IMPLEMENTATION OF SNOWPACK AND MACHINE LEARNING MODELS IN AN OPERATIONAL FORECASTING WORKFLOW

Chad Brackelsberg[1*], Paige Pagnucco[1], Travis J. Morrison[1]

1 Utah Avalanche Center, Utah, USA

ABSTRACT: Snowpack modeling and machine-learning–based decision-support tools offer increasing potential for operational avalanche forecasting, yet their complexity can limit practical adoption in time-constrained forecast environments. The Utah Avalanche Center (UAC) has integrated SNOWPACK, instability metrics, and danger-level models into its daily avalanche forecasting workflow using an operations-first development framework focused on usability and forecaster needs.

The primary objective of this effort is to incorporate advanced model output into routine forecasting without overly increasing cognitive load or disrupting established workflows. System design was guided by four operational goals: improving forecasting efficiency, supporting forecast accuracy, providing clear and interpretable visualizations, and delivering information not readily available from field observations alone.

To guide implementation, we identified more than 35 questions commonly considered by avalanche forecasters during daily forecast development, including assessment of snowpack structure, tracking temporal changes in instability, and identifying spatial trends across elevation and aspect. Model outputs and visualizations were explicitly designed to address these questions, allowing forecasters to quickly contextualize model information within their existing decision-making process. This question-driven approach supported consistent interpretation and facilitated routine use during daily operations.

At the UAC, SNOWPACK and machine-learning–derived models are not intended to replace field observations or professional judgment. Instead, they function as an additional decision-support layer to help forecasters synthesize complex information, recognize emerging patterns, and prioritize limited time and field resources. Particular emphasis has been placed on transparency, consistency, and building forecasters' trust through iterative refinement and evaluation in operational settings.

KEYWORDS: snowpack, machine learning, operational forecasting, workflow

1. INTRODUCTION

1.1 Snowpack models as operational forecasting tools

Development of the SNOWPACK model by the Swiss Federal Institute for Snow and Avalanche Research began in the 1990s (Lehning et al., 1999). Its physical and numerical foundations were comprehensively documented in a three-part series in *Cold Regions Science and Technology* in 2002 (Bartelt and Lehning, 2002; Lehning et al., 2002a, b).

Morin et al. (2018) examined the operational use of snowpack models and identified challenges associated with implementing research-oriented tools in operational forecasting. Despite these challenges, snowpack models can provide information that cannot practically be obtained through field observations or may be missed by field observers. Snowpack models can provide continuous temporal information at locations where direct observations are limited, helping forecasters identify processes or changes that may otherwise go undetected until subsequent observations become available. Models can also diagnose conditions conducive to rapidly evolving near-surface processes, including surface-hoar growth and faceting through diurnal or radiation recrystallization, potentially alerting forecasters to emerging weak layers before they are documented by field observations (Birkeland et al., 1998; Adams et al., 2009; Horton et al., 2014).

More recently, machine-learning approaches have expanded model-derived information to include estimates of snow instability (Mayer et al., 2022) and avalanche danger (Pérez-Gullién et al., 2022), further increasing both the potential value and volume of information available to forecasters.

* *Corresponding author address:*
Chad Brackelsberg, Utah Avalanche Center, Salt Lake City, UT 84109;
tel: +1-801-365-5522
email: chad@utahavalanchecenter.org

Since 2018, operational interest in snowpack modeling has increased substantially, with both research organizations and forecast centers investing in implementation. Yet, there is still often a mismatch between research-oriented implementations and the operational usability of model outputs.

In addition to this mismatch, these models were developed and evaluated primarily in Switzerland, where climate and snowpack differ substantially from those of the Intermountain West of the United States. The models must therefore be validated to characterize model performance, identify limitations, and establish confidence before broad operational deployment and adoption. Validation frameworks incorporating both manual and automated processes are being developed for model validation and verification (Horton et al., 2026; Talty et al., 2026).

1.2 The operational problem

New model development, increased model implementation, and more model data do not necessarily translate into greater operational value. Operational avalanche forecasters face the simultaneous challenges of synthesizing ever-larger volumes of information while still working with substantial spatial and temporal gaps in direct observations. Snowpack modeling can help address both challenges, but the output data is often technically complex, information dense, and designed around model variables rather than forecaster questions. The outputs are difficult to interpret quickly and separate from the established forecaster workflow (Morin et al., 2018).

