

Response to review #1

We thank anonymous referee #1 for the constructive review and helpful comments that have greatly helped us to improve our work in the revised manuscript. The main improvements are summarized as follows:

- We have performed a series of sensitivity tests which show that our findings are robust, and which have reduced noise in some plots and thereby their visual appearance
- We have improved the description of the methodology section and the visualisation of the results
- We have added further background literature on the use of causality methods in atmospheric science

We have taken into account all suggestions made and a point-by-point response to each comment is reported below. Please note that in the following text the referee's comments are highlighted in bold font, while our answers are in regular font.

Specific comments

1. **There is no literature review of other approaches to identifying causal relationships in climate data. One example is I. Horenko, S. Gerber, T.J. O’Kane, J.S. Risbey and D. Monselesan (2017) On inference and validation of causality relations in climate teleconnections, (In Nonlinear and Stochastic Climate Dynamics. Cambridge University Press, Eds. C. Franzke and T.J. O’Kane)**

We thank the anonymous reviewer for this suggestion. We will include in our revised version of the manuscript a paragraph briefly describing other causal approaches applied to atmospheric sciences: *“In recent years, several approaches have been applied to identify causal relationships in climate and atmospheric sciences (Runge et al., 2019b), ranging from Granger causality (McGraw and Barnes, 2018, 2020; Samarasinghe et al., 2019) to causal (Bayesian) graphical models (Pearl, 2000, Ebert-Uphoff and Deng, 2012a, 2012b; Horenko et al., 2017) and conditional independence-based network discovery methods for time series (Runge et al., 2019a). These studies have shown the ability of causal discovery tools in improving the understanding of several atmospheric circulation interactions such as Arctic – mid-latitudes connections (McGraw and Barnes, 2020; Samarasinghe et al., 2019), synoptic-scale disturbances between boreal summer and boreal winter (Ebert-Uphoff and Deng, 2012a) and the relationship between ENSO and surface temperature in the American continent (McGraw and Barnes, 2018).”*

2. **The initial application of MCA appears to perform a basic dimension reduction. The authors assert that “expert knowledge” is required in choosing the particular variables to calculate the cross covariances however there is no indication that any other combinations were examined. For example, OLR could be replaced with velocity potential – as in indices for the MJO – with similar results.**

We thank the anonymous reviewer for suggesting this interesting test and we have expanded our analysis by considering the results obtained when other variables are used. We 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 focussing on variables representing tropical convection when attempting to provide a comparable analysis. We originally selected OLR because it captures strong convective clouds (which is a smoother signal than 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. In the revised version of the manuscript, we will provide a series of sensitivity tests for the identified MCA patterns by substituting OLR with velocity potential or vertical velocity (a

proxy of convection). Figure R1 (in this document) shows the first two MCA patterns for mid-latitude Z200 paired with tropical vertical velocity (note that upward motion has a negative sign since vertical velocity is expressed in Pa/s), while Fig. R2 shows the same for Z200 paired with velocity potential. 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 original text), demonstrating the robustness of the original MCA results obtained with OLR. When we use velocity potential (Fig. R1), the MCA 1 pattern strongly resembles that originally obtained using OLR (with a wave-5 pattern in Z200 and low velocity potential over the Indian summer monsoon region). The MCA 2 pattern however shows less agreement: It correctly captures the OLR pattern in the western Indian Ocean but does not represent the WNPSM convective activity patterns. A reason for this discrepancy is that velocity potential provides a much smoother proxy for upper-level divergence than OLR, which is very strong in the Indian monsoon region, and apparently less pronounced in relation to the WNPSM. We briefly comment on this in the revised main text: *“We also investigate whether the obtained MCA patterns are sensitive to the choice of OLR in representing tropical convective activity. Using vertical velocity, another proxy for tropical convection where strong convective activity is represented by enhanced upward motions, shows qualitatively the same patterns as those in Figs. 2b,d (see Fig. S5 in the Supplementary Material). When velocity potential is used instead of OLR, the first MCA mode still closely resembles the OLR/Z200 MCA mode 1, while the second MCA mode only partly captures features in the western Indian Ocean (see Fig. S6 in the Supplementary Material).”*

