SHORT COMMENT (combined for both Reviewer #1 and #2)

We thank the two anonymous reviewers for their valuable comments and constructive reviews. Both reports have been very helpful for us to identify those sections in the manuscript that require further improvements. Since the two reviews share several inputs and raise similar points, we provide here an initial short response addressing comments from both reviews. This is *not* intended to be our final response, but a pitch of how we intend to address the reviews, and in the spirit of open discussion provided by Weather and Climate Dynamics, we hope to inspire a fruitful conversation with the reviewers to further improve our work. In our final response, we will describe in full detail how we plan to take all comments into account and implement them in the revised version of the paper. In particular:

- Both reviewers indicate that the causal maps appear "noisy" and that "a discussion on the sensitivity of results to data-length would also be useful". We have already done a range of sensitivity tests (see below) showing that the identified large-scale patterns are indeed robust. This step also makes the causal maps less noisy such that the robust patterns emerge better. This thus improves the visual appearance and interpretability of figures 3, 4 and 5.

- Further, we will improve on the design of the causal maps to make it easier for the reader to identify those regions in the maps that are described in the text

- We will improve the literature review and provide more context along with additional references to similar techniques previously applied to similar research questions

- We will improve the explanation of both the causal discovery algorithm and causal maps in the Methods section to help the readers better understanding each step

- We will highlight better why using causal discovery algorithms gives an advantage with respect to simple correlation techniques and why our work adds information on the tropical – mid-latitude interaction topic

- In general, we will work on the main text to improve the readability throughout the entire manuscript and make it is easier for the reader to follow the explanation of both the Methods and the Results sections

We have already implemented part of the suggested additional analyses, to verify the robustness and sensitivity of the results, and we provide some of the results below.

ROBUSTNESS OF CAUSAL MAPS

Both reviewers have pointed out that the causal maps appear noisy and that a test of robustness on how sensitive these results may be to different temporal periods is needed. This is an important point, and hence we addressed it immediately.

To address both issues, we have pursued the following strategy: we calculate causal maps several times (10x) by removing each time the 10% of the total time series length. Here, we have 40 years of data, we thus remove in each step 4 consecutive years. As an example, in the first iteration we remove years 1979-1982, in the second iteration years 1983-1986 and so on. As a result, we obtain

10 causal maps. Then, collapse the information to one single map (Fig. 1 in this document) showing (for each grid-point) the fraction of times in which a significant causal link has been found based on the reduced data sets. Figure 1 reported below shows the results of this process for MCA 1. Regions that show a dark purple colour correspond to those identified in all ten causal maps. We can then set a threshold, for example 0.7 (meaning that a certain region is identified in 70% of the times) and use this as a mask for our causal maps (See Fig. 2 in this document). This results in a less noisy plot, where the plotted regions also represent the robustness of the analysis.