In avalanche forecasting, time matters. Forecasters only have two to three hours to synthesize all the information necessary to publish a forecast. Model products that require substantial additional interpretation or data synthesis may therefore provide little operational benefit even when the underlying information is scientifically valuable.

1.3 Utah Avalanche Center (UAC) approach and project objectives

The Utah Avalanche Center (UAC) turned the approach around. Instead of asking, “What model output should we show forecasters?”, we asked, “What questions do forecasters need to answer?”

When the UAC began operational SNOWPACK simulations during the 2022–2023 season, we first identified what we wanted snowpack modeling to accomplish within the forecasting workflow. Four operational goals became the guardrails for subsequent development:

- Improve forecasting accuracy
- Increase forecasting efficiency
- Provide clear and interpretable visualizations of data
- Provide insights and trends not readily available from field observations

Although the UAC serves as the initial operational environment, the system has been designed for broader deployment. The underlying architecture is intended to allow other avalanche forecasting centers and operational programs to configure model inputs, geographic domains, and data sources while retaining a common visualization and decision-support framework.

This paper describes the development process, visualization design principles, and operational lessons learned from integrating physics-based and machine-learning models into an active avalanche forecasting program. We discuss challenges encountered, strategies for successful adoption, and implications for other forecasting operations seeking to incorporate advanced snowpack modeling tools into daily practice. The central contribution is an “operations-first” framework in which forecaster questions, rather than available model variables, determine which model information is presented and how it is integrated into the forecasting workflow.

2. OPERATIONAL CONTEXT

2.1 UAC forecasting environment

The UAC provides a daily avalanche forecast from November through April for eight regions across Utah. Forecasts are published by 7:30 AM each day. Forecast regions range from the data-rich Salt Lake region to data-sparse regions like the Skyline. Data-rich regions have a broad array of weather stations, frequent and numerous daily observations from professionals and the public, and readily accessible terrain. The data-sparse regions have limited weather station information, more difficult to access terrain, and may not receive professional or public observations for several days. The data-rich Salt Lake region has access to 24 automated weather stations, while some rural regions have as few as three. The UAC receives, on average, 1,800 professional and public observations each season. On average, 48% of these are from the Salt Lake forecast region. This variability creates substantially different information environments for forecasters and provides an important operational use case for model-derived information, particularly in regions where direct observations are sparse.

As illustrated in Figure 1, model output enters the same forecaster data-synthesis process as weather, field observations, and other operational information; it does not bypass the forecaster or independently determine forecast products.

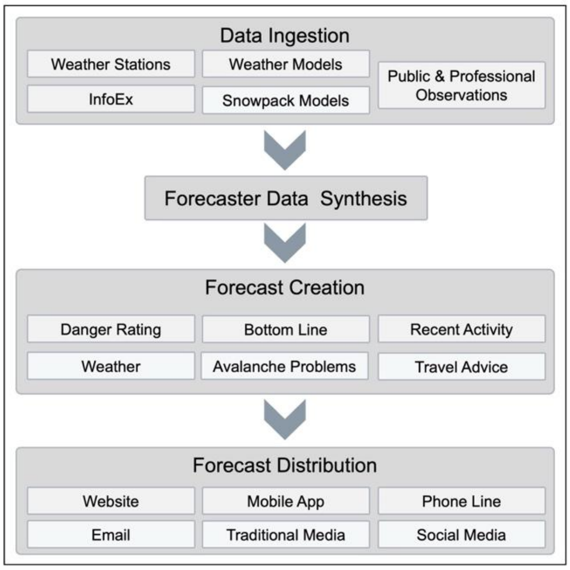

Figure 1: UAC forecaster workflow. Note that the forecaster remains the interpretation layer between the multitude of data streams and the final forecasting product.