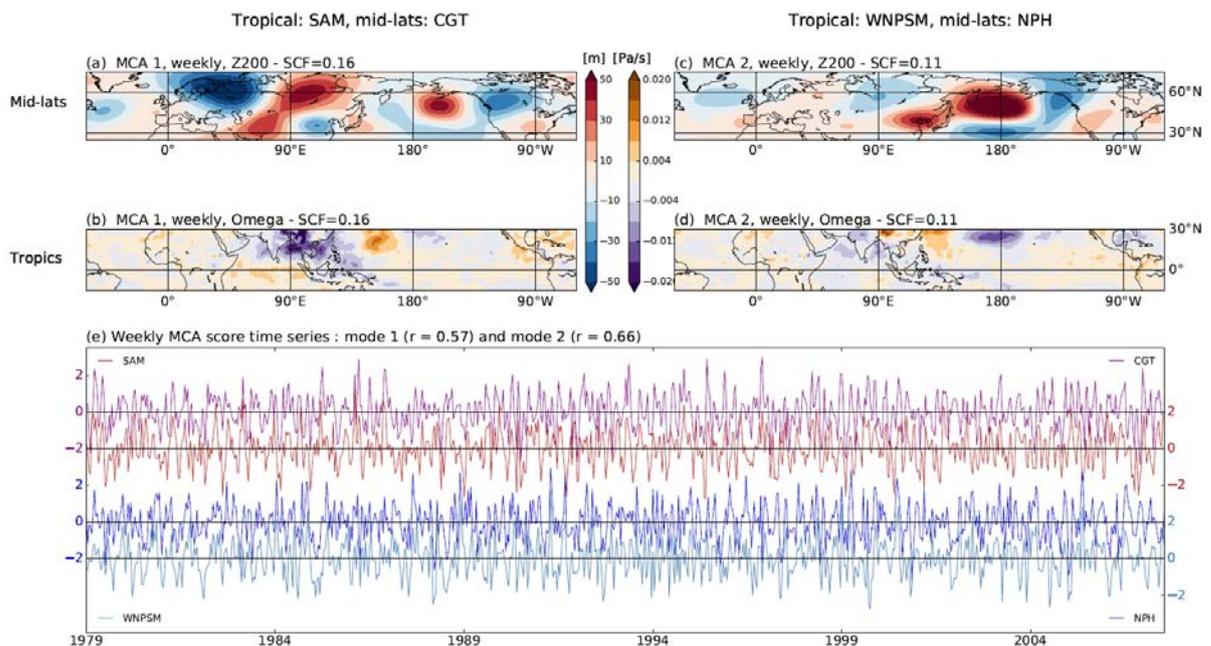


Figure R1. As for Fig. 2 in the original manuscript but for vertical velocity at 500 hPa (Omega).

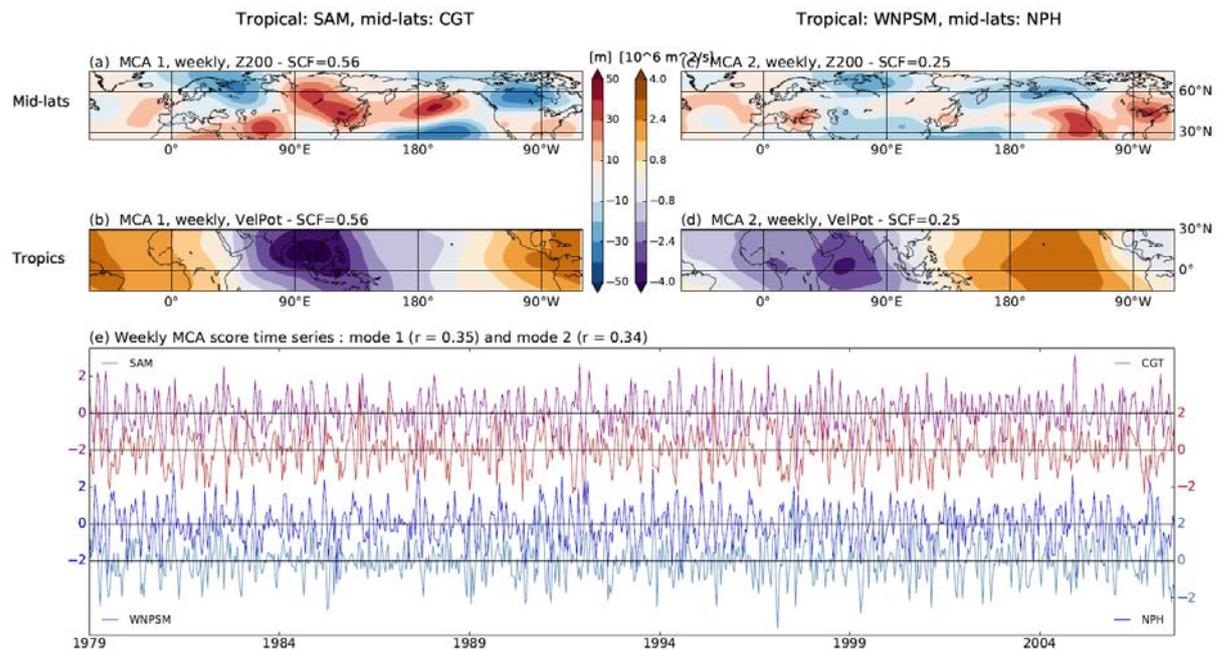


Figure R2. As for Fig. 2 in the original manuscript but for velocity potential at 200 hPa (VelPot).

