results show that machine learning models, such as the random forest model used here, are capable of generalizing well to entire mountain ranges. As Figure 1 shows, the training data is only from one valley (Bunderitsa valley), while the validation data is from the remaining part of the mountain range. Despite this, both approaches ATES 1 and ATES 2 perform very well on the validation data. The precision, recall, and F1 scores are high for most classes, as well as the overall accuracy when looking at the pixel and polygon level. This result can be expected, as machine learning models have the strength that they are able to learn local learn from local patterns and then apply what they learned successfully to other, unseen-before regions The Bunderitsa valley, where training occurred, is representative of the entire Pirin mountain range because it has very diverse terrain with both open, steep slopes,

forested steep slopes, and high-mountain type alpine terrain. Therefore, a model trained there can generalize well to the new terrain it is applied to.

Table 2: Confusion matrices for the per-polygon classification for both ATES approaches. In total, there are 36 validation polygons. The results are almost identical with the exception of 1 polygon misclassified as simple in approach ATES 2, when its true label (according to the expert classification) is challenging.

True ATES label						
	Approach ATES 1	Simple	9	0	0	0
		Challenging	0	7	0	0
		Complex	0	2	11	0
		Extreme	0	0	0	7
	Approach ATES 2	Simple	9	0	0	0
		Challenging	1	6	0	0
		Complex	0	2	11	0
		Extreme	0	0	0	7
			Simple	Challenging	Complex	Extreme
Predicted ATES label						

Another main conclusion that can be drawn is that machine learning models are robust and can adapt to a variety of input features fed into them. This is a major advantage over traditional, rule and threshold-based approaches, such as AutoATES 2.0 There, are new set of features would mean that experts have to define new rules and thresholds for each new feature and then figure out how to combine them into ATES classes. In contrast, in this study, approach ATES 2 was very easily run and only required a retraining. As the figures in the results section indicate, the final ATES maps and statistical validation using the validation polygons show that performance is very similar with approach ATES 1. This proves the robustness of ML models for ATES classification and that as long as meaningful, representative features are fed into the model, it can learn the relationship between them and the ATES classes and the end result is more or less the same. Manual rule-based models struggle with this and are significantly influenced by changing just one small rule or threshold, as demonstrated in Panayotov et al. (2024). The ATES maps produced from those approaches can vary significantly just from one minor change.

The best performance is for the simple and extreme terrain classes for both approaches, which can be expected, as these terrain types represent either one or the other end of the spectrum and are more easily distinguishable from the rest, causing less confusion. The challenging terrain class has often been the most difficult to classify (Markov et al., 2025), as it represents a wider variety of terrain forms. This could be flat terrain underneath a large, steep slope, very steep terrain that is forested, or short, moderately steep open sections. This variability is what makes classifying such terrain more difficult. Approaches ATES 1 and ATES 2 both experience more difficulties with challenging terrain than the other classes, but the agreement with the expert-labelled polygons is still relatively strong. One example of challenging terrain classification can be seen from the bottom part of Figure 2 (panels c-f). The validation polygon there is located at the bottom section of the steep, sparse forest, where it starts to level off as it enters the valley and is appropriately labelled as challenging terrain. Both approaches capture the challenging terrain in this region fairly well (labels 1 and 3). ATES 1 also captures all the steep forest openings and gullies higher upslope well, labelling them as complex terrain (label 2). Approach ATES 2 labels less regions as complex in the same slope. Since the original percent canopy cover layer captures most of these openings correctly, the reason for the differences in the final ATES map likely stems from how the forest cover is handled in both approaches. In ATES 1, the percent canopy cover is considered as a continuous percentage value when supplied to the PRA calculation, while in ATES 2, it is required to be a binary mask (where 0% is cast to non-forest, everything else – forest). This casting may lose some information in narrow forest gullies and openings such as those in this slope, where there are openings but the PCC is not 0%. The original Swiss Large Scale Hazard Indication Mapping use a more sophisticated approach for modelling forests, taking into account tree heights, forest cover and gap dimensions, which may result in more similar final results to approach ATES 1. Also, approach ATES 1 takes the continuous PRA value as one of its input features for classification, while ATES 2 takes the binary PRA layer (as that is the format it works with), which again may contribute to some loss of information in narrow forest openings.