2.2 *Models used in the UAC workflow*

The UAC workflow uses (i) SNOWPACK, (ii) instability, and (iii) danger level models. SNOWPACK (Lehning et al., 1999), which is the foundation model for the work, is a physics-based model that can be driven by either numerical weather model or weather station data. As discussed in Section 1.1, because this model was primarily developed and evaluated in Switzerland, we sought to minimize additional uncertainty associated with meteorological data. For this reason, we chose to run SNOWPACK with observed weather-station data for current-condition simulations and use HRRR weather model data for 12-, 24-, 36-, and 48-hour forecast intervals. This methodology reduces uncertainty associated with modeled meteorological runs. At each modeled weather-station location, simulations are produced for a flat field and for 35° slopes on north, east, south, and west aspects.

A snow instability model was implemented following Mayer et al. (2022), and an avalanche danger model followed the work of Pérez-Gullién et al. (2022). These are machine-learning models that use input data from both SNOWPACK and meteorological observations.

The instability model provides information about modeled snow instability and its association with weak layers across the simulated aspects, helping forecasters compare snowpack structure and instability between locations, terrain configurations, and time. The danger-level model provides an independent estimate of avalanche danger. At the UAC, this output is used as decision support rather than as an automated forecast; the forecaster considers it alongside observations, weather, snowpack information, and professional judgment when assigning the final danger rating.

The contrast between UAC's data-rich and data-sparse forecast regions also provides an opportunity to evaluate and tune model performance where observations are abundant, while assessing how those products may support forecasting where observations are limited.

These model outputs form the basis of the operational visualizations described in the following sections. The visualizations were designed around the goals identified in Section 1.3 and, more specifically, around the forecaster questions described in Section 3.1.

2.3 *The integration challenge*

As discussed in Section 1.2, the system needed to support forecaster decision-making without creating additional information overload. Models can be run at a very granular scale, providing dozens of data points per location, per aspect, per slope angle, per hour. A single weather-station location evaluated every 15 minutes across five terrain configurations produces 480 modeled states per day before individual variables or snowpack layers are considered. Extending the same temporal and terrain resolution across a 3-km grid covering the Salt Lake forecast region would produce more than 64,000 modeled states each day. Using our operations-first approach, the operational interface presents 12-hour summaries for four selected weather-station locations and five terrain configurations, resulting in 40 primary outputs for a forecaster to review each day. The operations-first approach does not require minimizing the amount of data generated by the models; it requires being selective about the information presented to the forecaster. In this implementation, substantially reducing the number of outputs requiring direct review limits information overload and makes model-derived information easier to incorporate into the existing forecast workflow.

The resulting design problem was therefore not how to expose the full range of available model output, but how to determine which information warranted a forecaster's limited attention. To make that determination, we began with questions forecasters routinely need to answer during forecast development.

3. OPERATIONS-FIRST DEVELOPMENT FRAMEWORK

3.1 *Start with forecaster questions*

The UAC is a public avalanche forecast center. Forecasters have limited time each morning to synthesize weather, snowpack, avalanche, and observational information; operational utility rather than research capability drove the design. The development process was designed around operational forecasting needs rather than available model outputs. Our goal was therefore to develop a tool that could function as a "second forecaster in the room": a decision-support system capable of quickly providing additional context when a forecaster was evaluating a specific question.

We compiled questions routinely raised during forecast development from forecaster discussions, weekly forecaster meetings, and interviews with individual forecasters. These questions defined the information needs that model products could help address. We initially identified more than 35 questions for which SNOWPACK or a machine-learning model output could potentially provide decision support.

Table 1: Example forecaster questions

Forecaster Question	***Information Needed***	***Model***
How has the instability of the most unstable layer changed over the last 24 hours and how is it forecasted to change over the next 24 hours?	Stratigraphy, Instability	SNOWPACK / Instability
How has the danger level changed over the last 24 hours and how is it forecasted to change over the next 24 hours?	Danger Level	Danger Level
Have conditions existed in the past 24 hours, or are they forecasted in the next 24 hours, that would support development of surface hoar or facets in the upper 30 cm of the snowpack?	Stratigraphy	SNOWPACK
How many weak layers are in the snowpack?	Stratigraphy	SNOWPACK
How unstable are the weak layers?	Stratigraphy, Instability	Instability
If there was an avalanche, how deep would it be?	Stratigraphy	SNOWPACK
If an avalanche stepped down to another weak layer, how deep would it be?	Stratigraphy, Instability	SNOWPACK / Instability
What is the modeled danger level by aspect?	Danger Level	Danger Level

Although individual questions required different model variables, the questions generally focused on three operational information needs: snowpack structure, instability and weak-layer evolution, and avalanche danger and its temporal or spatial trends.

