# Answer RC1

François Doussot, Léo Viallon-Galinier, Nicolas Eckert, Pascal Hagenmuller

## 1 Specific comments

### Comment 1

> **Predictors across elevation bands**: The rationale for selecting predictors from multiple SAFRAN elevation bands is not fully clear. As most predictors relate to avalanche release, conditions at start-zone elevations should be more relevant than those near runout elevations. Taking certain variables taken from a single elevation (e.g., wind direction, persistent weak layers) and others from multiple elevations accentuates an inconsistency. Snowpack conditions at lower elevations are likely secondary influences (e.g., sufficient snow depth to cover ground roughness, snow available for entrainment, or weak layers for step-downs). Mixing these predictors with those from start zone elevations may weaken model performance and, as noted in the discussion, adds complexity due to elevation-dependent climate change effects. A more consistent approach that either directly targets start-zone elevations or more clearly justifying the inclusion of lower elevations would improve interpretability.

We thank the reviewer for this comment. We agree that the predictor selection across different elevations was insufficiently explained in the original manuscript and that some choices may appear inconsistent. In response, we revised the predictor set so that the variables directly related to avalanche release processes (e.g., wind speed) are now provided only at the representative release-zone elevation (2700 m). We also added a detailed justification for the variables retained across multiple elevations, explaining their potential influence on avalanche flow and runout, which is particularly relevant given that the dataset includes only avalanches that meet a predefined observation threshold. The corresponding explanation has been added to the Methods section:

*Most predictors were selected to characterize snowpack conditions in the release area and therefore target avalanche release processes, such as wind slab formation or persistent weak-layer instability. Accordingly, variables directly related to release conditions (e.g., wind speed and strength-to-stress ratio) were provided only at 2700 m, which approximately corresponds to the mean elevation of the avalanche release zones. In contrast, a subset of variables was extracted at 1800, 2100, 2400, and 2700 m because conditions along the avalanche path may influence avalanche flow and runout altitude. This distinction is particularly relevant because the dataset includes only avalanches that meet a predefined observation threshold. Although these factors are expected to play a secondary role compared with release-area conditions, they may affect whether a released avalanche is recorded in the dataset. For example, sufficient snow depth or available transportable snow at lower elevations may be required along the path for an avalanche to reach the threshold. Solid and liquid precipitation were also provided at all elevations because precipitation can affect both avalanche release and avalanche flow. Wet-snow-related variables were similarly provided at all elevations because snowpack wetting generally progresses from lower to higher elevations. Accounting for this elevation-dependent evolution allows a more complete characterization of wet-snow conditions throughout the avalanche path. We retained variables with a plausible physical link to avalanche activity rather than applying a strict pre-selection procedure.*

*Gradient boosting methods are generally robust to the presence of weakly informative predictors, allowing physically relevant variables to be included without substantially degrading predictive performance. This approach preserves the physical interpretability of the model while allowing the algorithm to identify the most informative predictors among the provided features.*

This new feature selection leads to minor changes in the evaluation section: while the annual and seasonal MAE are slightly higher compared to the last version of the manuscript, the "worst week" MAE is now significantly lower. These results are updated in the new version of the manuscript and do not require major modifications to the discussion.

**Comment 2**

> **Results by aspect**: It would be interesting to report model performance by aspect (e.g., mean absolute error for each aspect) to better understand the skill of the model chain. Since the model predicts avalanches by aspect, this also raises the question of whether past and future trends differ by aspect (for example is there a stronger decrease on south-facing slopes than north-facing slopes?). While a full additional analysis may be beyond scope of this paper, briefly reporting performance by aspect and discussing potential aspect-dependent trends would add value.

Because of the valley topology, the number of avalanche paths varies substantially among aspect sectors. Eastern and western sectors contain relatively few avalanche paths, whereas northern and southern sectors each include 36 paths. We therefore focused our exploratory analysis on the comparison between northern and southern aspects.

