REVISION #2

The manuscript has improved substantially, with a larger number of snow-monitoring stations, the inclusion of the requested workflow diagram, the addition of a station map, and a more detailed explanation of the applied methodology, particularly in Section 2.3.1. The Results section has also been strengthened.

Answer: We really appreciated the comment of the referee. We checked all the suggestions made e we corrected as follow.

At line 138, could the authors provide a clearer justification for the 30-day gap-filling window?

Answer: Thank you for raising this point. We clarify that the 30-day window is not used for averaging or smoothing the snow cover signal. Instead, it serves exclusively as a search constraint to identify the closest valid observations before and after each missing value. For each pixel and missing time step t, snow cover was reconstructed using only the two nearest valid observations at times t1 and t2. Linear interpolation was then applied between these two points (Eq. 2), without using any additional observations within the window.

Moreover, we have completely rewritten Section 2.3.1 to improve its clarity. We streamlined the text by removing redundant information and thoroughly rephrasing the explanation of the workflow and the individual equations. We believe the proposed methodology is now presented much more clearly.

At line 250, Sentinel-2 is described as having a temporal resolution of 5–10 days. Please verify this value, as it is closer to 3–5 days.

Answer: While the nominal revisit time of the Sentinel-2 constellation is 5 combined days (10 days per single satellite), the effective temporal resolution over Europe is often higher, typically around 2/3 days, due to orbital convergence and swath overlap at mid-latitudes. We have revised the part of the manuscript to clarify this effect. Despite nominal mission temporal resolution being 5 days, it has an effective revisit frequency of approximately 2/3 days over Europe.

The authors are still encouraged to make the Google Earth Engine (GEE) script publicly available.

Answer: Yes. We have uploaded the code via GitHub at the following link: https://github.com/francesco-parizia/Papers/blob/main/MDSA.md

The manuscript has substantially improved, but several minor editorial, typographical, and technical issues should still be corrected before acceptance. These do not affect the scientific conclusions but should be addressed to ensure clarity, consistency, and technical accuracy.

Abstract -> “at-the-pixel-level” → “at the pixel level”.
line 20 -> “1000m a.s.l.” → “1000 m a.s.l.”
l. 38 --> . In (Cati, 1981) dataset --> In Cati (2018) dataset
l. 53 --> SAR Definition acronym.
l. 53 analisis -> Analysis

l. 53 new spectral index -> new spectral indices.
l.102 -> Why MOD16A1? - > Should be the MOD10A1 product, isn't it? Please correct it also in figure 2.
l.119 separate 500m^2 to 500 m^2
l. 257 --> synthesizes
l. 320 aquifers.

Answer: we checked and corrected all the indicated issues

Report 2

REV1

For the Introduction part that the authors mentioned: "Fugazza et al. (2021) as we cited as a key study in ……………..*"*

I like to point out, the total number of days with snow on the ground per year is not a new output, even if you extended the record to 2023, because we already have operational products such as Daily Fractional Snow Cover 2016–present (raster 20 m) (https://land.copernicus.eu/en/products/snow/fractional-snow-cover) or gap-filled MODIS snow data MOD10A1F / MYD10A1F and summing snow-covered pixels over one year is a straightforward outcome of those datasets:

Answer: The extraction of number of days with snow is not an aim of the paper, in fact we cited Comitato Glaciologico Italiano (1964) that published an historical map of snow cover 1959-1960. Days of snow (we called it "Snow Persistence") is just a way to validate our process. As explicated in the manuscript "our goal is not to track snow phenology but to quantify the accumulated loss of intensity of snow cover for a single mean snow event".

Moreover, regarding this part: "Compare with Landsat/Sentinel snow cover data."

You do not necessarily need to compare with the entire 2000–2023 period using higher-resolution sensors. You can compare with the existing daily snow product that I mentioned earlier for the period which the dataset is available.

Answer: We can certainly recognize the scientific need and curiosity for a comparison of MODIS derived results with the ones from Landsat and Sentinel products. But, in spite of this more than proper suggestion by reviewer, that we can certainly consider for future works, we have to kindly stress again that the goal of this work was to specifically investigate the potentialities of the MOD10A1 product for well declared reasons: the first is that, given the fast dynamics of snow cover (especially in the transitional phases of the year) the potentially daily acquisitions of MODIS is far to be satisfied by Sentinel or Landsat data. The second is that, a higher geometric resolution (like the one of Landsat and Sentinel 2) is not intended for regional/sub continental scale analysis, especially because of the asynchronicity of the acquisitions when a large area, mainly developing East-West like the Alps are, has to be investigated. Whereas Sentinel 2 or Landsat products were used, each sector of the Alps belonging to a different tile, would refer to a different date, introducing further approximation in the analysis that are difficult to take care about.

We have added text to the manuscript to clarify our rationale for using meteorological stations

For these reasons, we confirm our will of avoiding any further analysis in a direction, as indicated by the reviewer, that shifts our work from its native goals. Moreover, if the need is the validation of our approach, this was already achieved in the only way we assume as definitive and absolute, i.e. the comparison of our results with ground measures. These are daily and do not suffer from gaps nor interpolation; and, above all, they have been officially validated by the Regional agencies.