3.2 *Translate questions into model products*

For each forecaster question, we first identified the information required to answer it and then mapped those information requirements to available model outputs. This created an explicit connection between an operational question and the underlying SNOWPACK, instability, or danger-level information.

This mapping was not necessarily one-to-one. Some forecaster questions required information from multiple model outputs or comparisons across time, locations, or terrain configurations. Conversely, individual model variables could contribute to multiple forecaster questions. The objective was therefore not to expose individual model variables, but to combine and summarize model information in a form appropriate to the operational question.

The mapped information was then translated into visualization prototypes. Products were designed to provide a rapid first-level interpretation while retaining access to more detailed model information when needed. This allowed forecasters to move from summary information to supporting detail without requiring detailed model output to be reviewed during every forecast cycle. Maintaining access to the underlying model information also provided sufficient transparency for forecasters to understand the basis for the displayed results and build trust in the model products.

3.3 Visualization principles

Visualization design proved to be one of the most challenging components of the development process. Individual forecasters differed in how they preferred to interpret information, including preferences for graphical versus tabular displays and the amount of detail initially presented. Horton et al. (2020) identified the need to develop clear design principles in conjunction with forecast staff. In some cases, this required presenting the same underlying information in multiple formats. Existing model visualizations were used when they met operational needs; when they did not, new visualizations were developed specifically for the forecast workflow.

Visualizations followed these design principles:

- Flow from high-level to more detailed
- Easy to interpret in a short amount of time
- Minimize unnecessary screen navigation
- Facilitate navigation between regions, weather stations, and aspects
- Prioritize trends over individual model values
- Make temporal changes obvious
- Avoid requiring forecasters to understand model internals
- Maintain sufficient transparency for forecasters to understand why a model produces a result

These principles required balancing simplicity with transparency. Forecasters should not need to understand the internal mechanics of each model to use its output operationally, but they must have sufficient access to the underlying information to evaluate whether a model result is plausible and appropriate for the current conditions and available observations.

3.4 Iterative forecaster development

For each forecaster question, two to three candidate visualizations were developed and presented to forecasters. Forecasters evaluated the alternatives for interpretability, relevance, and compatibility with their normal forecasting workflow. Based on this feedback, visualizations were revised and returned for additional evaluation. Products that met these criteria were then tested during operational forecasting. Operational use frequently identified additional issues, returning the product to the evaluation and revision cycle before final integration into the forecasting workflow (Figure 2).

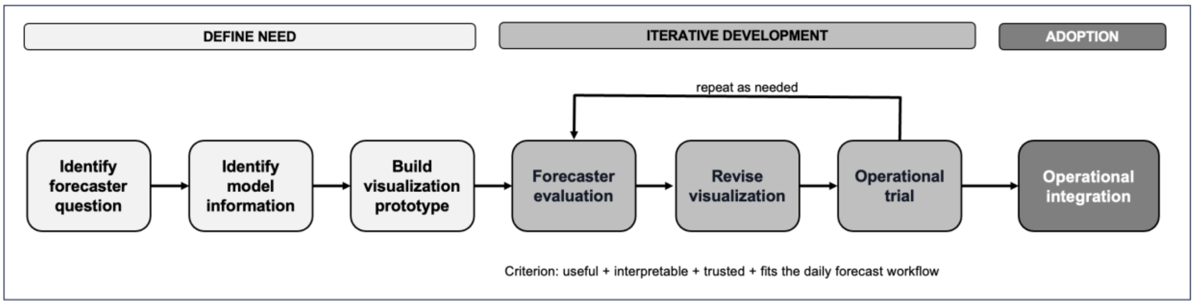

Figure 2: “Operations-first” iterative development process was used to translate forecaster questions into operational model visualizations. Forecaster evaluation and revision occurred repeatedly during operational testing prior to integration into the forecasting workflow.

