Title: Arctic Climate Response to European Radiative Forcing: A Deep Learning Approach

Authors: Sina Mehrdad, Dörthe Handorf, Ines Höschel, Khalil Karami, Johannes Quaas, Sudhakar Dipu, and Christoph Jacobi

Summary: This study utilizes deep learning techniques to investigate the Arctic atmospheric circulation responses to European radiative forcing, which is both intriguing and significant for advancing our understanding of applying convolutional auto-encoder frameworks in Arctic climate change research. Overall, the manuscript is well-written, and the authors extensively discuss the consistent dynamical responses, although interpreting causality remains challenging. Below, I provide a few major and minor comments for the authors' consideration.

Thank you very much for your insightful review. We greatly appreciate your recognition of the research topic's intrigue and significance. We acknowledge your point regarding the challenges in interpreting causality within our study. Indeed, one of the key contributions of our work is to highlight that deriving causality in the context of regional forcing on Arctic climate is complex and multifaceted.

As you have pointed out our presentation of the results and the mechanisms derived from our analysis might not have been as clear as necessary. We will revise the manuscript to better highlight these aspects. Specifically, we will clarify our methodology and its ability and shortage to identify the mechanism through which Arctic responded to the forcing. We will highlight our findings more clearly in the updated version of the manuscript.

Major comments:

Baseline machine learning model comparison. This study presents an innovative application of convolutional autoencoders (AEs), which appears to be one of the first attempts in studying Arctic response to radiative forcing. However, I am still unclear about the motivation behind choosing convolutional AEs. For example, the authors mentioned that in Lines 538-539 that the FAE is a compelling method for generating a concise and informative representation, but did not elaborate on compare to what. Have the authors compared the performance of convolutional AEs with other machine learning or statistical methods? For example, similar clustering analyses could be conducted using self-organizing maps (SOMs), which are computationally less expensive than training convolutional AEs. SOMs have been used in studying atmospheric moisture transport in the Arctic or large-scale atmospheric circulations (e.g., Skific et al. 2009; Lee 2017). Additionally, principal component analysis (PCA) or empirical orthogonal function (EOF) analysis could be employed for clustering tests. Why not start with these simpler methods before diving into complex deep learning models? However, if the authors can demonstrate that convolutional AEs outperform SOMs or PCA/EOF, it would strengthen the justification for using convolutional AEs in this study. Perhaps the authors could consider quickly implementing these simpler methods and comparing the results with those obtained from convolutional AEs.

Thank you for these comments and the opportunity to clarify our motivation to choose our deep learning approach. Your insights have encouraged us to more thoroughly compare our deep learning method with other established conventional techniques.

Our initial motivation for employing a deep learning approach, particularly convolutional operations, comes from their documented success in meteorological applications, as outlined in lines 66-80 of our manuscript. The effectiveness of convolutional neural networks (CNNs) arises from the way they treat the dataset, taking into account the local context (such as 2D and 3D neighborhood structures) which, in our case, is crucial for understanding the spatial configurations of weather patterns and their changes over time. They further exploit hierarchical feature extraction. This focus is discussed in lines 194-202 in our paper. In particular, we highlight our motivation to employ a convolutional auto-encoder architecture, inspired by its success in weather prediction by Weyn et al. (2019, 2020). These studies illustrate the architecture's proficiency in capturing spatiotemporal patterns within weather datasets, a point we emphasize in lines 526-527 of our manuscript.

The flexibility of deep learning methods, which allows us to customize the latent space to align with our specific research goals, is a further motivation to use this approach in our research (line 238 of our manuscript). Specifically, our aim is for data points yielding similar PMSET patterns to lay more closely in the MCAE latent space, as detailed in lines 251-254. In the discussion section lines 549-559, we evaluate the MCAE's latent space representation of the dataset. We compare this with the traditional data space representation, which is dominated by the seasonal cycle. Additionally, we compare it to the PMSET index representation, which focuses solely on PMSET similarities. This level of flexibility in data representation, unachievable with conventional methods, further motivate us to use the deep learning approach.

In the opening paragraph of the discussion (lines 528-539), we aimed to assess the data representation obtained from the FAE. Particularly, in lines 536-539, we compare unsupervised clustering results performed on the data space and the FAE's latent space. This comparison reveals that the FAE's latent space allows us to discern datapoints corresponding to transitional seasons, such as autumn and spring. This observation leads to our concluding remark in lines 538-539, where we state that the FAE offers an effective method for creating a concise and informative representation.