The methodology applied here seems to be unable to answer if a sufficient set of covariates has been chosen apart. How, for example, do you test if the combination of actors is sufficient or even parsimonious? Can some form of information theoretic approach be applied for example Akaike or Bayesian?

We consider causal discovery here and not a prediction task of any of the actors, for which criteria such as those mentioned are indeed important. Hence, the choice of included actors is subject to the hypothesis underlying the analysis setup. One could, however, phrase causal discovery, as in Granger's work, as a prediction problem. On the other hand, a causal interpretation rests on a number of assumptions and we discuss limitations related to causal sufficiency and other assumptions made in the discussion in the revised manuscript: *"Finally, it should not be forgotten that in the context of the present work, causal interpretation rests upon several assumptions, such as the causal Markov condition, faithfulness, causal sufficiency, stationarity of the causal links and assumptions about the dependence-type (Runge, 2018). These assumptions can be violated in a real system and it is important to be aware of the associated typical challenges for causal discovery in Earth system sciences (Runge et al., 2019). Causal sufficiency requires that all relevant actors in a specific system are accounted for. Here, given the limited set of actors analysed, we cannot rule out that other excluded actors may act as important (common) drivers. Therefore, the obtained links can be considered causal only with respect to the specific set of actors used here. However, the absence of a link can still be interpreted as a likely indication that no direct physical connection among the respective variables exists. Moreover, we assume linear dependencies and stationarity for the detection of the causal links. While linearity has been shown to be a useful assumption in previous work (Di Capua et al., 2020), monsoon dynamics behaves partly nonlinearly and therefore, our causal networks only capture some part of the underlying mechanisms by construction. Also, the SAM teleconnections might well behave in a nonstationary manner on decadal time-scales (Di Capua et al., 2019; Robock et al., 2003). We therefore cannot rule out that (multi-)decadal oscillations such as the Pacific Decadal Oscillation may influence our results. However, the amount of reliable data is limited and this prohibits the application of nonlinear measures or study of effects of nonstationarity."*

3. Given the leading two modes of MCA appear to be in quadrature, how does MCA compare to EOF/PCA or even k-means?

We thank the anonymous reviewer for raising this point. We have now performed a comparison between MCA patterns and EOF patterns. In Fig. R3, which will be included in the revised Supplementary Material, we show the first 5 EOF patterns for both Z200 and OLR. We calculate the spatial correlation between all EOF and MCA patterns. For Z200, MCA 1 shows the strongest correlation with EOF 2 ($r \sim 0.8$). This is consistent with previous literature showing that the circumglobal teleconnection pattern (as captured by Z200 of MCA1), is linked to the second EOF of Z200 (Ding and Wang 2005, Di Capua et al. 2020). MCA 2 has a strong spatial correlation ($r \sim 0.6$) with EOF 1. For OLR, MCA 1 shows the strongest correlation with EOF 2 ($r \sim 0.5$), while MCA 2 has the strongest correlation with EOF 5 ($r \sim 0.4$). Thus, with only the exception of OLR MCA 2, all MCA patterns are closely related to the first two EOFs for both Z200 and OLR. This comparison shows that the identified MCA patterns are also on a regional level important in explaining the variability. Note that the fraction of variance explained is relatively low (for all EOFs), but this relates to the prior removal of interannual variability, thus leaving only the disturbances from the year-specific mean state. In our present work, we are interested in identifying those patterns that evolve simultaneously (due to the dynamical coupling between the two fields), and therefore we applied MCA to identify those patterns that can explain *shared* covariance, which is not captured by separate EOF analyses. We briefly comment on this in the main text: “We compare the patterns obtained with MCA with those obtained with EOF analysis of Z200 and OLR fields (see Fig. S4 in the Supplementary Material). We find that the closest match of the Z200 MCA mode 1 pattern is with Z200 EOF 2 (spatial correlation ~ 0.8), while the closest match of Z200 MCA mode 2 is with EOF 1 (spatial correlation ~ 0.6). OLR MCA mode 1 has the closest match with EOF 2 (spatial correlation ~ 0.5), while OLR MCA mode 2 has the closest match with EOF 5 (spatial correlation ~ 0.4). Thus, in general our MCA patterns also reflect the first two EOFs of Z200 and OLR indicating that they explain an important fraction of the regional variability. Nevertheless, here we are interested in those patterns that can explain shared covariance, which cannot be achieved by using EOF analysis alone. Therefore, we use the MCA-defined patterns for the following part of the analysis.”