Although the precision, recall, and F1 score are similar in both cases for the challenging terrain class, there is a larger difference in the area under the precision-recall curve. This means that

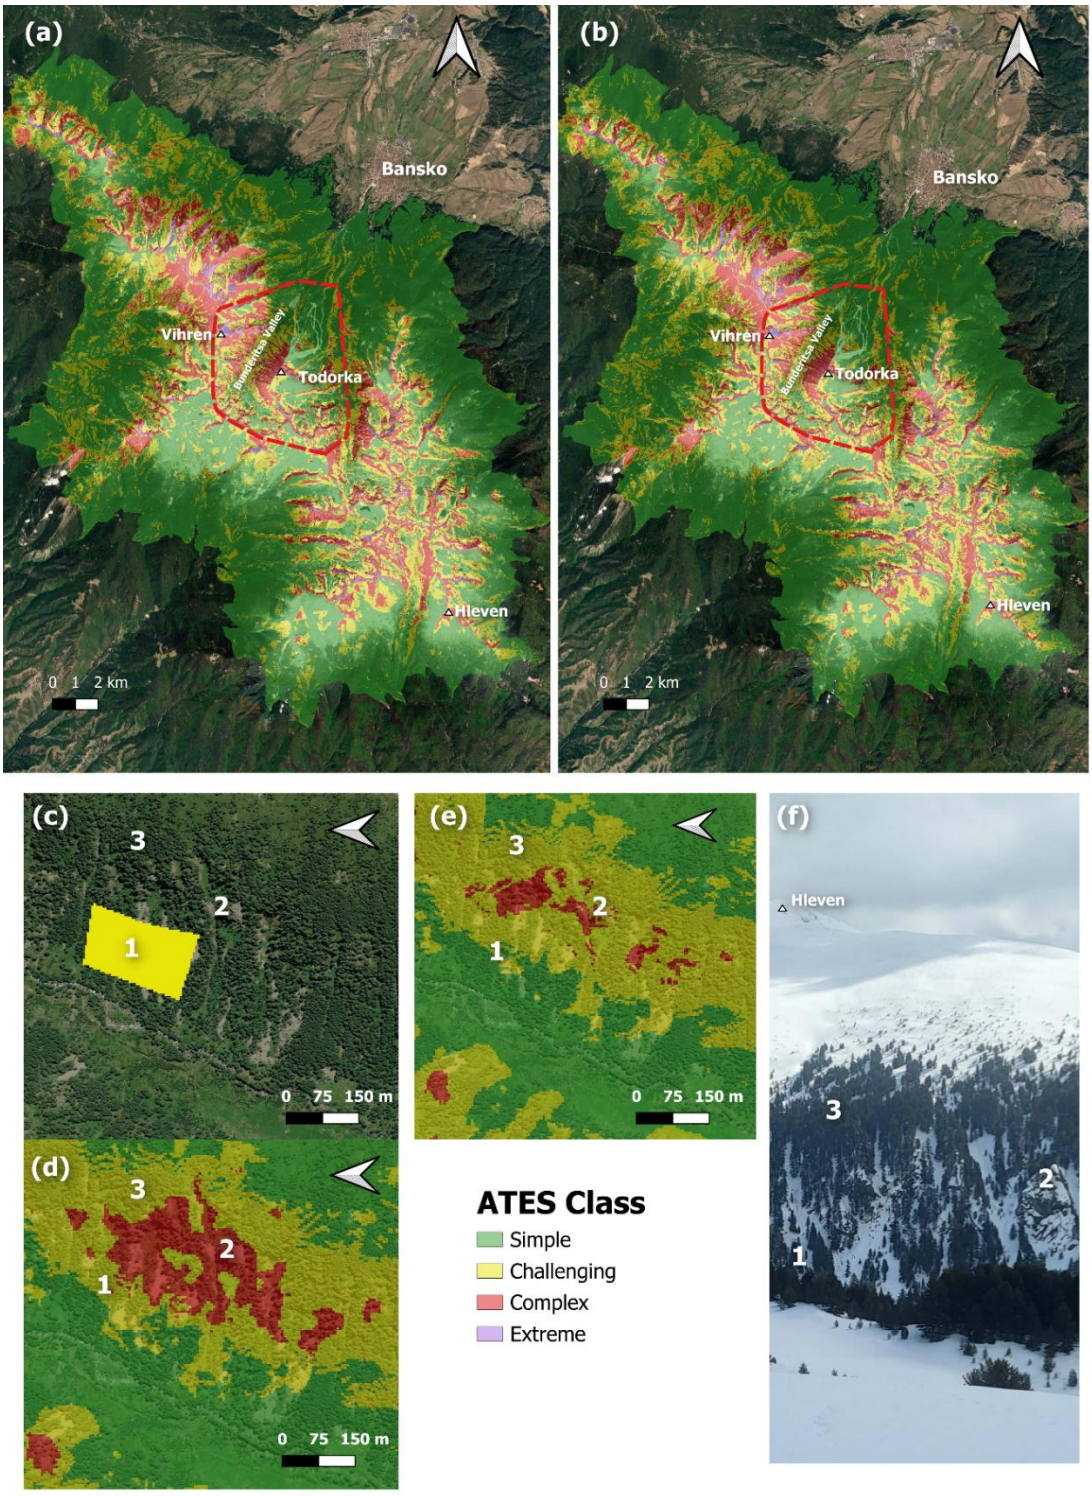


Figure 2: ATES maps of the entire Pirin National Park from approaches ATES 1 (a) and ATES 2 (b). The dashed lines outline the Bunderitsa Valley, where the training polygons are located. The bottom part of the figure (c-f) shows a snippet of the ATES classification for a steep, forested region on the west face of Hleven Peak, where (c) shows the aerial imagery and the challenging class validation polygon located