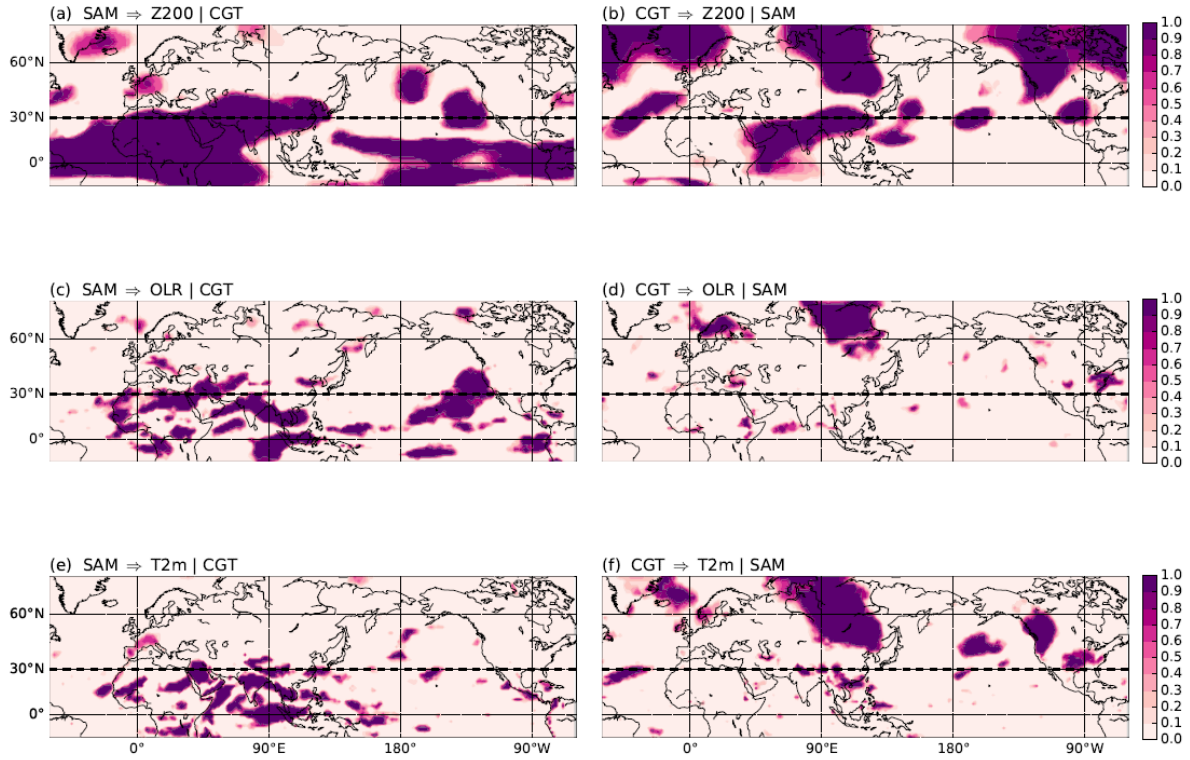

Fig. 1: Robustness test for causal maps as shown in Fig. 3 of the main manuscript. Dark purple shows regions that show a significant causal link at alpha = 0.05 (after applying the false discovery rate correction) in all the ten causal maps obtained by iteratively removing a set of 4 consecutive years each time, with a value of 1 meaning that a certain regions is always identified.