We acknowledge the importance of comparing our convolutional auto-encoders with simpler, less computationally expensive methods such as SOMs and EOF analysis. Therefore, we implemented unsupervised clustering using SOM and EOF and compare their performance with our clustering analysis. The EOF analysis was applied to the whole Control and Experiment runs data points treating them as one time series. The first five leading EOFs were only used for spanning the reduced representation. The corresponding PC deriving from these EOFs serves as the features within the reduced representation. The five leading EOFs account for 64.7% of the total variance within our dataset. The centers of these classes, along with their monthly occurrence frequencies, are depicted in Figure 1 and Figure 2, respectively. Moreover, we configured the SOM to categorize the data into six classes, arranged in a 3x2 grid, over 100 iterations. Due to the sequential nature of SOM training, leveraging the parallel processing capabilities of GPUs offers challenges, necessitating reliance on the sequential processes of CPUs. Consequently, the SOM training was time-consuming, requiring approximately one week on our local server, due to the extensive dataset and its large input shape. This extensive training time limited our ability to fine-tune the SOM parameters to optimize performance. The class centers of these SOM classes, along with their monthly occurrence frequency are illustrated in Figure 3 and Figure 4, respectively. The class centers and occurrence frequencies resulting from these conventional clustering approaches are similar to those obtained by performing k-means clustering directly on the data space. We will elaborate on this comparison in the Appendix section of the revised manuscript.

Additionally, we will include t-SNE visualizations on different data representations in the Appendix to further argue the advantages of our chosen methodology over traditional approaches. This visualization aims to provide a compelling argument for the proficiency of our method by highlighting the distinct data embeddings and the insights they offer. This addition will not only enhance the robustness of our argument but also provide a comprehensive evaluation of the convolutional auto-encoder effectiveness in capturing complex spatiotemporal patterns in response to the excreted

We appreciate your constructive feedback, which has guided us to further validate our methodology. We believe that these additional comparisons will emphasize the importance of using convolutional autoencoders in our research.

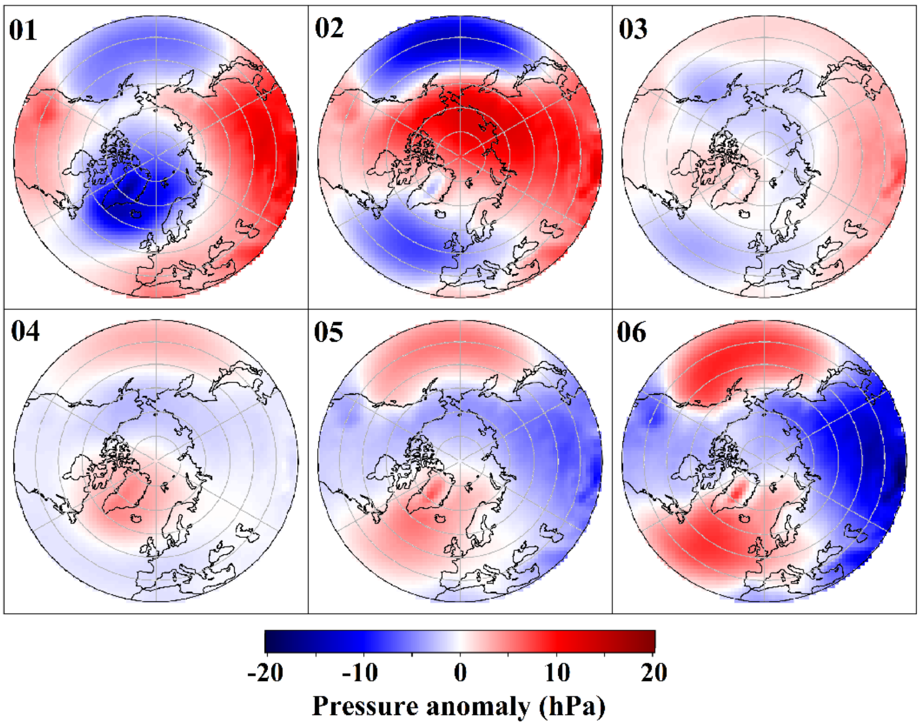

Figure 1: Mean MSLP fields for clusters derived from k-means clustering (6 clusters) on PC time series derived from the 5 leading EOFs. The fields for each cluster are plotted as anomalies relative to the mean MSLP field for all data points, encompassing both the Control and Experiment runs.

