A CLUSTERING TECHNIQUE TO IDENTIFY SPATIAL PATTERNS IN SNOW COVER MODEL OUTPUT

Simon Horton[1,2*], Florian Herla[1], and Pascal Haegeli[1]

[1] *Avalanche Research Program, Simon Fraser University, Burnaby, BC, Canada*
[2] *Avalanche Canada, Revelstoke, BC, Canada*

ABSTRACT: Spatially distributed snowpack models are increasingly being adopted by avalanche forecasting agencies, but forecasters could benefit from simpler methods to interpret the complex model output. We introduce a statistical clustering technique that summarizes model output by predicting forecast regions with similar hazard characteristics. The method derives five metrics from simulated profiles to summarize different components of hazard: snow depth, new snow, wind-drifted snow, persistent weak layers, and wet snow. These metrics, along with spatial arrangement information, are combined into a distance metric that is fed into a fuzzy clustering algorithm to group small regions into larger spatially contiguous regions with distinct hazard characteristics. Our application of the method to western Canada during the 2023-24 winter produced regions that closely aligned with the regions in Avalanche Canada's daily public forecasts. The method's flexibility in considering various snowpack properties and spatial patterns makes it an attractive option for decision support.

KEYWORDS: snowpack modelling, clustering, forecast regions, spatial variability

1. INTRODUCTION

Numeric snowpack models have the potential to improve avalanche forecasts by augmenting information in data-sparse regions, providing continuous spatial coverage, and forecasting future conditions. However, incorporating these models into operational workflows requires presenting their complex output in simple and informative formats that can be processed quickly and easily (Morin et al., 2020).

Recent advancements have significantly enhanced snowpack models for avalanche forecasting. High-resolution numerical weather prediction (NWP) models offer relatively reliable meteorological inputs, post-processing tools can simplify outputs by comparing, averaging, and clustering snow profiles (Herla et al., 2021; Herla et al., 2022; Horton et al., 2024), and post-processing or machine learning models improve the connection to stability and avalanche hazard (Mayer et al., 2022; Pérez-Guillén et al., 2022).

Statistical clustering methods provide an effective way to analyze complex datasets. Bouchayer (2017) first clustered modelled snow profiles based on vertical profiles of specific surface area. Herla et al. (2021) expanded on this by incorporating multiple layer properties, such as grain type, hardness, and deposition date, into a distance metric for clustering. Horton et al. (2024) further developed this approach to account for spatial and temporal patterns, enabling the clustering of snow profiles into coherent forecast regions. However, these techniques are computationally intensive, as they process full snowpack stratigraphies, which limits their operational use.

To address this issue, this study builds on Horton et al. (2024) by deriving avalanche hazard metrics from snow profiles and clustering them into forecast regions. By clustering these metrics instead of full profiles, the process becomes more computationally efficient and scalable for large operational domains.

2. DATA

2.1 Study area

We applied the clustering method to 91 subregion polygons used by Avalanche Canada for public avalanche forecasts in the 2023-24 winter (Fig. 1). Forecasters group these subregions daily into larger forecast regions based on their assessment of regional-scale hazard patterns. The subregions cover a diverse range of terrain and avalanche climates (Shandro and Haegeli, 2018).

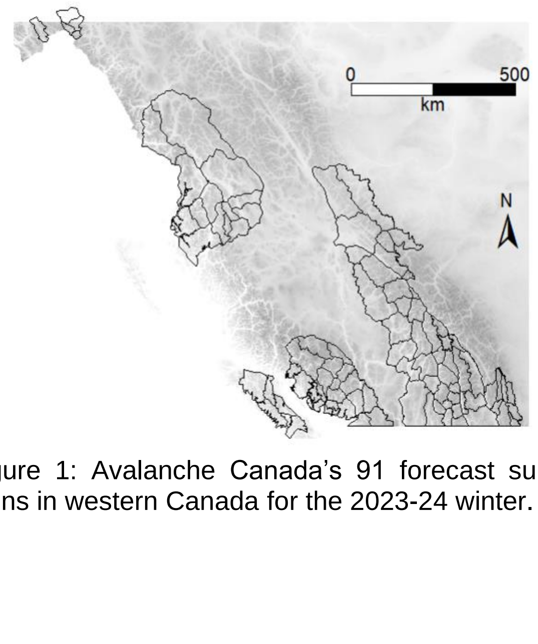

Figure 1: Avalanche Canada's 91 forecast subregions in western Canada for the 2023-24 winter.