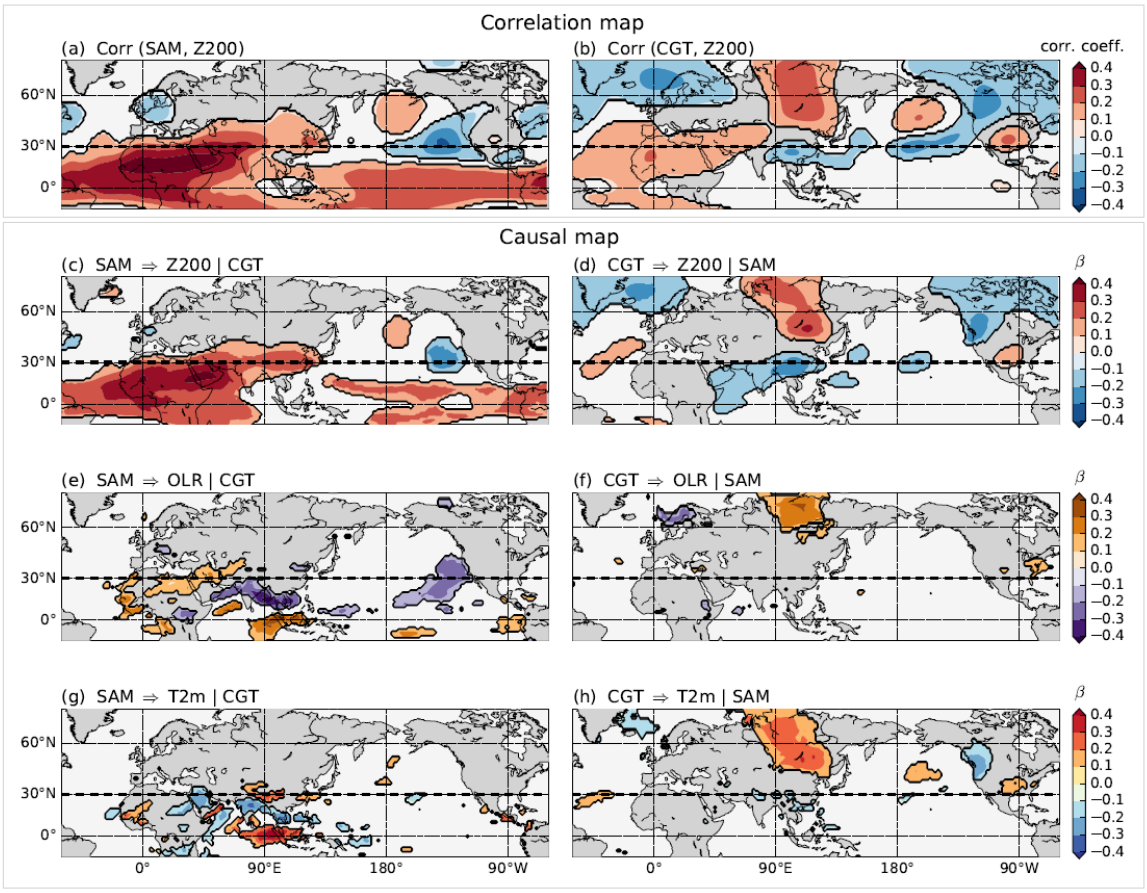

Fig. 2: As for Fig. 3 in the main text but masking out all regions that appear less then 70% of the times (see Fig. 1 in this document).

COMPARISON WITH EOF PATTERNS

We address the comment n3 by reviewer #1 on how our MCA patterns compare to EOF patterns. In Fig. 3, we show the first 5 EOF patterns for both Z200 and OLR. We calculate the spatial correlation between EOF and MCA patterns. For Z200, MCA 1 shows the strongest correlation with EOF 2 (r ~ 0.8), unsurprisingly since this pattern represents the circumglobal teleconnection pattern, which has been previously shown to be linked to the second EOF of Z200 (Ding and Wang 2005, Di Capua et al. 2020). MCA 2 has a stronger spatial correlation (r ~ 0.6) with EOF 1. For OLR, MCA 1 shows the strongest correlation with EOF 2 (r ~ 0.5), while MCA 2 has the strongest correlation with EOF 5 (r ~0.4).

Thus, with the only exception of OLR MCA 2, all MCA patterns are closely related with the first two EOFs for both Z200 and OLR. This comparison is useful to show how important each pattern is when a dimension reduction analysis is applied on Z200 and OLR separately. Note that the amount of variance explained is relatively low, but this depends on the fact that the interannual variability has been removed, thus leaving only the disturbances from the year-specific mean state. However, since in our present work, we are interested in identifying those patterns that evolve *simultaneously* (due to some dynamical coupling between the two fields), we applied MCA to identify those patterns that can explain shared covariance, which is an objective that cannot be addressed by using EOF analysis alone.