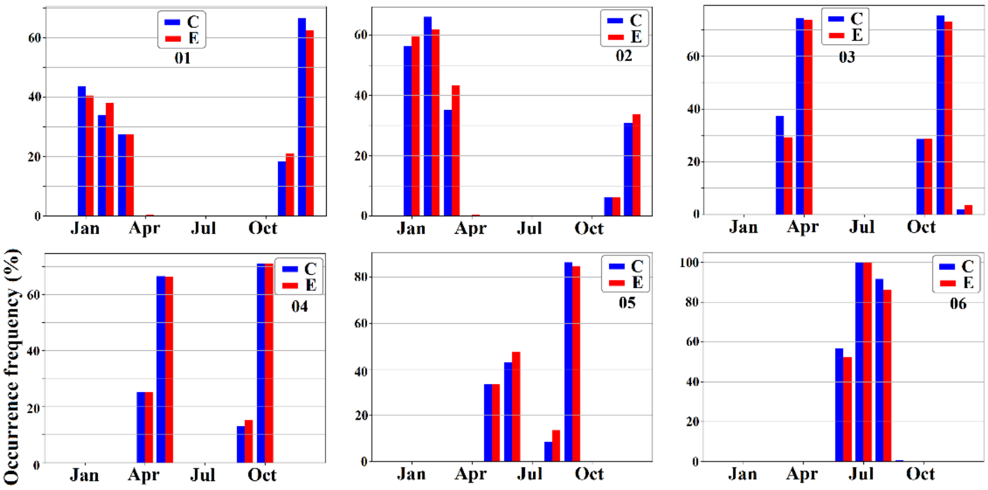

Figure 2: The monthly occurrence frequency of the 6 clusters showed in Figure 1 for the Control (blue) and Experiment (red) runs.

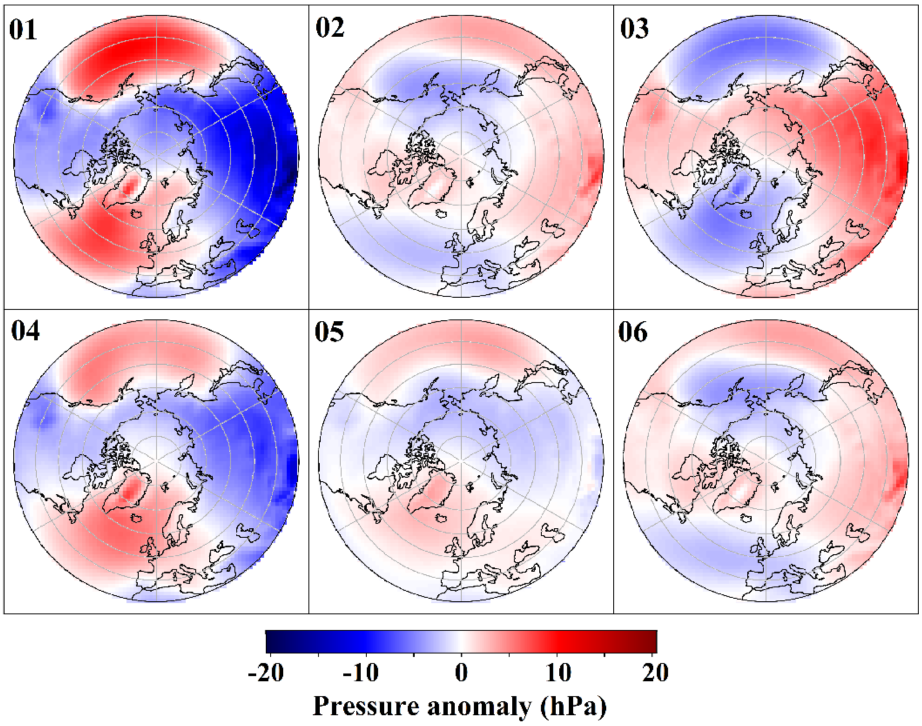

Figure 3: Same as Figure 1, but the clusters were derived from applying SOM clustering on the data space.