* *Corresponding author address:*
Simon Horton; Avalanche Canada
email: shorton@avalanche.ca

2.2 Avalanche forecasts

We analyzed Avalanche Canada's public forecasts from December 1, 2023, to April 30, 2024. Over this period the 91 subregions were grouped into 9 to 16 forecast regions per day, with a median of 15 regions per day. We examined the arrangement of subregions into forecast regions, the danger rating, and avalanche problem types at the treeline vegetation band. To simplify the analysis, we grouped the avalanche problem types into analogous European Avalanche Warning Services problems (Table 1).

Table 1: Avalanche problem groups and hazard metrics.

European problem	*North American problem(s)*	*Hazard metric*
New snow	Storm slab Dry loose	Height of new snow in the past 72 hours
Wind-drifted snow	Wind slab	Average wind speed over the past 24 h multiplied by the skier penetration depth
Persistent weak layers	Persistent slab Deep persistent slab	Persistent weak layer depth multiplied by the probability of layer being unstable
Wet snow	Wet slab Wet loose	Total liquid water content in the profile

2.3 Snowpack model simulations

We used snow profile simulations from Avalanche Canada's operational snowpack modelling chain, which forces the Swiss SNOWPACK model with NWP forecast data from the Canadian High-Resolution Deterministic Prediction System (Horton et al., 2023). The model chain selects three grid points from each subregion polygon with a stratified sampling approach to get representative weather conditions at treeline elevations. SNOWPACK simulations are updated daily at these grid points using the latest NWP forecast, and then the averaging algorithm of Herla et al. (2022) produces a single snow profile for each subregion. Our analysis used average treeline profiles from each subregion, which were available for 147 days over the study period.

3. METHODS

We apply a clustering method that groups the 91 subregion polygons into forecast regions based on their geographic arrangement and simulated hazard properties. Following the approach of Horton et al. (2024), this method computes a distance matrix that quantifies the pairwise similarity between each subregion. The subregion distance (*dist*) is calculated as:

$$dist = 0.33\ dist_{spatial} + 0.67\ dist_{hazard}$$

where *dist* is a weighted average of the spatial distance ($dist_{spatial}$) and the hazard distance ($dist_{hazard}$). Weights of 0.33 and 0.67 produced forecast regions that were mostly spatially contiguous for our study area (see Horton et al., 2024 for details about optimizing spatial weights). A matrix of pairwise subregion distances is then inputted into a fuzzy clustering algorithm. Fuzzy clustering methods use a predetermined number of clusters (*k*) and then assign each data point a probability of belonging to each cluster.

3.1 Spatial distance

We designed the spatial distance ($dist_{spatial}$) to promote geographically contiguous forecast regions by reducing the distance between neighbouring subregions and increasing the distance for separated ones. Using a binary neighbourhood-based approach, subregions sharing borders were assigned a distance of 0, while those without shared borders were assigned a distance of 1. Our initial explorations found that the binary approach created forecast regions that follow the elongated shape of mountain ranges better than Euclidean distances.

3.2 Hazard distance

We designed the hazard distance ($dist_{hazard}$) to characterize predominant patterns in the simulated snow profiles. Simple scalar metrics were chosen to quantify the severity of each avalanche problem. Several potential metrics for each problem type were derived from the profiles. The best metric for each problem was selected using the Wilcoxon rank-sum test to compare days with and without corresponding avalanche problems across the entire dataset (Table 1).

New snow problems were quantified using the height of new snow in the past 72 hours. To create a normalized metric between 0 and 1, we chose an upper limit of 30 cm and linearly mapped values from 0 to 30 cm to the 0 to 1 range.

Wind-drifted snow problems were quantified using the product of average wind speed over the past 24 hours and the skier penetration depth. To normalize the metric, we set upper limits of 4 m/s for wind speed and 30 cm for skier penetration depth before the multiplication.

Persistent weak layer problems were quantified using the product of weak layer depth and the probability of instability. First, potential weak layers were identified using the quantitative module of avalanche hazard described by Herla et al. (2024). This method identifies unstable layers with the random forest model developed by Mayer et al. (2022). Then, unstable layers with persistent grain types were grouped by burial date. In cases with multiple potential weak layers, we used the median depth and probability. To normalize the metric, we set an upper limit of 100 cm for weak layer depth.