For this part the authors mentioned: "We acknowledge the existence of the gap-filled products, MOD10A1F and MYD10A1F, however, our primary computational environment was the Google Earth Engine (GEE) platform…" As I mentioned before, GEE is not a reference for data downloading, and it also has the option of uploading data. Therefore, the argument that the selection is constrained mainly by the GEE public catalogue does not fully address the scientific concern (you could process externally or upload if needed). Because variable like snow changes rapidly and varies strongly across space. The gap-filled products MOD10A1F / MYD10A1F already exist and are specifically designed to handle this issue.

Answer: We thank the reviewer for this observation. Our decision was driven, also, by specific scientific motivations. Recent literature has identified known issues with the standard MODIS filtered product (e.g., Hao et al., 2022 - https://doi.org/10.5194/hess-26-1937-2022) They showed as products derived directly from MOD10A1 have high accuracy (86%–97%) under clear conditions. MOD10A1F has similar accuracy in areas with short-lived clouds, but may overestimate snow in, regions with long-lasting clouds or, persistent cover. Therefore, to ensure greater scientific rigor, we chose to generate our own filtered product. Moreover, as reported (new part added in 2.3.1 to clarify the concept), we decided to follow this precise gap-filling worklow to avoid also oversmooth or generate oscillatory behaviours that many "complex" interpolation techniques could create, as mentioned by Hird and McDermid (2009) and by Jönsson and Eklundh (2004)

Moreover, the authors mentioned: "nT, which quantifies the change of a snow covered area in mean annual snow event over the study period," the "mean snow event" is mathematically defined but it may still sound non-physical. Also, elevation still drives snow processes; normalization may reduce baseline effects, but it does not remove physical controls.

For this part: Snow Cover Trends Quantification: while the authors clarify that trends are derived from a full linear regression, the normalization by the intercept remains sensitive to the choice of temporal origin and may be influenced by anomalous early years. The explanation does not fully address this concern. The authors should either justify the robustness of the intercept-based normalization or consider alternative normalization strategies or sensitivity analyses.

Answer: We merge these two questions to better explain the concept:

We agree that the term "mean snow event" may sound ambiguous and potentially non-physical, and we revised the sentence to clarify its statistical meaning. In this study, the normalized trend nT is defined as the ratio between the slope of the linear regression of MDSA over time and its intercept. The intercept represents the estimated mean MDSA at the reference time of the regression and is used just as a local normalization factor. As such, nT expresses a relative rate of change (dimensionless) with respect to the local mean snow covered area, rather than describing an individual physical snow event.

The reference year 2000 is not selected arbitrarily nor interpreted as a physically meaningful benchmark. It corresponds to the temporal origin of the regression model. Using a common reference allows trends to be expressed relative to a consistent baseline and ensures spatial comparability across areas with markedly different snow regimes (e.g., lowlands versus high-elevation regions). Any other reference year would lead to equivalent normalized trends, provided that the same centering is applied consistently.

We also fully agree that elevation remains a primary physical control on snow processes. Normalization does not aim to remove these control, but rather to reduce baseline differences in mean snow cover, so that relative changes can be compared across heterogeneous environments that characterize the test area *territories. Physical drivers such as elevation, climate, and topography continue to govern the snow dynamics reflected in the observed trends.*

For the part Discussing The Trends In High-Elevation Areas and the physical mechanisms driving snow cover changes, I still have this question: why do reductions in perennial snow lead to a net decrease in MDSA rather than cancelling out in the iSCA/SP ratio? This needs a clearer, metric-consistent explanation given the definition MDSA=iSCA/SP.

Answer: The absence of cancellation in the iSCA/SP ratio directly follows from the different definitions of snow persistence and snow covered area. As defined in Eq. (5), SP is obtained by summing a binary snow mask over time and therefore accounts only for the presence/absence of snow within a pixel, independently of its spatial extent. In contrast, iSCA (Eq. 6) integrates the actual snow-covered area over time and is thus sensitive to sub-pixel variations in snow extent. When perennial or near-perennial snow retreats in high-elevation areas, snow-covered area during long summer periods is reduced, leading to a substantial decrease in iSCA. However, SP may remain unchanged or decrease only slightly as long as some snow persists within the pixel (i.e., the binary mask remains equal to 1). As a result, iSCA decreases more rapidly than SP, and their ratio (MDSA; Eq. 7) exhibits a net decline rather than cancelling out.

Therefore, the observed reduction in MDSA reflects a real contraction of snow-covered area during periods of high persistence, fully consistent with the metric definition.

We added follow sentence in the paper: "Because snow persistence is derived from a binary mask, it is insensitive to sub-pixel reductions in snow-covered areas. Consequently, losses of perennial snow primarily affect iSCA and lead to declining MDSA values even when snow-covered days remain numerous."