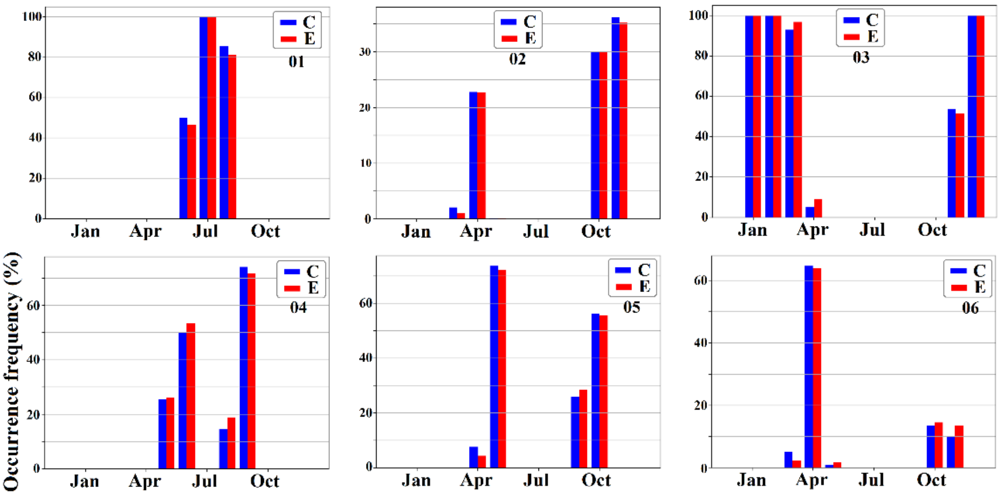

Figure 4: The monthly occurrence frequency of the 6 clusters showed in Figure 3 for the Control (blue) and Experiment (red) runs.

What new physical or dynamical insight do we learn? I am curious about the new insights or knowledge gained from the new clustering method employed in this study. Lines 566-567 mention that well-established large-scale circulation patterns (e.g., NAO, AO, PNA) are identified, and consistent dynamical responses in the troposphere and stratosphere can be demonstrated. I assume that similar conclusions may be drawn from other clustering methods as well. Could we uncover new dynamical pathways in which the Arctic responds to European radiative forcing differently from previous understandings based on stratosphere-troposphere coupling? It would be helpful if the authors could create a table summarizing the dynamical responses associated with each cluster, indicating which dynamical responses are already known and which are new. Similarly, Lines 324-327 discuss the seasonality changes. What do these seasonality changes signify physically, and what can we learn from them?" In addition, what is the separation of WCVC and FSDC components brings us new insights?

Thank you for your insightful feedback. To respond to your comment first we need to note a key aspect of different clustering methodologies, which is the dependency of clustering outcomes on the chosen representation space. Different representation spaces can influence which data points are considered close to each other according to the specific geometry of that space. Clustering algorithms like k-means then group data points based on their proximity in the chosen representation space. For instance, the clustering performed on PMSET indices (Figure B1) is guided solely by similarities in PMSET indices. That is, data points with similar PMSET indices lie closer together in this representation. Conventional clustering methodologies only focus on data point similarities using standard measures of similarity, such as Euclidean pixel-wise distance between data points or PCA-reduced (Principal Component Analysis) space representation. One of the main contributions of this paper is to introduce a deep-learning based method that generates an effective representation of data points. This representation accounts for both the similarities among the data points and their alignment with a target objective similarity, here the PMSET patterns, to effectively distance the data points.

In the offered representation, both the similarity in data points' spatiotemporal patterns and the associated PMSET pattern play a role in the data points' proximity in the representation and consequently determine the clustering results. The expression in lines 566-567 emphasizes the fact that with a focus on the similarity in data points spatiotemporal patterns and the associated PMSET pattern, we were able to have patterns with a similar focus as the well-established circulation regime. While similar patterns might be recognizable through other clustering techniques, our method enhances the interpretability and relevance of these patterns by grounding them in the specific context of PMSET dynamics. Thus, the new insight we offer is not merely the identification of these patterns but a deeper understanding of their dynamical significance, facilitated by a clustering approach that integrates physical relevance with data-driven techniques.

Thank you for pointing out the need for a more detailed discussion on the physical significance of seasonality changes and the insights gained from separating the WCVC and FSDC components in our analysis. Our findings indicate that the applied forcing may not necessarily induce new circulation patterns. Instead, we explored how alterations in the grouped circulation regimes within the Experimental run have contributed to observed anomalies in critical Arctic climate variables. We attribute these changes within the grouped circulation regimes to the applied forcing, as it is the only transient perturbation in the Experiment run.