Wet snow problems were quantified with the total liquid water content in the profile with an upper limit value of 1 kg/m^2.

We used a total snow depth metric alongside the four avalanche problem metrics to create regions that align with major snow climates, which is a favourable attribute for forecast regions. We normalized snow depth by setting the maximum snow depth across all profiles as the upper limit.

We computed hazard distance by calculating Euclidean distance matrices for each metric and then computing their weighted mean with 25% weight on new snow, 25% on persistent weak layers, 8% on wind-drifted snow, 8% on wet snow, and 34% weight on snow depth. These weights were informed by a linear regression model that predicted danger ratings across the entire dataset and then further tuned with trial and error to produce realistic shaped regions.

3.3 Fuzzy analysis clustering

We used the fuzzy analysis clustering method included in the *cluster* package for R (Kaufman and Rousseeuw, 2009), testing *k* values from 5 to 20 to span the number of human-forecasted regions over the period. The method requires a fuzziness parameter that controls the degree of fuzziness in the clusters, so for each *k* value, we tried fuzziness parameters between 1.1 and 1.5 and then selected the solution that maintained the target number of clusters while maximizing fuzziness.

We applied the clustering method over the entire season, selecting the optimal number of regions each day according to the method outlined by Horton et al. (2024). This method calculates three cluster validation metrics for each clustering solution and selects the solution with the smallest value of *k* that meets a threshold for each metric. The final solution is the rounded average *k* value among the three metrics. Our threshold values were chosen to produce an average of 15 regions per day, which for this study were an average silhouette width of 0.36, a within-between ratio of 0.48, and a Pearson gamma of 0.50.

3.4 Cluster evaluation

We first present a case study for a single day where we compare the public forecast regions, danger, and problems with the clustering results, hazard metrics, and representative profiles from each cluster.

To compare the overall patterns between the clustering method and public forecasts, we counted how often each pair of subregions was grouped together over the season with each method. This created distance matrices, which we clustered using the fuzzy clustering method with *k* = 15 to determine the most common arrangements of forecast regions.

4. RESULTS

4.1 Case study on March 2, 2024

To illustrate the capabilities of the clustering method we present the results for March 2, 2024 (Fig. 2). On this day heavy snowfall in southern regions led to elevated danger ratings and avalanche problems related to new snow. In northern regions, the danger ratings were relatively lower, and the primary problems were related to wind-drifted snow. Persistent weak layers were the primary problem in central regions but also existed in most regions except for some coastal areas.

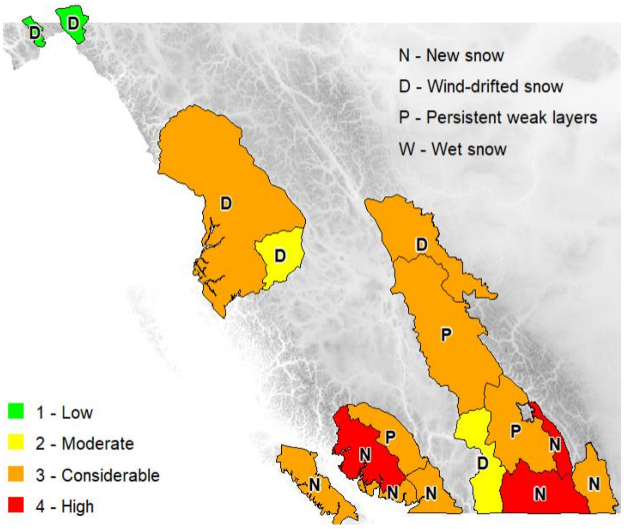

Figure 2: Public forecasts on March 2, 2024 with forecast regions coloured by treeline danger rating and labelled by primary avalanche problem type.

The hazard metrics derived from snowpack simulations show similar spatial patterns in avalanche problem distribution and severity (Fig. 3). Therefore, combining these metrics into a single distance metric is expected to effectively characterize hazard patterns across the study area.

The clustering results for different numbers of clusters on March 2, 2024 demonstrate the method's ability to group regions at various resolutions (Fig. 4). The transparent subregions in these plots highlight areas where the fuzzy clustering algorithm was uncertain about cluster membership, which were often near the boundaries of human forecast regions. The *k* = 15 clustering in Fig. 4 closely matches the public forecast regions in Fig. 2, with many boundaries aligning or nearly aligning with the public regions.