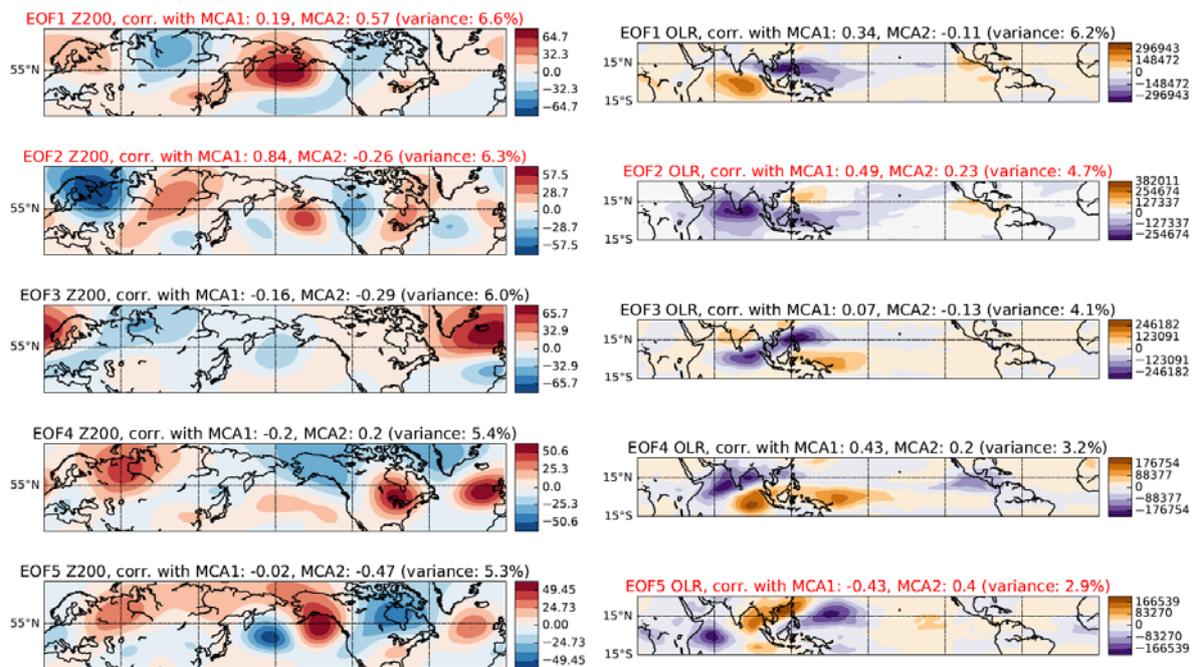


Figure R3. EOF analysis. The right column shows the first five EOF patterns for Z200, the left column shows the first five EOF patterns for OLR. In the title of each panel, the spatial correlation

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.

Apparently, many of the underlying assumptions are the same i.e stationarity etc It would help greatly if the authors could indicate if their approach is causal in the sense of Grainger given there appears to be no underlying stochastic model?

Our definition of causal graphs follows Pearl's causal Bayesian networks (Pearl 2000) and our approach to estimate these graphs from data comes from the constraint-based causal discovery framework (Spirtes 2000), here adapted to time series (Runge et al. 2019). In the constraint-based causal discovery framework, the existence (or absence) of causal relations is based on conditional independencies among subsets of the lagged variables together with a number of assumptions (as listed in our Discussion section). If Granger causality is only applied to pairs of variables, Granger causality does not account for common drivers or indirect links as is the case in our framework. Further, the constraint-based causal discovery framework *in general* goes beyond Granger causality since it can also account for contemporaneous causal links. Here we only focus on lagged links. If Granger causality is meant in a full multivariate setting, our approach is asymptotically equivalent to Granger causality, but for finite samples Granger causality has much lower detection power since it does not deal well with the curse of dimensionality as investigated in detail in Runge et al. (2019).