there (yellow rectangle), (d) shows the corresponding ATES map from approach ATES 1, (e) shows the corresponding ATES map from ATES 2, and (f) shows a photograph of the same region, looked at from the same angle. The numbers 1, 2, and 3 indicate corresponding locations in (c) – (f).

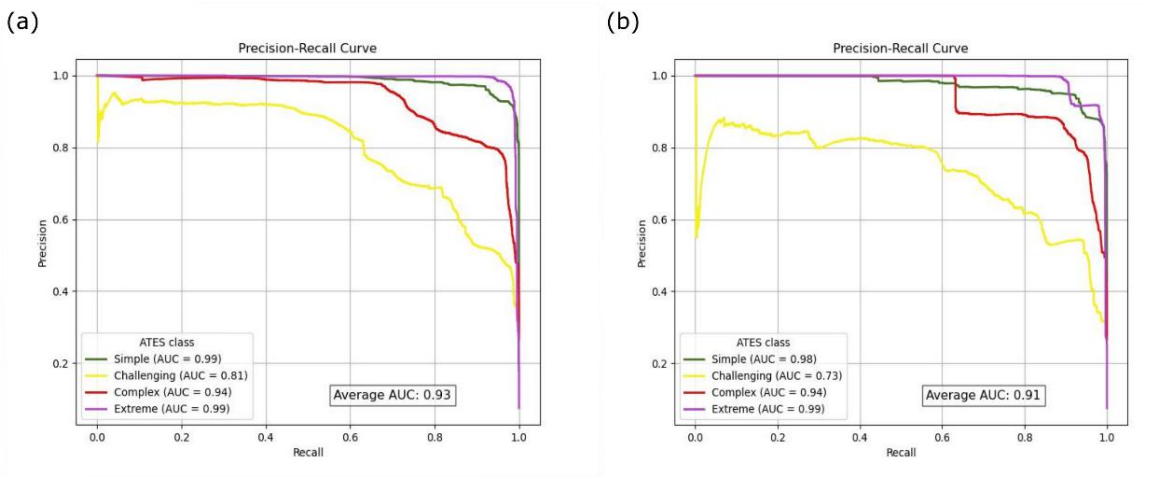


Figure 3: The precision-recall curve per ATES class for both approaches ATES 1 (a) and ATES 2 (b). AUC stand for "area under the curve".

ATES 2 is slightly less confident in its predictions for the challenging terrain class. This may also stem from the differences in forest structure and PRA modelling described earlier.