Despite having the same number of avalanche paths, the two sectors exhibit markedly different avalanche activity. On average, since 2006/2007, northern aspects have experienced approximately 57 avalanches per year, compared with 32 on southern aspects. This difference in baseline activity should be kept in mind when interpreting aspect-dependent results, as avalanche counts may be affected by threshold effects, especially when avalanche counts are low.

As suggested, model performance was evaluated separately for each aspect sector using the same metrics as those presented in the manuscript. Performance was lower than for the aspect-aggregated model. For example, the annual mean absolute error (MAE) reached 33% for northern aspects and 34% for southern aspects, compared with 23% for the aggregated model. These results indicate greater uncertainty in aspect-specific predictions and, therefore, reduced confidence in aspect-level trend estimates. In addition, the model exhibited contrasting biases in the predicted annual number of avalanches between aspects, with a positive bias of 12% for northern aspects and a negative bias of 29% for southern aspects. These results indicate greater uncertainty in aspect-specific predictions and suggest that model errors are not homogeneous across aspect sectors. Consequently, comparisons of temporal trends between aspects should be interpreted with caution.

The Bayesian framework allows direct comparison of temporal trends across aspects via the posterior distributions of trend slopes. The results are discussed here based on the S2M reanalysis (not for the ADAMONT simulation). The posterior median estimates suggest a stronger decrease in the annual number of avalanches on northern aspects (-4.4% per decade; 95% credible interval: [-9.3%, 0.5%]) than on southern aspects (-3.1% per decade; 95% credible interval: [-8.4%, 2.4%]). However, both posterior distributions largely overlap (shown in Fig. 1), reflecting substantial uncertainty in the estimated trends. The posterior probability that the trend is more negative on northern than on southern aspects is 60%, indicating only weak evidence for an aspect-dependent difference in trends. Moreover, this probability does not account for the differing biases and predictive performance of the aspect-specific models. Therefore, while the median estimates point toward a stronger decrease on northern aspects, the evidence remains weak and does not support a robust conclusion regarding aspect-dependent trends. Given the lower predictive performance of the aspect-specific models compared with the aggregated model, a dedicated analysis would be required to assess such differences more reliably.

While this additional analysis provides some insight into potential aspect-dependent differences, the associated uncertainties remain large and the results do not lead to conclusions that differ from those obtained with the aggregated analysis. To maintain the focus and conciseness of the manuscript, we therefore chose not to include this exploratory analysis in the revised version.

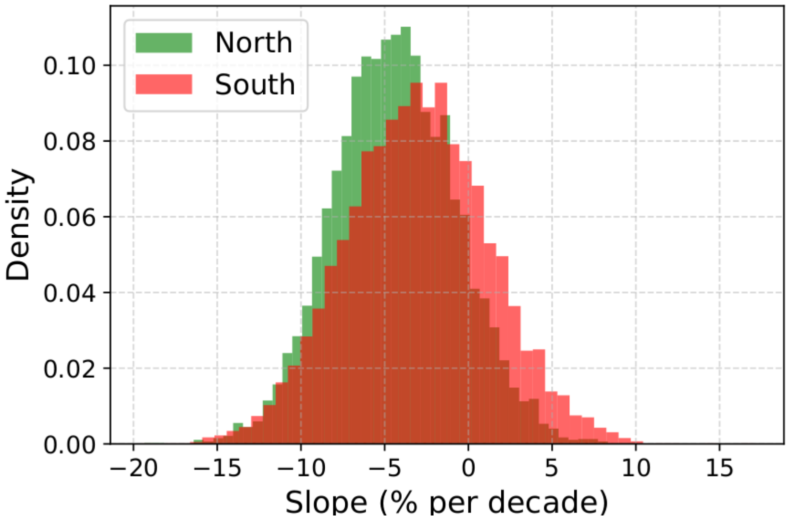


Figure 1: Computed slopes per aspect of the annual number of avalanches based on the S2M reanalysis

**Comment 3**

> **Underpredicting biggest cycles**: This is an interesting result that leads to questions about whether it is a limitation of the physical models, the selected predictors, or the extreme value statistics models. This could deserve a bit more direct attention in the discussion.