One of the main goals of the paper is to identify the contributions of various classes to the anomalies observed as a consequence of the applied forcing. We tried to understand the dynamic interactions of each pattern that led to the observed anomaly. In addressing them, we developed our class contribution formulation. We reformulated the anomalies to attribute them to different circulation groups. We identified two main mechanisms through which a circulation cluster can contribute to the anomaly: 1) through changes in the occurrence frequency of that cluster as a result of the forcing (FSDC), and 2) through slight adjustments in the cluster's mean characteristics as a result of the forcing (WCVC). This framework helps us delineate how changes in circulation patterns, whether in occurrence frequency or in the cluster's mean characteristics, contribute to the observed anomalies in the key climate variables. This approach not only enhances our understanding of the dynamical interactions leading to the observed anomalies but also offers valuable insights into the complex ways in which atmospheric circulation responds to external forcing. It underscores the importance of considering both the occurrence and the mean characteristics of circulation patterns in assessing their impact on climate anomalies.

Regarding uncovering new dynamical pathways, we utilized this advanced methodological approach to understand the role of different circulation patterns in forming the anomalies. These patterns were grouped based on similarities in their spatiotemporal behavior and their associated PMSET patterns. This grouping and the class contribution formulation helped us understand how these interactions lead to observed anomalies resulting from negative radiative forcing over Europe. Among the main findings are:

- In most cases, the change in a cluster's main behavior is more significant than changes in the occurrence frequency of that cluster when considering the cluster's contribution to the anomaly. For example, the decreased occurrence frequency of a cluster associated with a high-pressure system over Northern Eurasia and Scandinavia in autumn, observed in the Experiment run, led to reduced upward wave propagation. However, a slight adjustment in the mean behavior of this cluster in the Experiment run resulted in an increase in upward wave propagation. This increased wave propagation dominates the cluster's contribution to the observed anomaly in upward wave propagation.

- Our method's capability to attribute observed anomalies, such as the SIC loss in the Barents-Kara Seas in autumn, the warm Arctic and cold mid-latitude in winter, and others.

We will highlight this argument more in the revised version. Moreover, we will try to put all the most meaningful results in a more compact format like a table. This approach will not only make the information more concise but also emphasize our contributions more effectively. These modifications will undoubtedly make the paper more accessible and emphasize the novelty and significance of our contributions.

Linking the responses to European radiative forcing. I noticed that the discussion on the results seems to focus less on the direct response to radiative forcing and more on the subsequent atmospheric circulation responses and PMSET. For example, how does the European radiative forcing lead to increased upward EP flux for cluster 3 in SON (Figure 17)? Or how does the localized radiative forcing in Europe give rise to changes in 2m temperature across the entire Northern Hemisphere, as depicted in Figure 13? Some of the temperature increases appear contradictory to the cooling effect of aerosol negative forcing (or specifically here, the cloud forcing).

Thank you for highlighting the need for a clearer connection between the response to European radiative forcing and the observed atmospheric circulation change in our analysis. We considered the large-scale circulation regime as the main mechanism by which our local forcing influences remote regions such as the Arctic. This is why we performed clustering to group the similar large-scale circulation regimes, and analyzed the impact of changes in these grouped circulation regimes on the observed anomaly, as a consequence of the applied forcing. Given our focus on Arctic climate variables, we utilized the PMSET pattern associated with each circulation as an additional target similarity measure for our clustering. The PMSET pattern, indicative of how each circulation regime transfers energy into the Arctic, is crucial for determining Arctic climate conditions.

Our analysis reveals that while the forcing may not introduce new circulation patterns, it can modify existing circulation patterns in two significant ways: by altering the occurrence frequency of a pattern (the FSDC component of the class contribution) or by changing the mean characteristics of the circulation pattern (the WCVC component of the class contribution). Consequently, we calculated each class's contribution to the observed anomaly. As mentioned in line 287 and equation 2 of the manuscript, the sum of these class contributions yield the observed anomaly. Each class's contribution can then be decomposed into two parts: WCVC and FSDC. This decomposition enables us to attribute different aspects of the anomaly to specific clusters and to the particular type of change within that cluster, whether it's a change in frequency or in mean characteristics. In response to your comment, we will enhance the manuscript to more explicitly draw these connections between European radiative forcing and its effect on Arctic climate via changes in atmospheric circulation.

Thank you for your insightful questions about the specific impacts of European radiative forcing on different atmospheric dynamics. These phenomena indeed highlight the complex and nonlinear responses of the climate system to localized radiative forcing. We demonstrated that the increased upward EP flux for C3 during SON is primarily attributed to the adjustment in the mean characteristics of this cluster in the experimental run (WCVC component), despite the reduced occurrence frequency of this cluster, which is commonly known to decrease upward wave propagation. However, this reduction is not the dominant factor in this class's contribution (for more details, see lines 609-629). Furthermore, we attributed the observed warm Arctic and cold midlatitude pattern to changes in the mean characteristics of C1, as detailed on lines 600-607 of the manuscript.