The *k* = 5 result in Fig. 4 is the easiest to interpret and compare with the hazard metrics in Fig. 3. The three southern regions are dominated by new snow problems and are distinguished by a main coastal region (purple), an interior region (orange), and a more scattered region in areas with lower snow depths and snowfall (green). The central interior region (blue) is dominated by persistent weak layer problems, while the northern areas (red) are where wind-drifted snow was the primary problem.

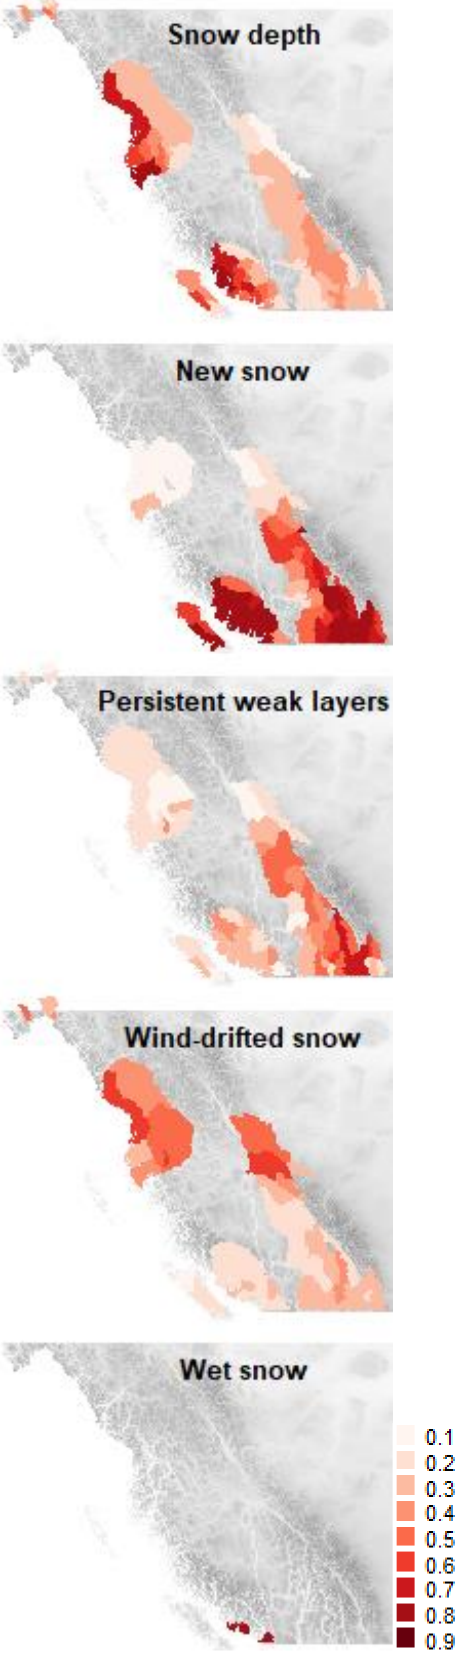

Figure 3: Hazard metrics derived from snowpack simulations on March 2, 2024. Each metric is normalized between 0 and 1 and displayed as a red heatmap.