First, the XGBoost model is primarily optimized to reproduce the more frequent conditions that dominate the observations. Such machine-learning approaches are therefore known to regress extreme predictions toward more common values, leading to an underestimation of the rare events, here the large avalanche cycles.

In addition, the proposed loss function implicitly accounts for uncertainty in avalanche release dates by distributing the observed avalanche counts among all potential release days. This generally improves model performance, as evidenced by the lower MAEs. A consequence of this approach is that intense avalanche cycles are partially spread over several days, which reduces the amplitude of daily peaks and tends to underestimate the highest daily avalanche counts.

Moreover, as discussed in Sec. 4.2, extreme precipitation is often due to eastern returns, which are heterogeneous over the massif (and are maximum in the considered municipalities). They are smoothed by SAFRAN assumptions (massif-scale homogeneity), thereby reducing the intensity of the predictor signals. However, this is also the case in the training dataset, making it difficult to evaluate whether this limitation limits the statistical model's ability to reproduce the largest avalanche cycles.

These additional explanations have been added to the new version of the manuscript.

### Comment 4

> Length and repetition: Some sections are longer than necessary and occasionally repeat similar points across different sections the manuscript. Certain details could be reduced without affecting the core message, results, or interpretation. In my opinion streamlining the text would improve clarity and strengthen the paper overall, but the current level of detail is also acceptable if the authors prefer it as is.

We thank the reviewer for this general comment on manuscript structure and clarity. However, given the paper's structure and the need to maintain coherence among the methodological description, results, and discussion, it is difficult to substantially reduce certain sections without compromising the clarity of the overall narrative. In several cases, the level of detail was intentionally maintained to ensure that the methodological choices and their implications are clearly justified, particularly for readers less familiar with the modelling framework.

We also note that the comment does not refer to specific sections, which makes it challenging to identify precisely where reductions would be most appropriate. Some minor changes have been made to the new version of the manuscript to simplify and clarify the text.

## 2 Technical comments

### Technical Comment T1

> Line 70-73: Sentence is confusing and should be clarified.

This sentence has been changed by : *"At the scale of entire mountain ranges, biases associated with individual observers may be negligible because large-scale trends may still emerge despite local observational artefacts (Eckert et al., 2010d). However, such biases can become critical when the observation record is based on only a few observers."*

### Technical Comment T2

> Line 90: What metric did Castabrunet use for future avalanche activity?

It has been clarified in this version of the article: *"they used a Composite Index of avalanche activity derived from the EPA database and from the MEPRA hazard index (Durand et. al,1999). This synthetic indicator, which cannot be associated to a direct observable avalanche activity index, was computed at seasonal and annual time scales over large regions of the French Alps."*

### Technical Comment T3

Line 133: Wind can also be a primary driver of snowpack variability by aspect, which I assume is also true for this study area.

> Exactly, this remark has been taken into consideration in the new version of the article : *"and sometimes on wind direction"*.

**Technical Comment T4**

Line 167: “Strong” activity is vague. Please specify (e.g., higher avalanche counts).

> This has been fixed in the revised text, using *"higher avalanche counts"*.

**Technical Comment T5**

Fig 2 and 3. Captions should clarify that these refer to the Haute-Maurienne subset, as “EPA dataset” could be interpreted as covering all of France.

> This has been fixed in the revised text.

**Technical Comment T6**

Line 188: Elevation bands are not described as ranges. Are the bands centered around the reported values?

> Thank you for pointing this out. The term "elevation bands" was potentially misleading, as SAFRAN provides meteorological and snowpack conditions at discrete reference elevations spaced every 300 m. To avoid ambiguity, we have replaced the term *"elevation bands"* with *"discrete elevation levels"*.