As expected, the forcing does not always yield uniform results; its effects differ depending on the present circulation regime. Unlike linear analyses, which evenly attribute the anomaly across all clusters or data points, our approach captures the nonlinear behavior of the climate system. This means that the forcing leads to different contributions by each class depending on how the forcing change the cluster behavior, and the sum of these contributions is equal to the observed anomaly (see equation 2).

Thanks to your feedbacks, we acknowledge that the presentation of this content within the paper may not have been sufficiently clear. In general we are committed to revise the manuscript structure to better represent our analytical framework.

Minor comments:

Lines 58-60: perhaps the authors considering to cite two new studies on this topic: Xu et al. (2023) and Liang et al. (2024).

Thank you very much for recommending the studies by Xu et al. (2023) and Liang et al. (2024). We agree that these recent contributions are highly relevant to our work and can enrich the context of the manuscript. We will incorporate these studies in the revised version of our manuscript.

Line 123 and Figure 1: why there are statistically significant radiative forcing increase in eastern Siberia, Asia, and North Pacific?

You've raised an interesting point. These increases are indeed noteworthy and could be attributed to dynamical feedback mechanisms related to the atmospheric circulation change by European radiative forcing. While we acknowledge the importance and potential implications of these observed increases, a detailed exploration of their causes and consequences lies beyond the current scope of our research.

Lines 161-162: why 8 days? Does this indicate a certain physical process dominating?

The choice of an 8-day period for our data points was motivated by a combination of methodological considerations and preliminary testing to identify a timeframe that captures dynamic atmospheric processes. In our initial tests, we observed that an 8-day period yielded the best performance regarding the DL algorithm reconstruction loss and overfitting measure (line 163 of the manuscript). This performance could be due to its efficieny to capture atmospheric dynamics or just related to our deep learning architecture. The choice of 8 days aligns well with our deep-learning architecture, particularly the use of 3D convolutional pooling and upsampling within our autoencoder architecture. The binary nature of the number 8 (2^3) potentially enhances processing efficiency and effectiveness within this architecture.

Line 230: the maximum PMSET? Or both maximum and minimum PMSETs?

Thank you for your keen observation. Indeed, our analysis considers both the maximum and minimum PMSETs, not solely the maximum PMSET. We will correct this typo in the revised manuscript.

Lines 246-248: could the authors provide a figure to demonstrate this grid transformation?

Thank you for suggesting the inclusion of a figure to demonstrate the transformation applied to angular features in the MCAE's latent space (explained in lines 246-248). To address the continuity challenge between angles close to 0 and 360 degrees, we indeed transform each angular feature into two dimensions using sine and cosine functions. This approach effectively preserves the circular nature of angular measurements, ensuring that values close in angular space (for example, 355° and 5°) remain proximate in the transformed two-dimensional space, despite their apparent numerical distinction.

For instance, an angular feature represented by a value within the 0-360 degree range can create discontinuity issues where angularly adjacent values appear numerically distant. By converting each angle into a pair of coordinates, using the sine and cosine of the angle, we create a two-dimensional representation where the proximity of angles accurately reflects their true angular relationship. In Figure 5, points A and B are represented by their angles, θ1 and θ2, respectively. Despite θ2 being numerically far from θ1, their two-dimensional representations (using sine and cosine) on the unit circle are close, accurately reflecting the angular proximity.

We initially opted for a written description of the transformation in the paper. However, in response to your valuable feedback, we will consider the possibility of adding a similar figure that represents the transformation to enhance clarity and understanding.

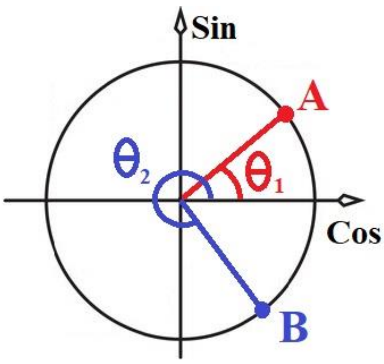

Figure 5 visualization of the transformation applied to angular features to address the angular continuity challenge

Lines 298 and 301: Are there any reference papers for WCVC and FSDC?