- 4. The analysis and attribution of the causal relationships is ultimately largely empirical, at times overly complicated and in some parts exceedingly verbose in description. The "causal maps" are very noisy and the reported relationships are very poorly represented from the patterns in the causal maps presented.**

We have taken the issue of noisiness raised by the anonymous reviewer very seriously, and combining this suggestion with the corresponding comment by anonymous reviewer #2, we have designed a robustness test that has removed much of the noise in the causal maps, greatly improving their visual appearance and interpretation. As a result, some of the more scattered regions that were described in the first version of the paper are now removed, and we can purely focus our description on the main, robust patterns. We describe this robustness in the revised manuscript: *"Finally, to test the robustness of our causal maps to the choice of time period, we calculate causal maps for a range of sub-periods. In 10 trials we removed 10% of the record (4 years). For ENSO-phase dependent causal maps, we have shorter time series and we thus remove one year in each trial, leaving a set of 14 causal maps for La Niña events and 13 causal maps for El Niño events. As a result, we obtain an ensemble of causal maps and apply the false discovery rate correction to p-values of each single map. Then, both for the full period (1979-2018) and for El Niño and La Niña years separately, we masked out areas where less than 70% of the trials indicated a significant causal link, giving an indication of the robustness of our findings and at the same time suppressing noise."*

This results in reduced noise in the new causal maps (see Fig R4-R6 in this document).

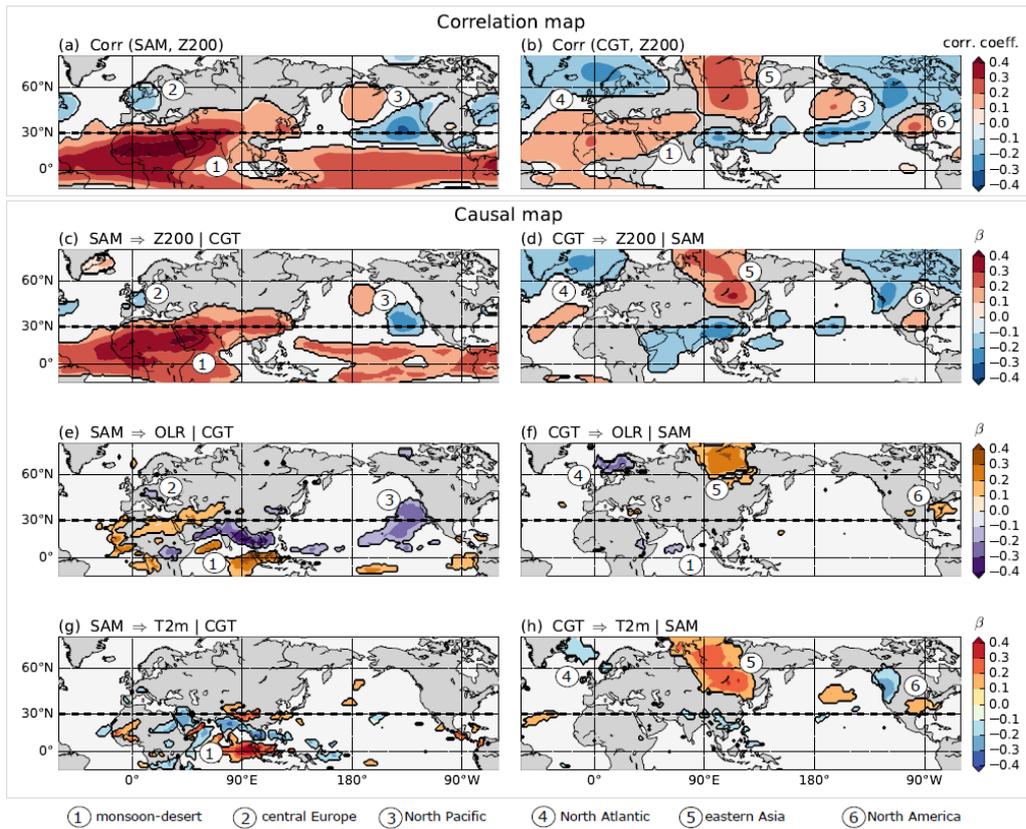


Figure R4. New Fig. 3.