**Technical Comment T7**

Line 195: Briefly define/explain quantile mapping.

> This has been fixed in the revised text: *"which is a statistical bias-correction method that adjusts the distribution of climate model outputs to match that of a reference dataset, which is here the SAFRAN reanalysis."*

**Technical Comment T8**

Line 201: Add a sentence describing the final weather forcing datasets (e.g., temporal resolution and how it is structured across elevation and aspect).

> It has been added in the new version: *"Crocus outputs were produced at a 6-hour temporal resolution"). The spatial resolution (elevation, aspect and slope) is now explicitly provided ("here 1800, 2100, 2400 and 2700 m [...], N, E, S, W [aspect] and 40 degrees slope").*

**Technical Comment T9**

Line 206: Clarify whether “50 layers” refers to snow stratigraphy layers and whether “5 variables” refers to snowpack properties.

> The *“5 variables”* has been deleted to avoid misunderstanding in the new version.

**Technical Comment T10**

Line 226: It would be worth noting in the discussion the limitations of using a single strength–stress ratio to represent persistent weak layer avalanches. These types of avalanche are among the most complex to model and forecast, and a simplified metric is likely to miss important processes. A brief acknowledgment of this limitation would strengthen the interpretation of results.

> We agree with this remark, it has been added in the discussion (lines 472-479): *The choice of predictors could also be further refined. For example, we only used the strength-to-stress ratio to characterize persistent weak-layer avalanche problems. It provides a very simplified description of snowpack stability. More physically based metrics, such as the critical crack length (Gaume et. al, 2017), could potentially offer a more realistic representation of avalanche release processes and improve predictive skill.*

**Technical Comment T11**

Sect 2.3.2: It appears predictors are computed daily at 18:00. This should be explicitly stated.

> It has been explicitly added in this section.

**Technical Comment T12**

Sect 2.4.2: This section does not clearly explain how the different metrics are combined into a loss function. While nmin and nmax are intuitive based on the observation windows, the role of $\Delta$T is unclear. Even with Fig. 4, its interpretation and purpose in normalizing residuals are difficult to follow. Clarifying how $\Delta$T functions within the loss formulation would improve readability.

> You are right, in the description, we did not clearly relate the residual to the loss function. It is now corrected : *"This distance is then distributed uniformly across all days within $\Delta T$ by dividing it by the duration of the interval (in days). The resulting quantity defines the daily residual $R$ used as the loss."*

**Technical Comment T13**

Line 276: Winter seasons are often labeled by the year they end, though conventions vary.

> It has been corrected in the new version of the paper, in all the figures and in the text.

**Technical Comment T14**

Line 405: Sentence is unclear and should be revised.

This sentence has been clarified to define what is a zero-inflated count distribution : *"Indeed, the observed daily avalanche counts follow a zero-inflated count distribution, characterized by integer-valued observations with an excess of zeros (Young et al., 2022)."*

**Technical Comment T15**

Line 419: Please be more specific about the uncertainty in older observations. If the concern is potential underreporting in earlier periods, this could be stated more directly rather than described as generic uncertainty.

It has been clarified in the new version : *"These results suggest that avalanche reporting was not consistent throughout the 1958–2023 period in the Haute-Maurienne valley, with possible under-reporting during the earlier decades"*

**Technical Comment T16**

Line 560: Consider starting a new paragraph when transitioning to future results.

This has been fixed in the revised text

## 3 Additional modifications

In the previous version of the manuscript, a single model was trained using the S2M reanalysis dataset over the 2006/2007–2022/2023 winter seasons and then applied to the entire 1958/1959–2022/2023 period to estimate daily avalanche activity. However, this approach introduces a methodological inconsistency because the training period is also included in the reconstruction period. As shown by the leave-one-year-out cross-validation, the model exhibits systematic biases. When the model is evaluated on the data used for training, these biases are artificially reduced, which may, in turn, affect the estimated long-term trend.