This iterative process was crucial and led to key tool features, such as the summary page showing changes over the previous 24 hours and predicted changes over the next 24 hours, as well as the implementation of layer tracking. Layer tracking provides an example of how forecaster feedback changed not only product design but also how model information was organized. Weak layers were assigned names corresponding to their formation dates because the forecasters already used those names during operational communication. Matching terminology in the model interface to established forecaster terminology reduced the effort required to relate modeled layers to the conceptual snowpack model used by the forecasters.

Initial danger-level visualizations displayed the modeled danger level for each aspect as a single color corresponding to the North American Public Avalanche Danger Scale (Statham et al., 2010). The danger-level model uses a random forest consisting of 1,000 decision trees, with each tree producing an integer danger-level prediction, resulting in the probability for each danger level. For example, if 510 trees compute danger level 2, this is interpreted as a 51% likelihood that the day is moderate danger. Initially, we presented the danger level with the highest probability for display. This representation led to confusion and erosion of forecaster trust because small numerical differences could produce a disagreement over danger level between the model and the forecaster. Simple biases were magnified rather than highlighting the model's uncertainty. For example, the highest probable danger level from the model on a particular day could be danger level 2 with 51% likelihood, with danger level 3 having 49% likelihood. A forecaster assessing Considerable danger could therefore interpret the first result as the model underestimating danger, even though only 2% separated the two modeled values. Through iterative evaluation with the forecasters, the single-category display was replaced with a box and whisker-style graph that communicates the distribution of modeled danger levels across the decision trees rather than reducing the model output to a single categorical value. This experience demonstrated that the way model output is summarized and visualized can influence forecaster interpretation and trust independently of changes in the underlying model.

3.5 *Designed for different operational environments*

Although the tool was initially developed within the UAC forecasting environment, transferability to other avalanche operations was a design requirement. Potential users include public avalanche centers, ski areas, guiding operations, transportation departments, railways and utilities, and avalanche consultants. These organizations differ substantially in geographic extent, available observations, forecasting workflow, and operational objectives; consequently, the system was designed so that model locations, geographic domains, and data sources could be configured without changing the underlying visualization and decision-support framework.

During the 2025–2026 season, four additional avalanche operations used the system and provided operational feedback. This expanded the iterative-development process beyond the UAC and provided an initial test of whether products designed around forecaster questions and workflows were transferable to other operational settings.

Initial deployment requires access to meteorological input data of sufficient quality and completeness to drive the SNOWPACK simulations, together with a mechanism for automated data ingestion.

4. IMPLEMENTATION IN DAILY FORECASTING

4.1 *Daily workflow*

Model products are used at three points in the operational workflow: evening forecast preparation, morning forecast development, and targeted investigation during the day. During evening preparation, forecasters compare observations from the day with modeled snowpack conditions and evaluate expected changes over the next 12–24 hours. During morning forecast development, they review weak-layer evolution, newly developed layers, and modeled instability alongside observations and weather information when assessing avalanche problems and danger. During the day, more detailed model information can be used to investigate specific snowpack questions, such as whether an existing weak layer is strengthening or whether conditions favor faceting. Model outputs therefore function as an additional source of evidence within the existing forecasting process rather than as an independent forecast.

4.2 *Example forecasting case*

One of the goals of the project is to provide insights and trends not readily available from field observations. During each of the past four seasons, model products have provided the forecast team with insight into snowpack changes that later became associated with avalanche activity. One example of this was the spring of 2025. As depicted in Figure 3 (left), a melt-freeze crust formed on 25 February at approximately 265 cm above the ground. Beginning on 26 February, temperatures dropped for six days, during which facets developed above and below the crust at depths of approximately 245-255 cm above the ground (Figure 3, middle). On 3 March, we received 30 cm of snow with additional snow falling on 4 and 5 March. On 5 March, the models identified this facet layer as the most unstable layer in the snowpack at 245 cm above the ground (Figure 3, right). On the same day, a forecaster contacted the modeling team to question the validity and operational significance of the layer. At that time, the layer was not considered operationally significant. On 6 March, the UAC received reports of seven avalanches in the Salt Lake Region that failed on this layer. More snow fell on 7 and 12 March with additional avalanches failing on this layer reported from 7-8 March. These were soft slab avalanches 30-60 cm deep.