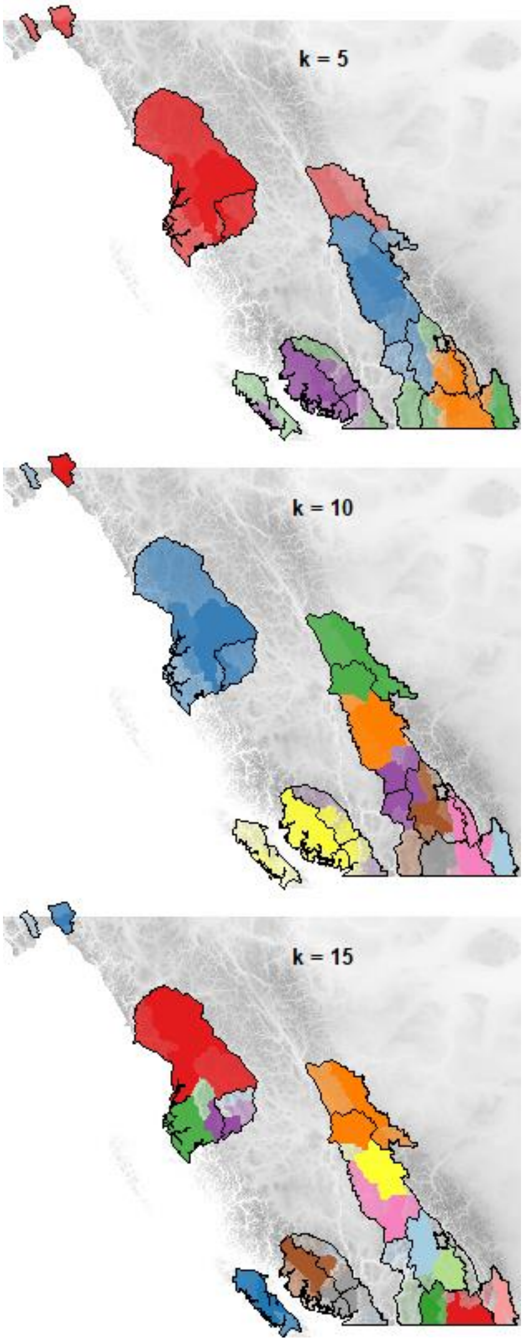

Figure 4: Clustering results for 5, 10, and 15 clusters on March 2, 2024. Subregion polygons are coloured by their primary cluster, with greater transparency when the probability of belonging to that cluster is low. The boundaries of the human-assessed forecast regions are outlined in black.

We can further characterize each cluster by plotting their simulated snow profiles. Fig. 5 shows the medoid profile from each of the k = 5 clusters shown in Fig. 4. The medoid is the profile with the smallest total distance to all other profiles in the group. The medoid profiles show representative snowpack characteristics for each cluster, such as thicker layers of precipitation particles in regions with new snow problems and surface hoar, facet, or depth hoar layers in regions with persistent weak layer problems.

Figure 5: Representative (medoid) snow profiles for each of the k = 5 clusters on March 2, 2024. The title colours match the corresponding regions in Fig. 4.

4.2 Seasonal trends

Figure 6 shows the temporal progression of clusters for five days following March 2, 2024 and illustrates how region boundaries changed with evolving conditions. The human-assessed forecast region boundaries often align with either a cluster boundary or an area of cluster uncertainty. The boundaries of human-assessed regions changed on March 4 near the brown cluster, and on March 5 near the brown, blue, and purple clusters. The clustering results changed the boundaries more frequently and suggest the need for more regions than the human forecasts in some areas (e.g., the northeast regions with green, orange, and brown clusters) or fewer regions in others (e.g., the northwest region with a blue cluster).

The most common arrangement of regions from the clustering method generally matched human-assessed regions, with only slight differences (Fig. 7). In the northwest, the human and model-derived regions were fully aligned, while in the southwest both methods split the area into four regions, with slight variations in where the boundaries exist. Model clustering tended to identify more regions than the humans in the northeast, and the regions had slightly different shapes in the southeast.