To avoid this issue, we revised the methodology. For the 1958/1959–2005/2006 period, avalanche activity is estimated using a model trained on the 2006/2007–2022/2023 dataset, as in the previous version of the manuscript. For the 2006/2007–2022/2023 period, we use the results of the leave-one-year-out method as done in Sec. 3.1, ensuring that each year is predicted by a model that was not trained on that year.

This revised approach provides a clearer separation between model training and application, thereby avoiding potential in-sample bias in the trend analysis. The correction leads to moderate changes in the results, with generally weaker estimated trends than those reported in the previous version, as shown in Table 1. The discussion is still consistent with these new computed trends.

| | Annual | DJF | MAM | Worst week |
|---|---|---|---|---|
| **Previous version** | -6.0% | -4.7% | -8.4% | -3.5% |
| **New version** | -4.7% | -1.5% | -8.4% | -1.7% |

Table 1: Changes in the computed trends from the reanalysis. DJF refers to winter months (December, January, February), and MAM to March, April and May.

# Answer RC2

François Doussot, Léo Viallon-Galinier, Nicolas Eckert, Pascal Hagenmuller

## 1 Specific comments

### Comment 1

> Fig 2 shows an increase in avalanche activity, including the period 2006-2022. Arguably climate change was already taking effect here, yet rather than decreasing the records increase. This increase is insufficiently explained in the methods, as well as the effect it can have on the results.

We agree that Fig. 2 shows an increase in observed avalanche activity over the 2006–2022 period. However, this period is relatively short for assessing climate trends. In mountain environments, avalanche activity exhibits strong inter-annual variability driven by fluctuations in snowfall, temperature, and snowpack conditions. Consequently, trends computed over a 17-year period may be strongly influenced by internal variability and are not necessarily representative of longer-term climate-driven changes, which are typically assessed over periods of at least 30 years.

The machine-learning model was trained on the observed avalanche record over this period under the assumption that the observation process remained sufficiently homogeneous. This assumption is supported by discussions with local avalanche observers, who did not report major changes in observation practices during the study period. While minor inconsistencies in the observational record cannot be completely excluded, no evidence was identified suggesting a systematic change that could explain the observed increase in avalanche activity.

Importantly, the objective of the model is not to reproduce the temporal trend present in the observations themselves, but rather to learn the statistical relationship between meteorological and snowpack predictors and the observed avalanche activity. The trained model is then applied to meteorological and snow conditions simulated by the reanalysis and climate model chain to reconstruct the expected avalanche activity associated with those conditions. In this framework, potential short-term fluctuations in the observed record primarily affect the calibration dataset but would not impose trends on the prediction.

Nevertheless, we acknowledge that any non-climatic changes in the observational record could influence the learned relationships and therefore represent a source of uncertainty. This limitation is discussed in the manuscript.

### Comment 2

> ln 195: Quantile mapping assumes distributions in historical data are preserved in future data. As such, they tend to impose past changes on future data, and are generally not well suited for the analysis of extremes. The authors would improve the paper by mentioning this in the methods, as well as the way the ADAMONT dataset aims to mitigate this. Now the reader has to wait for the discussion of this -in my opinion- important aspect.

We thank the reviewer for raising this point. Quantile mapping is known to have limitations regarding the preservation of quantile-specific climate change signals and the representation of extremes. Since this aspect is important for interpreting our results, we now introduce it in the Methods section, where the ADAMONT framework is presented, rather than postponing the discussion until later in the manuscript. We also further discuss the potential consequences of this limitation for our analyses in the Discussion section.

**Comment 3**

In the discussion of these shortcomings due to quantile mapping (sect 4.2), it could be elaborated upon how these shortcomings might affect the results (more than just that they affect the results).

We agree that it is important to discuss not only the existence of this limitation, but also its potential implications for our results.