The concepts of the Within-Cluster Variability Contribution (WCVC) and Frequency-weighted Seasonal Deviation Contribution (FSDC) introduced in our study represent novel contributions to the field. These were developed as part of our class contribution formulation, aimed at examining the dynamic interactions of circulation patterns and their roles in producing the observed anomalies in the Experiment run. This approach allows us to attribute anomalies to specific changes within circulation clusters, either through variations in their occurrence frequency (FSDC) or adjustments in their mean characteristics (WCVC) in response to radiative forcing.

Given the innovative nature of this framework, there are no direct reference papers that discuss WCVC and FSDC in the context we have applied them. To the best of our knowledge, our study is the first to formalize these concepts as distinct mechanisms through which atmospheric circulation clusters can influence climate anomalies.

In section 2.6 of our manuscript, we comprehensively explained these concepts. We believe that our work may serve as a foundation for future research in this area, potentially encouraging further studies that attribute anomalies to different components of the system.

Figure 8: it seems the signal-to-noise ration is quite small. How could the authors say these responses are important?

As emphasized in lines 339-344 of the manuscript, in Figure 8, our primary objective is to demonstrate the effectiveness of our clustering methodology in identifying and classifying similar atmospheric patterns across different simulations and throughout different seasons. The class centers were calculated relative to the seasonal mean fields depicted in the figure's first row. The low variation in class center patterns across seasons underscores the robustness and consistency of our clustering approach, even without explicitly incorporating information about the seasonal cycle into the network.

Regarding the signal-to-noise ratio mentioned, it's crucial to recognize that the subtlety of the variations among the class centers highlights the method's sensitivity and precision in capturing the dynamics of the climate system. We appreciate this opportunity to clarify the intention and findings related to Figure 8.

Figures 9 and Figure 10: combine these two figure into one figure?

Thank you for the suggestion to combine Figures 9 and 10. We see the value in combining these figures, and we will undertake this in the revised version.

Figure 12: the arrows are hardly seen. Perhaps the authors can try to enhance the visibility of the arrows.

Thank you for pointing out the visibility issue with the arrows in Figure 12. We appreciate your suggestion and will enhance the visibility of the arrows in the revised version of the figure.

Lines 458-460: but the seasonal distribution changes?

Thank you for bringing attention to the subtle change in the seasonal distribution of sudden stratospheric warming (SSW) occurrences. Table A3 indicates a subtle shift in the seasonal distribution of SSW occurrences in the Experiment run, with a decrease by one occurrence in autumn and a compensatory increase by one occurrence in winter.

As illustrated in Figure 16, the autumn season shows an increase in the zonal mean zonal wind in the high-latitude upper stratosphere and a decrease in the lower stratosphere. This pattern of variability does not directly correlate with the slight decrease in SSW occurrences in autumn, suggesting that the observed changes in zonal wind patterns cannot solely be attributed to the frequency of SSW events.

Similarly, during winter, we observed an increase in high-latitude middle atmospheric zonal mean zonal wind, coinciding with an increase in occurrences of SSW events. Generally, an increase in SSW events would be expected to have a converse effect, namely a decrease in the zonal mean zonal wind. Therefore, attributing the observed change in wind patterns directly to the singular additional SSW occurrence is not straightforward.

This analysis leads us to conclude that while the seasonal distribution of SSW occurrences experiences minor adjustments, these changes do not significantly alter our findings regarding our upper atmospheric analysis.

Figure 17: the EP flux divergence does not exactly match the pattern of zonal wind anomalies in some seasons and clusters. How could we relate the EP flux change to zonal wind change?

Thank you for highlighting this subtle observation. It is important to note that the EP flux divergences are the momentum deposited by the resolved waves in the model. In addition to the resolved Rossby waves the changes in the zonal wind can be influenced by the non-resolved parameterized gravity waves (GWs). Changes in the zonal mean zonal wind due to the resolved waves may affect the parameterized momentum depositions by the GWs. This in turn, could affect the prevailing zonal flow and induce changes in the residual circulation below the breaking level. Such dynamics might explain the anomalies of the zonal mean zonal wind far away from the regions of significant EP fluxes (Cohen et al., 2013; Limpasuvan et al., 2016; Chandran et al., 2014). We appreciate your insightful comment and will further clarify this in our revised manuscript.