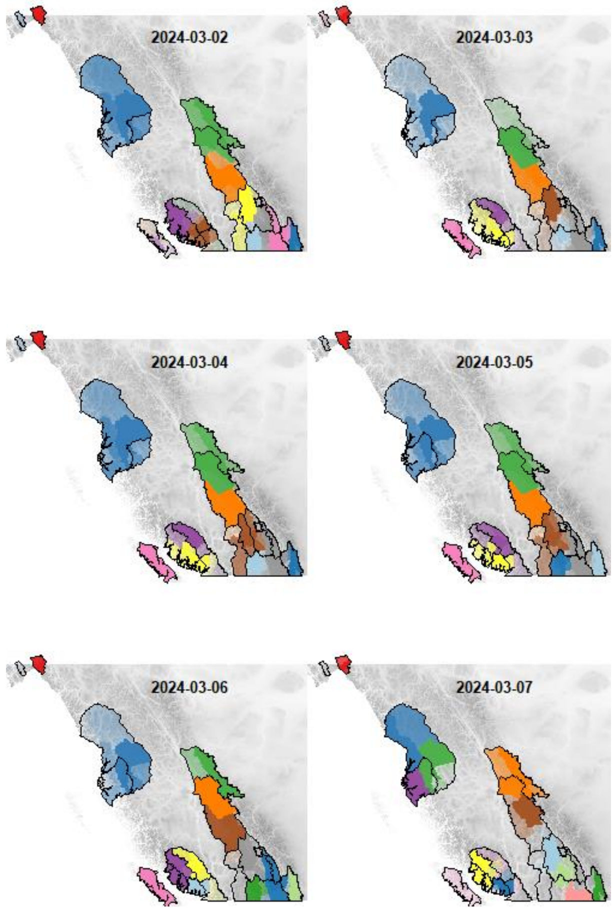

Figure 6: Temporal progression of clustering results from March 2 to March 7, 2024. The boundaries of the human-assessed regions are outlined in black.

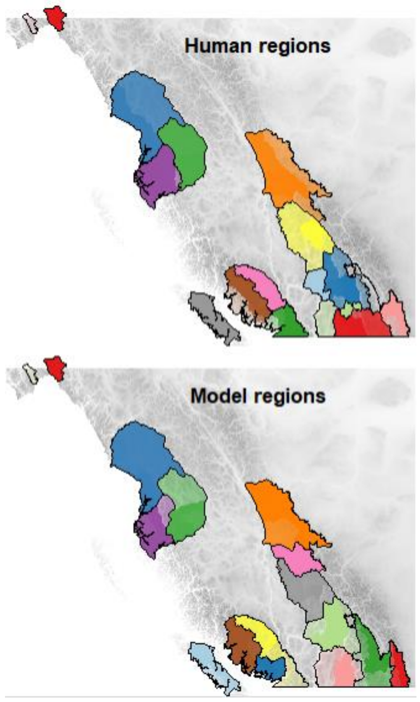

Figure 7: The 15 most common arrangements of subregions for the 2023-24 season according to the human-derived public forecasts and the model-derived clustering results.

5. DISCUSSION

The clustering method effectively captured major hazard patterns across western Canada during the 2023-24 winter. These results are consistent with our operational experience clustering snowpack model output at Avalanche Canada over the past three winters. Clustering provides simplified maps and a small number of representative snow profiles, giving forecasters a quick overview of potential hazard patterns. Forecasters at Avalanche Canada find cluster maps intuitive and useful for viewing model output and informing decisions about forecast region boundaries. The uncertainty from fuzzy clustering and the ability to adjust the number of regions are helpful features. However, due to challenges in validating snowpack model output, clustering is seen primarily as a data exploration tool rather than a fully automated solution for defining region boundaries.

Spatial constraints are important for producing coherent forecast regions, otherwise, distant regions with similar hazard characteristics could be grouped. Quantifying spatial relationships with the binary neighborhood approach was effective for subregion polygons, however, it could be interesting to quantify spatial patterns at smaller scales such as across aspects, elevations, or slopes.

Parameter tuning has been important when applying these methods over different domains. For example, different spatial weights were needed for this study area compared to the smaller study area in Horton et al. (2024). Similarly, the ideal fuzziness parameter in the clustering algorithm needs to be tuned to work with the distribution of values in the distance matrix.

The hazard metrics in this study were selected to correspond with four European avalanche problems. These problem types are more closely aligned with the physical processes resolved by snowpack models compared to the North American problems that are based on risk management factors. Quantifying snowpack differences with hazard metrics rather than comparing full stratigraphies resulted in dramatic improvements in computational efficiency and allowed clustering over larger domains. Comparisons with the stratigraphy-based approach of Horton et al. (2024) found similar region groupings for this domain and season (not shown). The differences between the hazard metric method and the stratigraphy-based method were within the same range as those resulting from subtle parameter tuning.

An additional constraint we propose for operational applications is a temporal distance that considers the previous day's regions, ensuring changes in region boundaries only happen in response to significant hazard variations. This approach guards against sensitivity to minor fluctuations in snowpack model output, as demonstrated by Horton et al. (2024).

6. CONCLUSIONS

Clustering simulated snow profiles is a promising method for presenting model output to forecasters. Our approach offers flexibility in considering diverse snowpack properties and spatial relationships. When combined with snow profile processing tools like averaging and weak layer detection, it can produce a concise summary of snowpack patterns. Developing tools that enable forecasters to explore model output in these ways has the potential to enhance the value of snowpack models for operational avalanche forecasting in a range of contexts.