As highlighted by Cannon et al. (2015), standard quantile mapping (QM) may modify the climate change signal simulated by the driving GCM, particularly for upper quantiles and extreme events. In their analysis of daily precipitation in Canada, QM was shown to substantially amplify projected changes in extreme precipitation (+500%) relative to the raw GCM output (+120%). More generally, QM does not guarantee preservation of quantile-specific future trends and may therefore alter projected changes in extremes. In this example, extreme precipitations are amplified by the QM method, but it could also be dampened depending on the tail distributions (of the simulated and reference dataset).

In our case, we are unable to quantify the influence of QM on our results, nor even determine the direction of any potential bias. Such an assessment would require a dedicated sensitivity analysis, for example, by comparing multiple bias-correction methods (e.g., QM, QDM, ADAMONT) and/or by analysing the raw climate model outputs. This is beyond the scope of the present study. Nevertheless, we agree that this represents an important source of uncertainty and have strengthened the discussion accordingly. However, we compared the results obtained using the reanalysis (so without QM) and the historical ADAMONT simulations on the same perid. The results in the climate trends are very similar (-4.7% per decade for the reanalysis, -5.9% with ADAMONT in the annual number of avalanches), which shows that the ADAMONT method does not lead to inconsistent climate trends.

We also note that the ADAMONT framework includes a specific treatment for values above the 99.5th percentile. Rather than applying a standard quantile-mapping correction beyond this threshold, a constant adjustment based on the 99.5th percentile is used to allow values outside the historical range (Verfaillie et al., 2017). This approach aims to reduce some of the limitations of classical QM when dealing with unprecedented extremes, although it does not completely remove the uncertainty associated with bias correction of future extremes.

Finally, Cannon et al. (2015) showed that quantile delta mapping (QDM) better preserves the relative changes simulated by the raw GCM. Future work could therefore investigate the sensitivity of our results to the choice of bias-correction and downscaling methodology, here directly imposed by the ADAMONT approach.

**Comment 4**

fig 5: from the figure it appears if the model underestimates the extremes in 2017 on all metrics. This should be related to the shortcomings due to quantile mapping mentioned above.

We agree that the model underestimates the extreme avalanche activity observed during the 2017–2018 period (Fig. 5), as discussed in the manuscript. In the answer to RC1, we discuss why the model tends to underestimate this type of events. More specifically, to answer this comment, we would like to clarify that Fig. 5 corresponds to simulations driven by the S2M reanalysis, and therefore does not involve any quantile mapping correction. The underestimation of these extreme events is thus not related to the quantile-mapping method.

This discrepancy is most likely explained by the rarity of such extreme avalanche conditions: the 2017/2018 extreme avalanche event has never been observed in the training dataset. The machine-learning model learns the relationship between meteorological/snowpack predictors and observed avalanche activity from historical data; however, very high daily avalanche counts such as those observed in January 2018 are poorly represented (or have never been seen) in the training dataset. As a consequence, the model tends to regress toward more frequently observed conditions and cannot fully reproduce extreme regimes that lie outside or at the edge of the training distribution.

This limitation is inherent in the approach used here and reflects the assumption that the training dataset provides a sufficiently representative sample of avalanche activity, including extreme events. When this assumption is not fully satisfied, extreme values are typically underestimated, particularly at the daily scale. To predict such events, the model would likely benefit from a larger training dataset, for example by extending the training period further back in time. However, as discussed in the manuscript, incorporating older observations introduces additional challenges related to the increased heterogeneity of the observational record, which may negatively affect the consistency and quality of the training dataset.

**Comment 5**

> ln 330; temporal linear trends are calculated (% per decade). Why the assumption that the changes should be linear in time, since warming or CO2 increase are not?

We thank the reviewer for this relevant comment. The assumption of a linear trend does not imply that the underlying climate forcing (e.g. temperature or CO2 concentration) evolves linearly in time. Rather, it provides a first-order approximation of the average rate of change in avalanche activity over the study period.