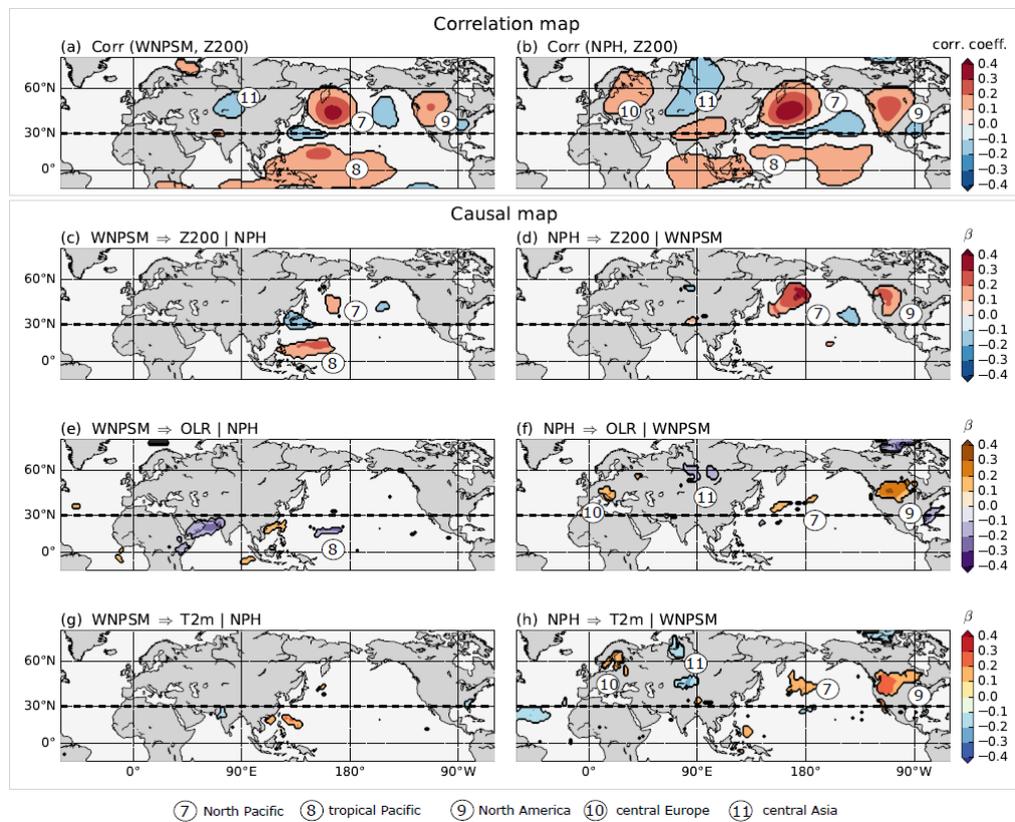


Figure R5. New Fig. 4.

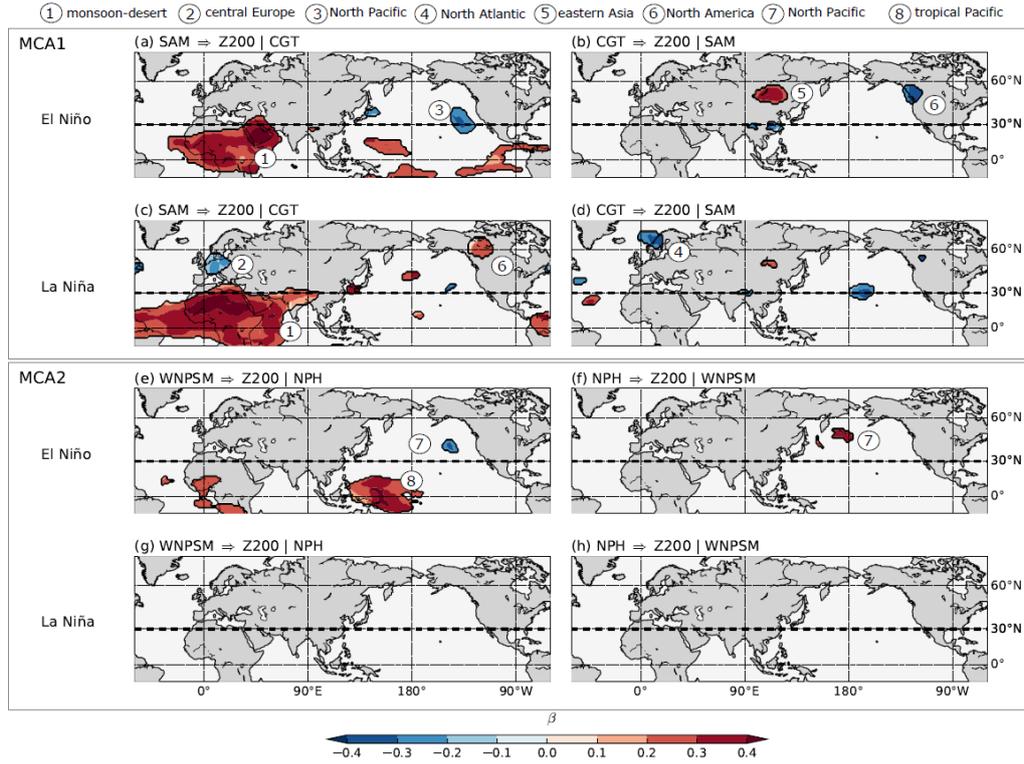


Figure R6. New Fig. 5.