Ongoing research into deriving hazard information from snowpack models (e.g., Mayer et al., 2022; Pérez-Guillén et al., 2022; Herla et al., 2024) is opening up opportunities to streamline the computation and interpretation of model outputs. This study serves as a proof of concept and has led to plans for refining the hazard metrics to better represent avalanche problems and integrate the method into Avalanche Canada's operational model.

ACKNOWLEDGEMENT

We thank the members and supporters of the SFU Avalanche Research program for their contributions to developing new snowpack model processing methods. Avalanche Canada forecast staff tested several iterations of clustering methods in operations over the past years, with the development team assisting in dataset preparation. This work was funded by NSERC, Mitacs, and Avalanche Canada.

REFERENCES

Bouchayer, C.: Synthesis of distributed snowpack simulations relevant for avalanche hazard forecasting, Master's thesis, University Grenoble Alpes, 2017.

Buhler, R, Horton, S., Schroers, B.: An analysis of Avalanche Canada's flexible forecast regions, in: Proceedings of the International Snow Science Workshop, Bend, USA, 9-13 October 2023, 158-165, 2023.

Horton, S., Haegeli, P., Klassen, K., Floyer, J., Helgeson, G..: Adopting snowpack models into an operational forecasting program: Successes, challenges, and future outlook, in: Proceedings of the International Snow Science Workshop, Bend, USA, 9-13 October 2023, 1544-1549, 2023.

Horton, S., Herla, F., and Haegeli, P.: Clustering simulated snow profiles to form avalanche forecast regions, EGUsphere [preprint], https://doi.org/10.5194/egusphere-2024-1609, 2024.

Herla, F., Horton, S., Mair, P., and Haegeli, P.: Snow profile alignment and similarity assessment for aggregating, clustering, and evaluating snowpack model output for avalanche forecasting, Geosci. Model Dev., 14, 239-258, https://doi.org/10.5194/gmd-14-239-2021, 2021.

Herla, F., Haegeli, P., and Mair, P.: A data exploration tool for averaging and accessing large data sets of snow stratigraphy profiles useful for avalanche forecasting, Cryosphere, 16, 3149-3162, https://doi.org/10.5194/tc-16-3149-2022, 2022.

Herla, F., Haegeli, P., Horton, S., and Mair, P.: A quantitative module of avalanche hazard - comparing forecaster assessments of storm and persistent slab avalanche problems with information derived from distributed snowpack simulations, EGUsphere [preprint], https://doi.org/10.5194/egusphere-2024-871, 2024.

Kaufman, L. and Rousseeuw, P. J.: Finding groups in data: an introduction to cluster analysis, John Wiley & Sons, Hoboken, New Jersey, USA, ISBN 0-471-73578-7, 2009.

Mayer, S., van Herwijnen, A., Techel, F., and Schweizer, J.: A random forest model to assess snow instability from simulated snow stratigraphy, Cryosphere, 16, 4593-4615, https://doi.org/10.5194/tc-16-4593-2022, 2022.

Morin, S., Horton, S., Techel, F., Bavay, M., Coléou, C., Fierz, C., Gobiet, A., Hagenmuller, P., Lafaysse, M., Ližar, M., Mitterer, C., Monti, F., Müller, K., Olefs, M., Snook, J., van Herwijnen, A., and Vionnet, V.: Application of physical snowpack models in support of operational avalanche hazard forecasting: A status report on current implementations and prospects for the future, Cold Reg. Sci. Technol., 170, 102910, https://doi.org/10.1016/j.coldregions.2019.102910, 2020.

Pérez-Guillén, C., Techel, F., Hendrick, M., Volpi, M., van Herwijnen, A., Olevski, T., Obozinski, G., Pérez-Cruz, F., and Schweizer, J.: Data-driven automated predictions of the avalanche danger level for dry-snow conditions in Switzerland, Nat. Hazards Earth Syst. Sci., 22, 2031—2056, https://doi.org/10.5194/nhess-22-2031-2022, 2022.

Shandro, B. and Haegeli, P.: Characterizing the nature and variability of avalanche hazard in western Canada, Nat. Hazards Earth Syst. Sci., 18, 1141–1158, https://doi.org/10.5194/nhess-18-1141-2018, 2018.