Given the relatively short time window considered in this study, a linear model offers a simple and robust way to summarize long-term tendencies while limiting the influence of interannual variability. This approach is widely used in climatological and cryospheric studies to express changes as rates per decade, which facilitates comparison with previous work.

We acknowledge that climate drivers and associated physical processes may evolve non-linearly. However, capturing such non-linear responses at a local scale, where the inter-annual variability is very high, would require either longer time series or explicitly time-varying models, which are beyond the scope of the present analysis. Here, the linear trend is intended as a descriptive metric rather than a mechanistic assumption about the system.

Future work could also investigate more flexible formulations, such as piecewise linear trends, to account for potential changes in the rate of evolution over the study period. This extension is beyond the scope of the present analysis but represents a promising direction for further research.

**Comment 6**

> fig 6: Not clear in black and white.( I don't know if that is still a requirement in this day and age, but:) This could be improved by different linestyles for observed and simulated trendlines (e.g. dashed and dot-dashed).

This has been corrected in the new version of the manuscript.

**Comment 7**

> In general, many studies have noted or predicted a decrease in avalanche activity due to climate change. However the question why is not always elaborated on sufficiently. The obvious explanation is that the winters are simply getting shorter, and as a result, less avalanches occur over a season. Section 4.3 (and possibly the introduction) could benefit from mentioning this point explicitly. IMHO the interesting question is whether the decrease in avalanche activity can be solely explained by the decrease in winter duration, or if there are other dynamics at play. I hate to be the reviewer suggesting that his own paper be cited, but I'm going to do it anyway, as I feel that in this case mentioning this work is relevant.

Thank you for this relevant comment. We agree that the role of the shortening of the avalanche season should be explicitly acknowledged. This point has now been added in the introduction. We analysed the seasonal distribution of the simulated avalanche activity (Fig. 7b), which shows a shortening of the avalanche season in average. It has now been explicitly added to Section 4.3, which mentions the work you suggested. However, we did not directly link the computed trend to the evolution of the physical drivers. While the reduction in winter duration is a key and physically intuitive explanation, since avalanche activity cannot occur in the absence of sufficient snow cover, other processes such as changes in snow stratigraphy may also contribute to snowpack stability. If we want to go further and link the results of the machine-learning model (number of avalanches) to the physical drivers, it would require a dedicated interpretable modelling framework (using shapley values for example), which is beyond the scope of the present machine-learning approach. In purpose, we therefore did not further interpret the physical drivers of the simulated trends.

# 2 Additional modifications

In the previous version of the manuscript, a single model was trained using the S2M reanalysis dataset over the 2006/2007–2022/2023 winter seasons and then applied to the entire 1958/1959–2022/2023 period to estimate daily avalanche activity. However, this approach introduces a methodological inconsistency because the training period is also included in the reconstruction period. As shown by the leave-one-year-out cross-validation, the model exhibits systematic biases. When the model is evaluated on the data used for training, these biases are artificially reduced, which may, in turn, affect the estimated long-term trend.

To avoid this issue, we revised the methodology. For the 1958/1959–2005/2006 period, avalanche activity is estimated using a model trained on the 2006/2007–2022/2023 dataset, as in the previous version of the manuscript. For the 2006/2007–2022/2023 period, we use the results of the leave-one-year-out method as done in Sec. 3.1, ensuring that each year is predicted by a model that was not trained on that year.

This revised approach provides a clearer separation between model training and application, thereby avoiding potential in-sample bias in the trend analysis. The correction leads to moderate

| | Annual | DJF | MAM | Worst week |
|---|---|---|---|---|
| **Previous version** | -6.0% | -4.7% | -8.4% | -3.5% |
| **New version** | -4.7% | -1.5% | -8.4% | -1.7% |

Table 1: Changes in the computed trends from the reanalysis. DJF refers to winter months (December, January, February), and MAM to March, April and May.

changes in the results, with generally weaker estimated trends than those reported in the previous version, as shown in Table 1. The discussion is still consistent with these new computed trends.