It is also important to note the general limitations in ATES mapping. ATES is a simple scale that is intended to be easy to understand, but its use of only a few terrain classes sometimes causes a loss of information. For example, a mountainous region labelled as the challenging ATES class could represent a short, moderately steep alpine section, but also a flat valley area underneath a large, steep slope. ATES users that are not familiar with the terrain have no way of knowing this looking at just the maps. That is why in some places, such as Switzerland, CAT and ATH maps are preferred instead of ATES (Harvey et al., 2024), as they depict more avalanche-related details.

Also, the validation performed in this study uses isolated, small polygons instead of a continuously drawn ATES map. Although it is a good baseline validation for how the models generalize to unseen territories, results may differ to some extent if evaluated against a continuous ATES map.

5. CONCLUSION

This study shows that random forest models can successfully and easily take avalanche expert knowledge from just one valley and generalize it reliably across an entire mountain range. The random forest approach proved to be highly flexible and robust. It adapted easily to completely different sets of input features with just a simple retraining, while still producing high-quality maps. This is a major advantage over older, rule-based systems that require experts to manually figure out new thresholds and logical rules every single time a new feature is added.

Both models behaved similarly by confidently classifying simple and extreme avalanche terrain while experiencing more difficulties with challenging terrain due to the higher variability in the types of terrain it encompasses, but still performing well there as well. Overall performance was very similar between models ATES 1 and ATES 2 and they both generalized well to new territories.

In conclusion, using automated machine learning pipelines creates highly-quality ATES maps that are easily adaptable to different regions and different input features fed into them. The resulting maps are more robust and less drastically influenced by changes to input features, as they don't use manual rules and thresholds (in contrast to previous AutoATES classifiers). These approaches can be applied to larger territories more easily and significantly help manage avalanche risk in backcountry regions.

ACKNOWLEDGEMENTS

The study was carried out in the framework of project "AvalForCE" funded by Multilateral Academic Projects (MAPS) of the Swiss National Science Foundation and the Bulgarian Ministry of Education grant IZ11Z0_230113. We acknowledge the use of the Copernicus Tree Density product.

REFERENCES

Bühler, Y., Bebi, P., Christen, M., Margreth, S., Stoffel, L., Stoffel, A., Marty, C., Schmucki, G., Caviezel, A., Kühne, R., Wohlwend, S., and Bartelt, P.: Automated avalanche hazard indication mapping on a statewide scale, Nat. Hazards Earth Syst. Sci., 22, 1825–1843, https://doi.org/10.5194/nhess-22-1825-2022, 2022.

Bühler, Y., Kumar, S., Veitinger, J., Christen, M., Stoffel, A., and Snehmani: Automated identification of potential snow avalanche release areas based on digital elevation models, Nat. Hazards Earth Syst. Sci., 13, 1321–1335, https://doi.org/10.5194/nhess-13-1321-2013, 2013.

Bühler, Y., von Rickenbach, D., Stoffel, A., Margreth, S., Stoffel, L., and Christen, M.: Automated snow avalanche release area delineation – validation of existing algorithms and proposition of a new object-based approach for large-scale hazard indication mapping, Nat. Hazards Earth Syst. Sci., 18, 3235–3251, https://doi.org/10.5194/nhess-18-3235-2018, 2018.

Christen, M., Bartelt, P., Kowalski, J., and Stoffel, L., 2010. Calculation of dense snow avalanches in three-dimensional terrain with the numerical simulation program RAMMS, Cold Reg. Sci. Technol., 63, 1–14, 10.1016/j.coldregions.2010.04.005

Copernicus: Tree Cover Density 2018 (raster 10 m and 100 m), Europe, 3-yearly, European Environment Agency (EEA) [data set], https://doi.org/10.2909/486f77da-d605-423e-93a9-680760ab6791, 2020

D'Amboise, C. J. L., Neuhauser, M., Teich, M., Huber, A., Kofler, A., Perzl, F., Fromm, R., Kleemayr, K., and Fischer, J.-T.: Flow-Py v1.0: a customizable, open-source simulation tool to estimate runout and intensity of gravitational mass flows, Geosci. Model Dev., 15, 2423–2439, https://doi.org/10.5194/gmd-15-2423-2022, 2022.

