We thank the reviewer for their thoughtful and constructive comments and suggestions, which has substantially improved the quality of the manuscript. We have addressed all the reviewer's concerns and revised the manuscript accordingly. Our point-by-point responses are in blue and the modifications to the manuscript are quoted in green.

Reviewer #1

In my first review, I had two main concerns about the paper. The first of these had to do with the superiority of the two-step process to reduce the number of SOMs over simply using a smaller number of SOMs to begin with. The authors have kindly provided this comparison in Fig. R3 of their response. I see two distinct differences between the two methods. As the author's point out, the two-step procedure in the top row has in the last panel a region of high gpm in the center-west portion of the domain which is not present in any of the lower set of panels for the straight 5-SOM method. However, the lower set of panels also has in panel 2 a pattern of low GPM in the southwest portion of the domain, which is not present in any of the panels in the top row. Which of the two is better? The author's claim that the top row is (subjectively) better, but that is not obvious to me. Is there a way to quantitatively measure this, through some metric of the greatest separation between the figure patterns? The authors also have decided not to include Fig. R3 in the paper, but to me it is an interesting result, worthy of inclusion and discussion. Without it, the reader is forced to simply accept the authors contention that the 2-step method is better, without being able to understand or appreciate the differences between the two.

We thank the reviewer for their detailed observations regarding the comparison between the two-step SOM/K-means method and directly applying a 5-SOM clustering approach. To address this, we provide a response from both physical and statistical perspectives:

Physically, the 10x10 SOM step (Figure 1 in the manuscript) generates a lower-dimensional representation of the full dataset, capturing finer-scale details of the large-scale meteorological patterns. For instance, the Fig R3 second panel's pattern (low geopotential height anomaly in the southwest) emerges clearly in the upper-right portion of the 10x10 SOM map. This pattern is subsequently aggregated into a land-ocean pattern during the second (K-means) clustering step. However, the high geopotential height anomaly over the central and western portion of the domain (representing the California high) observed in the 10x10 SOM prototypes is not captured when directly applying the 5-SOM clustering method. This demonstrates that the two-step approach preserves important physical features that the 5-SOM method misses, which are meteorologically significant for understanding regional wind variability.

Statistical, we calculated the Silhouette coefficient, a metric that measures the compactness of clusters and their separation. A higher silhouette coefficient indicates that clusters are well-defined and distinct. As shown in the figure below:

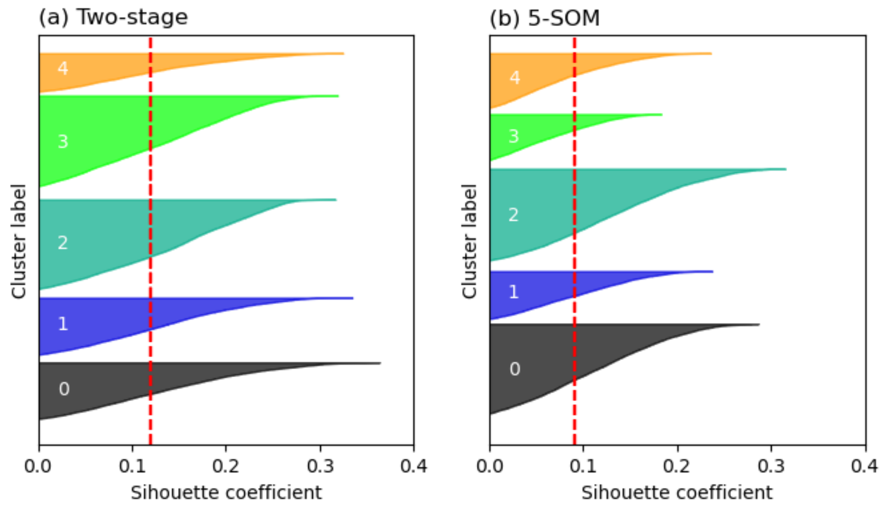

Figure Silhouette coefficients for each clusters in (a) two-state method and (b) 5-SOM clustering. The red dashed line indicates the mean Silhouette score.