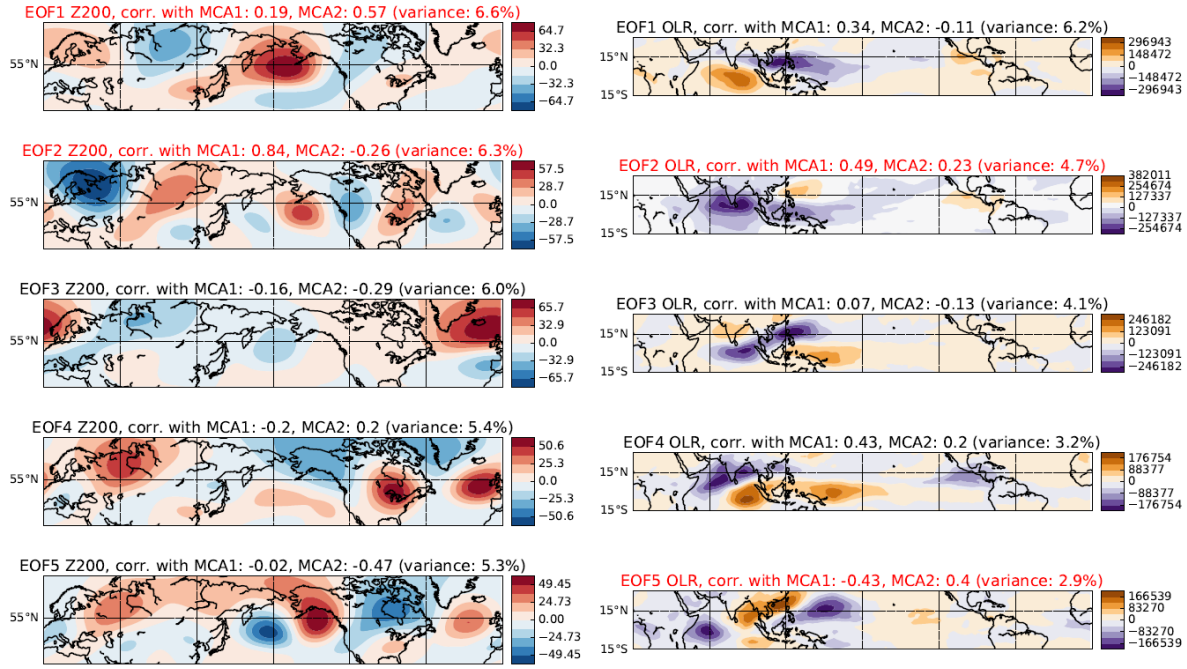

Fig. 3: EOF analysis. The right column shows the first five EOF patterns for Z200, the left column shows the first five EOF patterns for OLR. In the title of each panel, the spatial correlation values with the MCA patterns reported in Fig. 2 of the main manuscript are shown. Red font highlights those EOFs that exhibit the strongest overall correlation with the MCA patterns discussed in our manuscript.

MCA WITH VERTICAL VELOCITY AND VELOCITY POTENTIAL

Following the recommendations of reviewer #1, we provide a sensitivity test for the identified MCA patterns by substituting OLR with velocity potential and vertical velocity. Note that, we have originally applied MCA on mid-latitude Z200 and tropical OLR because we are interested in studying the relationship between mid-latitude circulation patterns and tropical convection. Thus, we are restricted to variables representing tropical convection when attempting to provide a comparable analysis.

We originally selected OLR both because it captures strong convective clouds (but in a more smoothed way than expected for direct rainfall estimates), and because OLR is also used, for example, to define the BSISO index that describes the essential evolution of convective activity over the Indian Ocean region. Figure 4 shows the first two MCA patterns for Z200 paired with vertical velocity, while Fig. 5 shows the same for Z200 paired with velocity potential.

As we are interested in a proxy for convective activity, we perform our MCA analysis using vertical velocity, another proxy of convection (as strong upward vertical motions occur along with strong convective activity). The MCA patterns obtained when pairing vertical velocity with Z200 show highly consistent results with respect to those found for Z200 and OLR (Fig. 2 in the main text), thus demonstrating the robustness of the original MCA results obtained with OLR. (Note that in Fig. 4 upward motion has a negative sign since vertical velocity is expressed in Pa/s).

When we use velocity potential (Fig. 5 in this document), the original MCA 1 pattern obtained using OLR is still well recovered (with a wave-5 pattern in Z200 and low velocity potential over the Indian summer monsoon region). The MCA 2 pattern however shows a less pronounced agreement, only

partly capturing the OLR pattern in the western Indian Ocean but failing to represent the WNPSM convective activity. A reason for this discrepancy is that velocity potential provides an even smoother proxy for divergence, which is very strong in the Indian monsoon region, and apparently less pronounced in relation to the WNPSM.