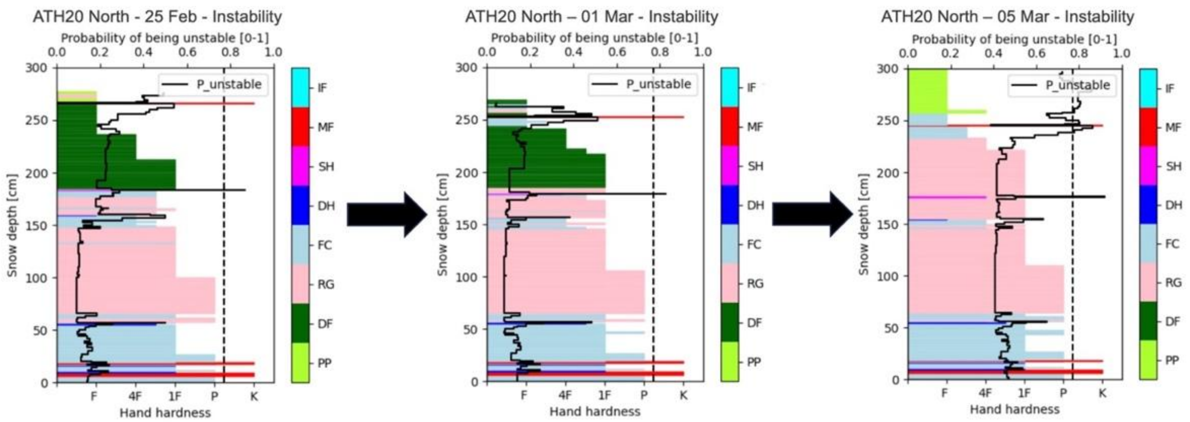

Figure 3: Modeled evolution of the 25 February melt-freeze crust and associated faceted layer during the March 2025 forecasting case. The sequence shows the formation of facets above and below the crust at approximately 250 cm from the ground surface. Moving from left to right subsequent loading highlights the faceted layer becoming the most unstable modeled layer.

This case demonstrates how modeled snowpack evolution can provide information that is not yet apparent from field observations. The model tracked faceting associated with the 25 February crust and identified the resulting layer as the most unstable modeled layer on 5 March, when its operational significance was still questioned by the forecast team. Avalanche observations beginning 6 March subsequently confirmed that the layer had become active. Similar cases during other operational seasons have also demonstrated the value of model products for identifying emerging weak layers before their significance is established through direct observations.

5. OPERATIONAL EXPERIENCE AND DISCUSSION

5.1 Designing around decisions rather than models

The operational value of a model product depends not only on the information produced by the model but on whether that information addresses a decision a forecaster is actively trying to make. Organizing the interface around forecaster questions reduced the amount of model information requiring direct review and allows model output to enter the same data-synthesis process as weather and field observations. This experience suggests that operational implementation should begin with information needs and decisions rather than with available model variables.

Because only a small fraction of model outputs warrant direct attention, focusing on operational value supports forecaster adoption and integration into their workflow.

5.2 Human-model interaction and trust

Operational trust developed through repeated use, transparency, and the ability to evaluate model output against other evidence. The danger-level visualization example demonstrated that trust can be eroded by how model uncertainty is represented and by not providing sufficient transparency into how the visualization was generated, even when the underlying model is unchanged. The March 2025 case also illustrates that disagreement between a model and a forecaster can itself be useful when it prompts closer examination of an emerging condition. The objective is therefore not model–forecaster agreement, but an interaction in which model output can be understood, questioned, and appropriately weighted within the broader forecasting assessment

5.3 Transferability and scalability to other operations

Initial use by four additional avalanche operations suggests that the operations-first framework is transferable beyond the UAC, although individual implementations must account for differences in available weather station and numerical weather-model data, climate conditions, geographic scale, forecast workflow, and operational objectives. The most transferable component may therefore be not a specific visualization or configuration, but the process of identifying forecaster questions, mapping them to model information, and iteratively evaluating products during operational use. Deployment also remains dependent on meteorological inputs of sufficient quality to support SNOWPACK simulations.