It would greatly help the reader if the methodology was described in sufficient detail and better placed in context with other approaches, both in terms of dimension reduction and causal inference. This, in combination with a more concise discussion of the physical properties of the modes would allow the reader to better judge the merits of the approach.

In the revised manuscript, we have improved the methodology section by adding a concrete example showing how the PCMCI algorithm works (also following the comments by the second reviewer, see point 1 in our response to reviewer #2):

“In this analysis, A and B represent the two MCA scores obtained for a selected MCA mode, while $C(lon, lat)$ represents the grid point time series of a 2D field, e.g. T2m or Z200. In its first step, PCMCI iterates through partial correlations with increasing cardinality of conditions to remove the influence of common drivers and indirect links and estimate a preliminary set of parents. The first iteration of PC (cardinality 0) calculates the correlation between a selected time series, e.g. $A_{\tau=0}$, and the past of any other available time series, $\{A_{\tau=-1}, B_{\tau=-1}, C(lat, lon)_{\tau=-1}, \dots, A_{\tau=-\tau_{max}}, B_{\tau=-\tau_{max}}, C(lat, lon)_{\tau=-\tau_{max}}\}$, including its own past $A_{\tau=-1}, \dots, -\tau_{max}$. For illustration purposes, we here provide an example for $C(lon, lat)$, where ρ denotes the correlation and τ is the lag that is being used in the network (in this example, $\tau_{max} = -2$):

$$\begin{aligned} \rho(C(lon, lat)_{\tau=0}, A_{\tau=-1}) &= 0.32, pval = 0.01 & (5) \\ \rho(C(lon, lat)_{\tau=0}, A_{\tau=-2}) &= 0.13, pval = 0.1 \\ \rho(C(lon, lat)_{\tau=0}, B_{\tau=-1}) &= 0.35, pval = 0.005 \\ \rho(C(lon, lat)_{\tau=0}, B_{\tau=-2}) &= 0.23, pval = 0.058 \\ \rho(C(lon, lat)_{\tau=0}, C(lon, lat)_{\tau=-1}) &= 0.41, pval = 0.01 \\ \rho(C(lon, lat)_{\tau=0}, C(lon, lat)_{\tau=-2}) &= -0.16, pval = 0.06 \end{aligned}$$

Applying a significance level $\alpha = 0.05$, only three actors are significantly correlated with $C(lon, lat)$ at the chosen time lag. These form the initial preliminary set of parents for $C(lon, lat)$ and are ordered by the strength of their correlation:

$$P_{C(lon, lat)}^0 = \{C(lon, lat)_{\tau=-1}, B_{\tau=-1}, A_{\tau=-1}\} \quad (6)$$

Next, partial correlations between $C(lon, lat)$ and each actor in $P_{C(lon, lat)}^0$ are calculated by conditioning on the strongest preliminary parent:

$$\rho(C(lon, lat)_{\tau=0}, C(lon, lat)_{\tau=-1} | B_{\tau=-1}) = 0.35, pval = 0.02 \quad (7)$$

$$\rho(C(lon, lat)_{\tau=0}, B_{\tau=-1} | C(lon, lat)_{\tau=-1}) = 0.28, pval = 0.03$$

$$\rho(C(lon, lat)_{\tau=0}, A_{\tau=-1} | C(lon, lat)_{\tau=-1}) = 0.25, pval = 0.04$$

Parents with significant partial correlations will enter the second set of preliminary parents:

$$P_{C(lon, lat)}^1 = \{C(lon, lat)_{\tau=-1}, B_{\tau=-1}, A_{\tau=-1}\} \quad (8)$$

Next, the partial correlation is calculated conditioning on the two strongest parents:

$$\rho(C(lon, lat)_{\tau=0}, C(lon, lat)_{\tau=-1} | B_{\tau=-1}, A_{\tau=-1}) = 0.31, pval = 0.03 \quad (9)$$

$$\rho(C(lon, lat)_{\tau=0}, B_{\tau=-1} | C(lon, lat)_{\tau=-1}, A_{\tau=-1}) = 0.23, pval = 0.04$$

$$\rho(C(lon, lat)_{\tau=0}, A_{\tau=-1} | C(lon, lat)_{\tau=-1}, B_{\tau=-1}) = 0.12, pval = 0.08$$