Harvey, S., Christen, M., Bühler, Y., Hänni, C., Boos, N., and Bernegger, B.: Refined Swiss avalanche terrain mapping CATV2 / ATHV2, in: Proceedings, International Snow Science Workshop, Tromsø, Norway, 2024, pp. 1637–1644, https://arc.lib.montana.edu/snow-science/item/3363, 2024.

Huber, A., Hesselbach, C., Oesterle, F., Neuhauser, M., Adams, M., Plörer, M., Stephan, L., Toft, H., Sykes, J., Mitterer, C., and Fischer, J.-T.: 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, OR, USA, 2023, Innsbruck, Austria, 1272–1278, http://arc.lib.montana.edu/snow-science/item/2989 (last access: 16 March 2025), 2023.

Markov, K. and Panayotov, M.: Defining of avalanche terrain with ATES modelling for the region of Bansko Ski Resort in Pirin, in: Avalanches in the Bunderitsa valley in Pirin Mountains, Intel Entrans, Sofia, ISBN 978-619-7703-60-3, 2024.

Markov, K., Huber, A., Panayotov, M., Hesselbach, C., Spannring, P., Fischer, J.-T., and Teich, M.: Machine learning for automated avalanche terrain exposure scale (ATES) classification, Nat. Hazards Earth Syst. Sci., 25, 4375–4403, https://doi.org/10.5194/nhess-25-4375-2025, 2025.

Markov, K., Panayotov, M., Tcherkezova, E., and Teich, M.: Identifying and mapping avalanche terrain using the ATES model for the region around Bansko Ski Resort, Pirin, Forest Science (Nauka za Gorata), 60, 99–122, 2024

Panayotov, M., Markov, K., Tsvetanov, N., Huber, A., Hesselbach, C., and Teich, M.: Avalanche hazard mapping using the avalanche terrain exposure scale (ATES) in the high mountain ranges of Bulgaria, in: Proceedings, International Snow Science Workshop 2024, Tromsø, Norway, https://arc.lib.montana.edu/snow-science/objects/ISSW2024_P12.5.pdf (last access: 21 April 2025), 2024.

Panayotov, M., Menteshev, B., and Mihaylova, S.: Avalanches in Bulgaria, BEFSA, ISBN 978-619-188-515-2, 2021.

Parks Canada Agency: Avalanche Terrain Exposure Scale – Avalanche Terrain Exposure Scale, Government of Canada, https://parks.canada.ca/pn-np/bc/glacier/visit/hiver-winter/ski/ates (last access: 29 October 2025), 2017.

Pirin National Park: Pirin National Park – World Heritage Property of Unesco, https://www.pirin.bg/?p=5591#more-5591 (last access: 28 August 2025), 2022.

Python Software Foundation: Python Language Reference, version 3.9.21 [code], https://www.python.org/ (last access: 22 March 2025), 2024

Rainio, O., Teuho, J., and Klén, R.: Evaluation metrics and statistical tests for machine learning, Sci. Rep.-UK, 14, 6086, https://doi.org/10.1038/s41598-024-56706-x, 2024.

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

Statham, G., McMahon, B., and Tomm, I.: The Avalanche Terrain Exposure Scale, in: Proceedings, International Snow Science Workshop 2006, Telluride, Colorado, https://arc.lib.montana.edu/snow-science/objects/issw-2006-491-497.pdf (last access: 5 May 2025), 2006.

Sykes, J., Haegeli, P., and Bühler, Y.: Automated snow avalanche release area delineation in data-sparse, remote, and forested regions, Nat. Hazards Earth Syst. Sci., 22, 3247–3270, https://doi.org/10.5194/nhess-22-3247-2022, 2022.

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

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

Veitinger, J., Purves, R. S., and Sovilla, B.: Potential slab avalanche release area identification from estimated winter terrain: a multi-scale, fuzzy logic approach, Nat. Hazards Earth Syst. Sci., 16, 2211–2225, https://doi.org/10.5194/nhess-16-2211-2016, 2016a

von Avis, C. D., Sykes, J. M., and Tutt, B.: Development of large scale automated Avalanche Terrain Exposure Scale (ATES) ratings in collaboration with local avalanche experts, in: Proceedings, International Snow Science Workshop 2023, Bend, Oregon, 982–988, http://arc.lib.montana.edu/snow-science/item/2998 (last access: 12 April 2025), 2023.