Lines 540-542: the authors mentioned that the external forcing can modify the circulation patterns. But in Lines 778-680, the authors conclude that the radiative forcing only alter the existing circulation patterns, and does not introduce new patterns. These two sentences seem contradictory somewhat. Could the authors clarify and reconcile these statements?

Thank you for highlighting the need for clarification between our discussion on the general effects of external forcings on circulation patterns and our specific findings related to negative radiative forcing over Europe.

In the general context (lines 540-543), we mention "External forcings can modify circulation patterns in complex and nonlinear ways (Gillett and Fyfe, 2013; Hannachi et al., 2017) by introducing new spatiotemporal patterns, changing the preferred circulation patterns, altering their frequencies, or by a combination of them." However, when focusing on the specific context of negative radiative forcing over Europe (lines 678-680), we mentioned "Our study revealed that negative radiative forcing over Europe, resembling the heterogeneous radiative forcing exerted by aerosols, influences the climate system by altering existing circulation patterns and their frequencies without introducing new patterns." Therefore, while external forcings, in general, have the capacity to induce a wide range of modifications in circulation regimes, our research specifically found that negative radiative forcing over Europe modifies the circulation by adjusting existing patterns and their frequencies, rather than introducing entirely new circulation patterns.

Line 689: upper troposphere lower stratosphere —> upper troposphere and lower stratosphere?

Thank you for pointing out this error. We will correct it in the revised version of the manuscript.

References:

Xu, M., Tian, W., Zhang, J., Screen, J.A., Zhang, C. and Wang, Z., 2023. Important role of stratosphere-troposphere coupling in the Arctic mid-to-upper tropospheric warming in response to sea-ice loss. npj Climate and Atmospheric Science, 6(1), p.9.

Lee, C.C., 2017. Reanalysing the impacts of atmospheric teleconnections on cold-season weather using multivariate surface weather types and self-organizing maps. International Journal of Climatology, 37(9), pp.3714-3730.

Liang, Y.C., Kwon, Y.O., Frankignoul, C., Gastineau, G., Smith, K.L., Polvani, L.M., Sun, L., Peings, Y., Deser, C., Zhang, R. and Screen, J., 2024. The Weakening of the Stratospheric Polar Vortex and the Subsequent Surface Impacts as Consequences to Arctic Sea Ice Loss. Journal of Climate, 37(1), pp.309-333.

Skific, N., Francis, J.A. and Cassano, J.J., 2009. Attribution of projected changes in atmospheric moisture transport in the Arctic: A self-organizing map perspective. Journal of Climate, 22(15), pp.4135-4153.

Reply References:

Chandran, A., Garcia, R. R., Collins, R. L., & Chang, L. C.: Secondary planetary waves in the middle and upper atmosphere following the stratospheric sudden warming event of january 2012. Geophysical Research Letters, 40 (9), doi: https://doi.org/10.1002/grl.50373, 2013.

Cohen, N.Y., Gerber, E.P., Bühler, O.: Compensation between resolved and unresolved wave driving in the stratosphere: Implications for downward control. Journal of the Atmospheric Sciences 70, 3780 – 3798, doi:10.1175/JAS-D-12-0346.1, 2013

Gillett, N. and Fyfe, J.: Annular mode changes in the CMIP5 simulations, Geophysical Research Letters, 40, 1189–1193, https://doi.org/10.1002/grl.50249, 2013.

Hannachi, A., Straus, D. M., Franzke, C. L., Corti, S., and Woollings, T.: Low-frequency nonlinearity and regime behavior in the Northern Hemisphere extratropical atmosphere, Reviews of Geophysics, 55, 199–234, doi: 10.1002/2015RG000509, 2017.

Limpasuvan, V., Orsolini, Y. J., Chandran, A., Garcia, R. R., & Smith, A. K.: On the composite response of the mlt to major sudden stratospheric warming events with elevated stratopause. Journal of Geophysi-cal Research: Atmospheres, 121 (9), 4518-4537, doi: 10.1002/2015JD024401, 2016.

Weyn, J. A., Durran, D. R., and Caruana, R.: Can machines learn to predict weather? Using deep learning to predict gridded 500-hPa geopotential height from historical weather data, Journal of Advances in Modeling Earth Systems, 11, 2680–2693, doi:10.1029/2019MS001705, 2019.

Weyn, J. A., Durran, D. R., and Caruana, R.: Improving data-driven global weather prediction using deep convolutional neural networks on a cubed sphere, Journal of Advances in Modeling Earth Systems, 12, e2020MS002 109, doi:10.1029/2020MS002109, 2020.