Since it is not possible to further increase the dimension of the condition set, from the PC step, the preliminary parents converge to:

$$P_{C(lon, lat)}^2 = \{C(lon, lat)_{\tau=-1}, B_{\tau=-1}\} \quad (10)$$

By repeating this step for each variable, preliminary sets of parents are estimated. Let's assume that in our example we also obtain:

$$P_A^3 = \{B_{\tau=-1}, A_{\tau=-2}\} \quad (11)$$

$$P_B^2 = \{B_{\tau=-1}\}$$

In the MCI step, partial correlation is calculated again between each pair of actors (at different time lags) conditional on the above estimated sets of preliminary parents, whereby both sets of parents are conditioned upon. To give one example, this would lead to:

$$\begin{aligned} & \rho(C(lon, lat)_{\tau=0}, A_{\tau=-1} | P_{C(lon, lat)}^2, P_A^3) = \\ & = \rho(C(lon, lat)_{\tau=0}, A_{\tau=-1} | C(lon, lat)_{\tau=-1}, B_{\tau=-1}, B_{\tau=-2}, B_{\tau=-4}) = 0.1, pval = 0.3 \end{aligned} \quad (12)$$

Note that the parents of $A_{\tau=-1}$ are shifted in time by $\tau = -1$. After repeating (12) for each pair of actors shown in (5) and for time lags from 0 to τ_{max} , those parents that are significant in the MCI test will then form the final set of causal parents for each actor. See Runge et al. (2019a) for a more detailed discussion and explanation of the algorithm design and extensive numerical experiments."

Moreover, we have improved the physical interpretation of each mode in the results section: "The first MCA mode explains 18% of the squared covariance (squared covariance fraction, SCF) and shows a CGT-like wave-5 pattern in mid-latitude Z200. The Pearson correlation between the two time series of MCA scores for the first mode is $r \sim 0.5$. The spatial correlation with the weekly CGT pattern,

as defined by Ding and Wang 2005, is 0.52 (Fig. 2a). The CGT pattern also represents the second most important pattern in boreal summer mid-latitude circulation (Di Capua et al., 2020; Ding and Wang, 2005). This wave-5 pattern is linked to the South Asian monsoon (SAM) activity via its positive centre of action east of the Caspian Sea (see Fig. 2a). Applying MCA, we find that the CGT pattern co-varies with a band of enhanced tropical convective activity that extends from the Arabian Sea towards Southeast Asia, with a peak of convective activity over the Bay of Bengal (Fig. 2b) (Kang et al., 1999).. Using OLR composites and the Kikuchi Boreal Summer Intraseasonal Oscillation (BSISO) index, we explicitly show that the temporal evolution of the SAM convective activity as defined in Fig. 2b at weekly time-scales resembles the evolution of the BSISO (Goswami and Ajaya Mohan, 2001; Saha et al., 2012) (see Figs. S1-S2 and further discussion in the Supplementary Material). Therefore, we explicitly link the region of low OLR identified in Fig. 2b over the northern Indian Ocean and the Indian subcontinent to the SAM activity as described in the literature. Note that we name each MCA pattern after a characteristic regional feature, but the analysis is applied to the larger geographical domains as shown in Figure 2.

The second mode of co-variability explains a SCF of 14% and is characterized by a region of strong positive Z200 anomalies located at $\sim 45^\circ$ N, over the western North Pacific, directly to the west of the dateline (i.e. the most prominent centre of action of the mid-latitude wave). The Pearson correlation between the two time series of MCA scores for the first mode is $r \sim 0.6$. We will refer to this pattern as the North Pacific High (NPH) (Fig. 2c). The NPH is the summer counterpart of the North Pacific subtropical high, which characterizes boreal winter. During summer, this high pressure region is displaced northward by the start of the monsoon season in the western Pacific Ocean and replaces the Aleutian Low (Lu, 2001; Riyu, 2002). The NPH is associated with a region of enhanced convection over the sub-tropical western North Pacific, related to the western North Pacific summer monsoon (WNPSM) convective activity (Fig. 2d) (Li and Wang, 2005; Nitta, 1987; Wang et al., 2001). The WNPSM core domain extends from 110° - 160° E and 10° - 20° N, while the boundary with the ISM is located over the South China Sea (Murakami and Matsumoto, 1994). The WNPSM is characterized by a late sudden onset (end of July) and a peak in rainfall activity during August and September, which is different to the ISM that features an earlier onset (in June) and peak rainfall activity during July-August.”.