6. CONCLUSIONS

Physics-based and machine-learning snowpack models can provide valuable decision support for operational avalanche forecasting, but increasing model capability does not inherently increase operational value. The UAC experience demonstrates that model products can be integrated into a time-constrained forecasting workflow by beginning with forecaster questions, selectively presenting information relevant to those questions, and iteratively refining products through operational use. Model outputs have provided temporal information about snowpack evolution that may not be available from direct observations, while experience with danger-level visualization and emerging weak layers demonstrates the importance of transparency and appropriate human-model interaction. Models remain one source of evidence within a broader forecasting assessment rather than replacements for observations or professional judgment. As these tools are deployed to additional operations, the operations-first framework provides a transferable approach for adapting increasingly complex model information to different forecasting environments.

ACKNOWLEDGEMENTS

We thank the UAC forecasters, Alta Ski Area, Deer Valley Resort, Park City Powder Cats, Utah Department of Transportation, and Avalanche Risk Consultants for supporting this project through weather-station data and operational feedback throughout the development process. This project was made possible through funding from the Utah Outdoor Adventure Commission and other interested donors.

REFERENCES

Adams, E. E., McKittrick, L., Slaughter, A., Staron, P., Shertzer, R., Miller, D., Leonard, T., McCabe, D., Henninger, I., Catharine, D., Cooperstein, M., and Laveck, K. 2009. Modeling variation of surface hoar and radiation recrystallization across a slope. Proceedings of the International Snow Science Workshop, Davos, Switzerland, 97–101.

Birkeland, K. W., Johnson, R. F., and Schmidt, D. S. 1998. Near-surface faceted crystals formed by diurnal recrystallization: A case study of weak-layer formation in the mountain snowpack and its contribution to snow avalanches. Arctic and Alpine Research, 30(2), 200–204.

Horton, S., Bellaire, S., and Jamieson, B. 2014. Modeling the formation of surface hoar layers and tracking post-burial changes for avalanche forecasting. Cold Regions Science and Technology, 97, 81–89.

Horton, S.,Herla, F., van Herwijnen, A., Meloche, F., Morrison, T., Quéno, L., Vionnet, V., Building trust in snowpack models: evaluation guidelines for researchers and forecasters. In: Proceedings of the International Snow Science Workshop, Whistler, BC, October 2024.

Horton, S., Nowak, S., and Haegeli, P.: Enhancing the operational value of snowpack models with visualization design principles, Nat. Hazards Earth Syst. Sci., 20, 1557–1572, https://doi.org/10.5194/nhess-20-1557-2020, 2020.

Lehning, M., Bartelt, P., Brown, B., Russi, T., Stöeckli, U., Zimmerli, M., 1999. SNOWPACK model calculations for avalanche warning based upon a new network of weather and snow stations. Cold Reg. Sci. Technol. 30, 145 – 157.

Lehning, M., Bartelt, P., 2002. A physical SNOWPACK model for the Swiss avalanche warning Part I: Numerical Model. Cold Reg. Sci. Technol. 35, 123-145.

Lehning, M., Bartelt, P., Brown, B., Fierz, C., Satyawali, P., 2002a. A physical SNOWPACK model for the Swiss avalanche warning: Part II. Snow Microstructure. Cold Reg. Sci. Technol. 35,147 – 167.

Lehning, M., Bartelt, P., Brown, B., Fierz, C., 2002b. A physical SNOWPACK model for the Swiss avalanche warning: Part III. Meteorological forcing, thin layer formation and evaluation. Cold Reg. Sci. Technol. 35, 169 – 184.

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

Pérez-Guillién, C., Techel, F., Hendrick, M., Volpi, M., van Herwijnen, A., Olevski, T., Obozinski, G., Pérez-Cruz, F., and Schweizer, J., 2022. Data-driven automated predictions of the avalanche danger level for dry-snow conditions in Switzerland, Natural Hazards Earth System Sciences, 22, 2031–2056.

Statham, G., Haegeli P., Birkeland K., Greene E., Israelson C., Tremper B., Stethem, C., McMahon B., White B., Kelly J., 2010. The North American public avalanche danger scale. In: Proceedings Snow Science Workshop, Squaw Valley, CA,17-22 October 2010, pp 117–123.

Talty, M., Morrison, T., Eskeland, J., 2026. An applied validation and verification framework for computer assisted public avalanche forecasting. In: Proceedings of the International Snow Science Workshop, Whistler, BC, October 2024.