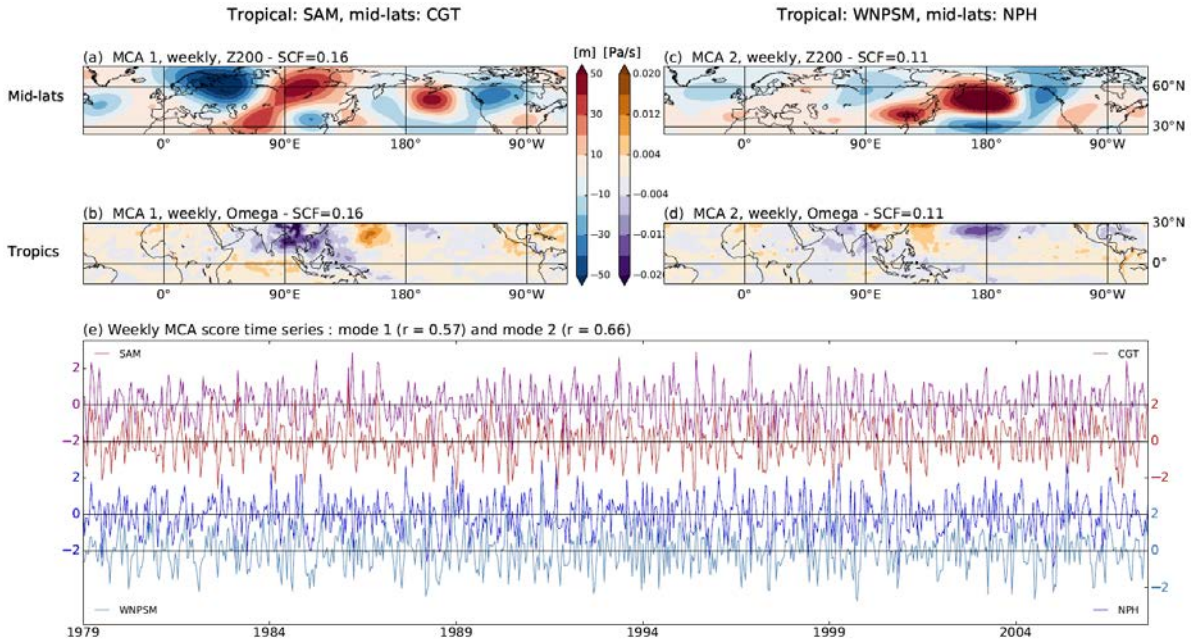

Fig. 4 MCA results for Z200 and velocity potential.

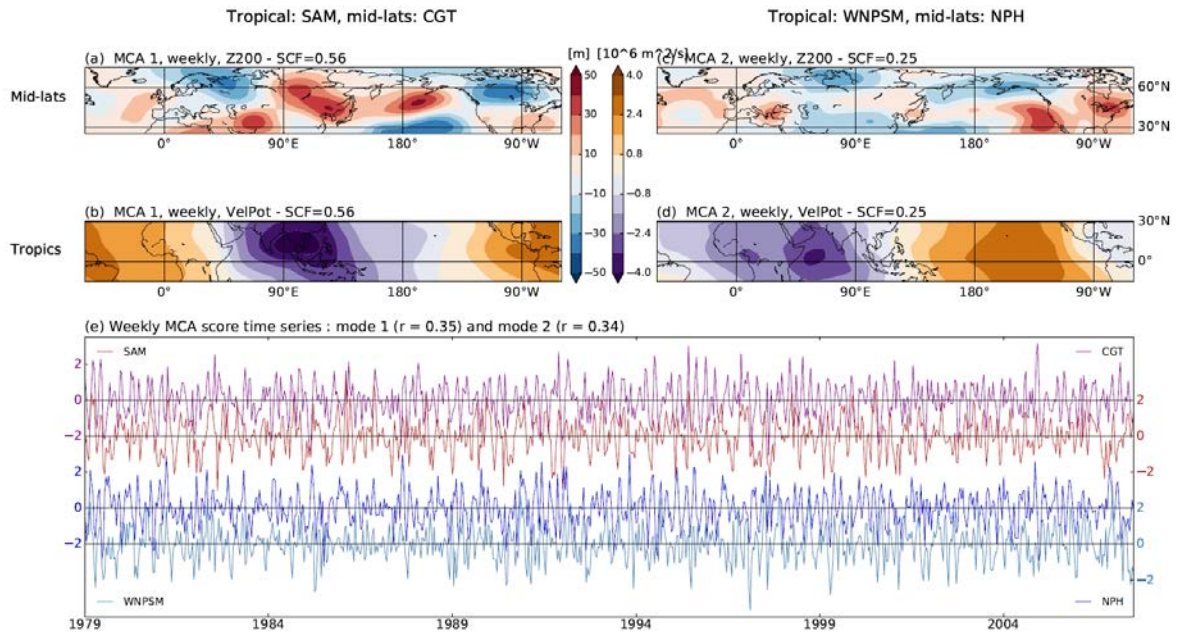

Fig. 5 MCA results for Z200 and vertical velocity.