The two-stage clustering approach produces clusters with higher average silhouette coefficients (0.12) compared to the direct 5-SOM clustering (0.09). The improved silhouette scores indicate that the two-stage approach results in more distinct and well-separated clusters, confirming its statistical superiority over the 5-SOM method.

Since the physical perspective of the two-stage clustering has already been documented in Section 2.4 of the manuscript, we have added the following text to **Lines 124–126** to provide a quantitative comparison:

We also compared the results of direct 5-SOM clustering with our two-stage method. The average silhouette coefficient is 0.12 for the two-stage method, compared to 0.09 for the direct 5-SOM clustering. The larger SS indicates that the resulting clusters are more well-defined and distinct.

The second main concern had to do with the "outlier" windspeeds present in the individual LSMP clusters shown in Fig. 4. The author's response to my concern is to note that in fact very different wind speeds can occur within a given cluster. That is in itself an important aspect of the method that should be emphasized more in the text. Also, it contradicts the statement in the abstract that "Distinct wind speed, wind direction, diurnal variation, and jet feature responses are observed for each LSMP and at both lidar locations." Fig. 4 shows that distinct wind speeds are not associated with each LMSP. Instead, within 8 of the 10 outlined LSMPs shown in the figure, the wind speed varies from ~4 to ~14 m/s. Also, multiple LSMPs appear by eye to have almost identical distributions of wind speeds. Again, I think that the paper would be stronger if these issues were discussed in the manuscript. No method is perfect, and the science community would benefit more if some of the limitations of the SOM/K-means method were acknowledged and addressed.

We appreciate the reviewer's insightful feedback regarding the variability of wind speeds within individual LSMP clusters. We agree that it is important to emphasize that while the LSMP-based classification method successfully links large-scale meteorological patterns to observed offshore winds, it does not guarantee a perfect one-to-one correspondence between an LSMP and a single "characteristic" wind speed. Indeed, different SOM prototypes within the same LSMP cluster may display a broad range of wind speeds, as noted by the reviewer.

This variability arises because offshore wind conditions result from the complex interplay between large-scale circulations and localized factors such as coastal topography, sea surface temperature gradients, or mesoscale phenomena.

Even within a single LSMP, subtle differences in weather features can lead to substantial variation in wind speed. Thus, while the LSMP clustering serves as a useful framework to understand and categorize overarching synoptic patterns, it does not eliminate the inherent variability present in real-world wind data.

In the revision, we have changed the sentence in the abstract mentioned by the reviewer to (**Lines 16-18**):

While each LSMP is associated with characteristic large-scale atmospheric conditions and corresponding differences in wind direction, diurnal variation, and jet features at the two lidar sites, substantial variability in wind speeds can still occur within each LSMP

In the Results section, we added the following paragraph (**Lines 238-246**):

It is important to note that the classification of LSMPs does not imply that each LSMP is associated with a narrowly defined wind speed. Rather, each LSMP reflects a dominant synoptic environment under which the prevailing direction and magnitude of offshore winds are modulated. Within any single LSMP, a range of mesoscale and local factors (e.g., frontal passages, varying thermal contrasts, topographic influences, and boundary layer structures) can lead to substantial variability in observed wind speeds. For instance, under post-ridge conditions, most SOM prototypes show a decrease in wind speed due to the induced southerly anomaly; however, a few prototypes exhibit unexpectedly high wind speeds when the southerly anomaly surpasses the prevailing northerly flow. This illustrates that while the LSMP framework provides a useful synoptic-scale context, it is primarily a classification tool rather than a deterministic method, and thus cannot eliminate the inherent complexity and spread in the local wind speed distributions.

And in the Conclusion section, we added (**Lines 318-320**):

Nonetheless, these results should not be interpreted to mean that each LSMP strictly enforces a single wind speed regime. A wide range of wind speed outcomes can occur, influenced by local and mesoscale processes.