

Response to Reviewer #1 (Chris Brierly):

This piece of work applies a new methodology to look at the decadal variability to understand the response across an ensemble of idealised paleoclimate simulations. I do not see anything incorrect in the work. But I'm not sure that the current layout will result in many citations. This is because the manuscript jumps between trying to describe two things simultaneously: a new method and some scientific results. As the manuscript is submitted to GMD, I presume that the authors consider the new method to be the primary innovation and will address my comments and recommendations accordingly.

R1.1 *I see that there are three important facets of the methodology that are outside of the standard approach. Firstly, there is the fact that the EOFs are computed by looking across the whole ensemble – rather than separately within each single model. This is a nice touch and should be made clearer within the text (currently this is mentioned briefly on L123, but not stressed as key aspect of the methodology). I happen to have adopted this approach myself before (Chandler et al, 2024, <https://doi.org/10.1175/JCLI-D-23-0089.1>) to look at regional models. However, we were looking at the mean climate, not variability, and the primary modes were detected between the models. I was therefore surprised that your approach does not pick up any inter-model variations. I guess that this comes from the application of the decadal filtering, which is in effect removing the mean climate from the individual models. You ought to explain in more detail what kind of filter is being applied, and its implications. I suspect that you are using a band-pass filter, and that if you instead used a low-pass filter you find fundamentally different EOF patterns (more akin to the EPP, we describe in Chandler et al). The labelling of Fig 1 and 2 implies that they show EC-Earth's EOF patterns – rather than full ensemble patterns. Only the PCs, directed graph and transition timeseries relate to the particular EC-Earth simulation.*

We thank the reviewer for his comments. Indeed, our methodology uses a multi-model merged dataset to extract the EOFs and project each simulation's principal components. We find the cited paper that uses a similar approach very informative. As mentioned in the reviewer's comment, the decadal filter helps in maintaining these EOFs focused on the mean representation of decadal climate and not in particular anomalies that may dominate the tropical and South Atlantic variability in some simulations. We have added comments to clarify that we calculated the ensemble using a merged dataset in the methodology, figure labels, and where needed. Some paragraphs that address this comment added to the revised manuscript are shown below:

(line 252) **"To make sure we extract physically meaningful patterns, we create the merged data set's phase-space using the Tropical South Atlantic three leading EOFs (Figure 3, 'a', 'b' and 'c')"**

(line 142) **"Low frequency filters."**

To investigate decadal variability in the Atlantic Ocean, decadal filters were applied across all datasets, calculated as simple decadal means from the original monthly time series. These filters serve two primary purposes: first, to examine the decadal coupling between precipitation and SST variables; and second, to extract consistent multi-model patterns that serve as the foundation for a shared phase space. Decadal filtering is applied to the precipitation and SST datasets before pattern extraction (see Methods, section 2.2.2) to effectively smooth out fine-scale structures and disparities arising from different model architectures. Since we define a single phase space to encompass the PCs of the merged dataset, this filtering process is essential to ensure that the leading PCs capture substantial variance across the integrated multi-model ensemble.”

(line 135) “... we extracted two distinct phase spaces (one for SST and another for precipitation) composed of PCs derived from the combined simulations of all models and scenarios. This process yields a unified phase space with a shared spatial structure for the entire ensemble, providing a consistent framework for analyzing and comparing variability across different simulations. This procedure is discussed and applied by Chandler et al. (2024), who construct a merged dataset across models and extract dominant modes of variability using principal component analysis. Their approach implicitly defines a shared low-dimensional representation for inter-model comparison, which is conceptually related to our phase-space construction.”

Furthermore, alongside the standard procedure of normalizing each PC by its standard deviation, our methodology uses the concept of maximum entropy threshold, which works as a filter of the principal components' amplitude. That way, Shannon's Entropy becomes a more reliable quantitative measure for multi-model comparison. A paragraph extracted from the revised manuscript that addresses this follows:

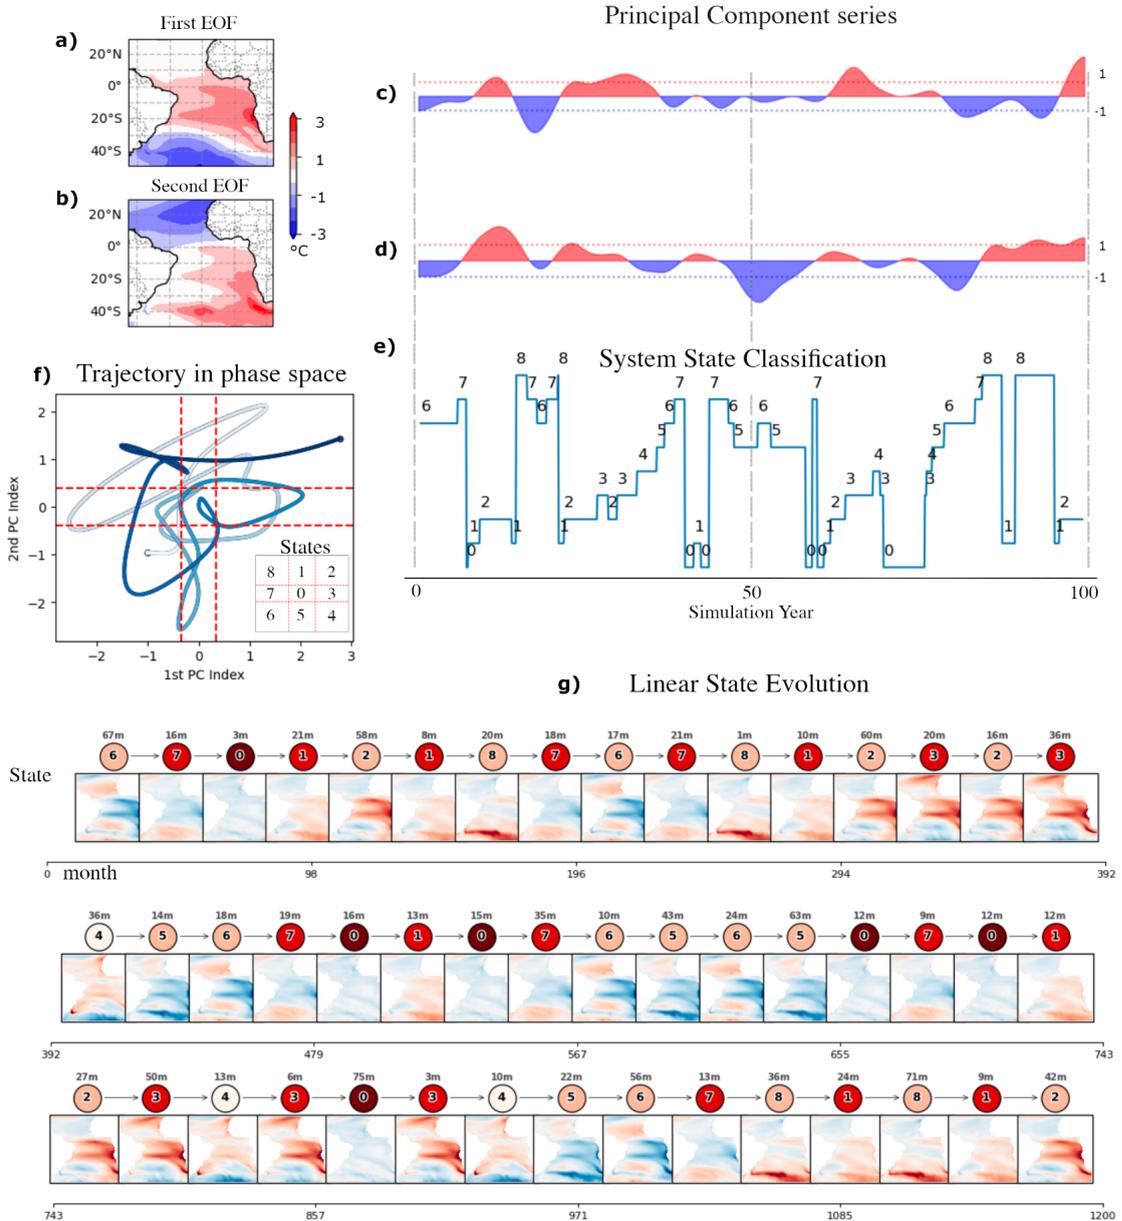
(line 233) “When we define a phase space using the leading PCs, we select a domain characterized by the patterns representing the system's greatest variance. By normalizing the indices and determining the threshold that maximizes Shannon Entropy, we effectively isolate temporal dynamics from the influence of amplitude variations. Our approach establishes a leveled ground for analyzing the temporal evolution of the system's state across its defining patterns, ensuring that the results remain independent of the specific amplitude differences simulated by various numerical models. In other words, although two models may reproduce the 1st PC with different amplitudes, they are both considered representations of the same climatic pattern.”

R1.2 *Your second innovation is the introduction of the directed diagrams. I confess that I find these hard to interpret, but I can see that they are important. Can you please spend a bit more time describing them? Perhaps thinking of them in a reduced dimension set would help; say by using the 2 ENSO modes of Ren & Jin (2011, <https://doi.org/10.1029/2010GL046031>) at then you can place the phase along the x and y axes. Can you also create some possible directed graphs for a system in which all the PCs are truly independent, through building a simple statistical model? These would then provide a suitable null-hypothesis and allow some statistical testing to be undertaken.*

The authors greatly appreciate the reviewers'suggestion. A new test example of the system's state evolution in a two-dimensional map was created using only the 2 main PCs of the Atlantic Ocean SST. Each axis represents one of the 2 dimensions from this simplified phase space, creating an obvious analogy to the direct graphs. The suggestion to shift the analysis from the Atlantic to the Pacific Ocean is also appreciated; however, we aim to retain the focus of this study on the Atlantic region. The new Methods section has the trajectory in phase space section updated as follows:

(line 163) "Projecting a simulation series onto the 'n' dimensions of our phase space generates a trajectory in that phase space. To elaborate more, we present a simplified 2-dimensional phase space for the SST decadal anomalies of the EC-Earth - PI experiment, using the two leading EOFs from the merged dataset. First, we apply the dimensionality reduction using the PC analysis to find the merged dataset's two leading EOFs (Figure 1 'a' and 'b'). Then, we project the EOFs into the simulation's Atlantic Ocean time series, generating its PCs (Figure 1 'c' and 'd'). The trajectory of our system in the 2-D continuous phase space is shown in Figure 1 'f'. Defining negative, neutral, and positive phases for each index, using a threshold linked to Shannon's Entropy (as described in Section 2.3), we coarse-grain this space into n states (9 possible states for this 2-D problem) and project our system on it (Figure 1 'e'). Finally, the trajectory can also be seen as a discrete system evolving in time, from one possible state to the next, as depicted in Figure 1 'g'.

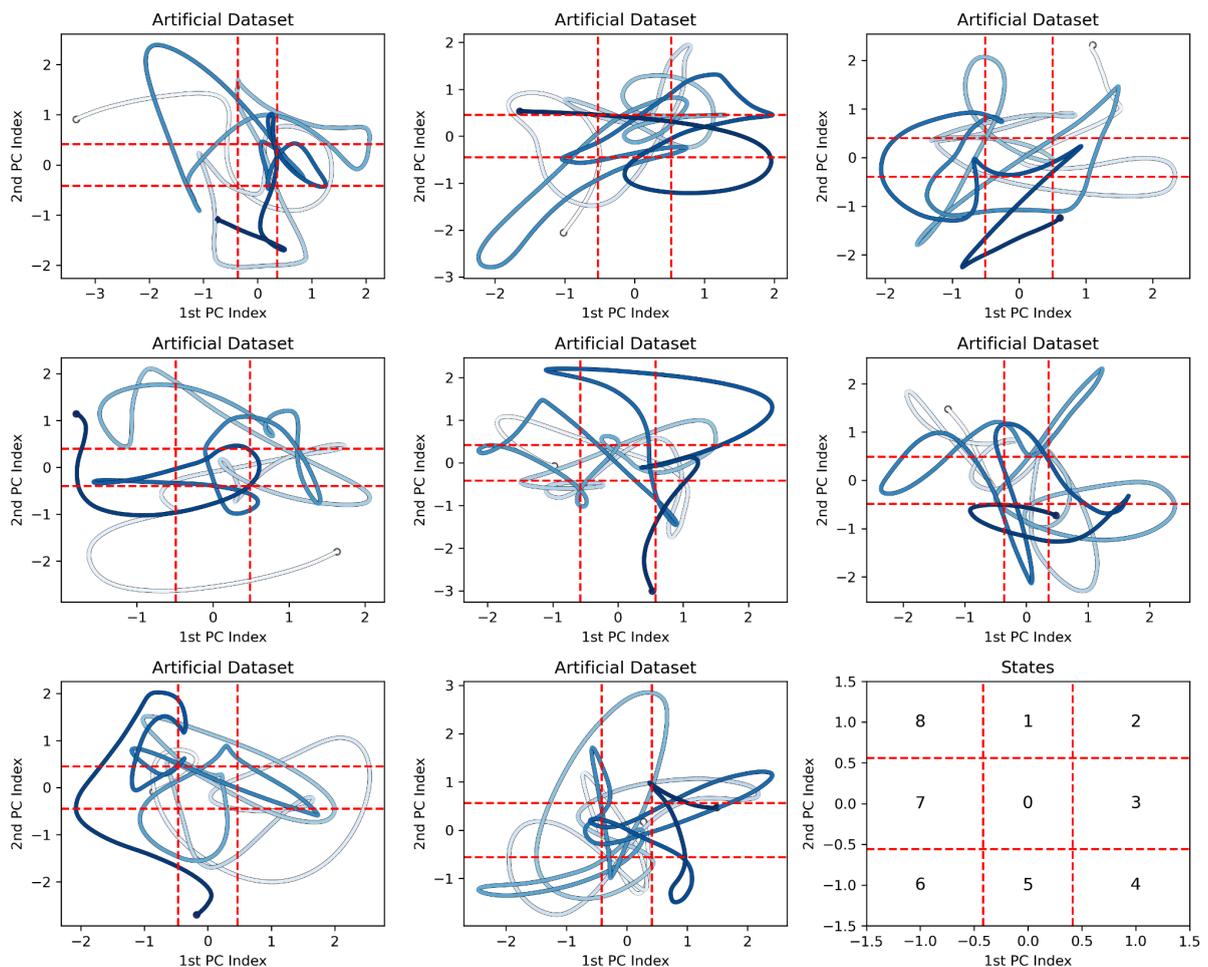
EC-Earth SST - PI experiment

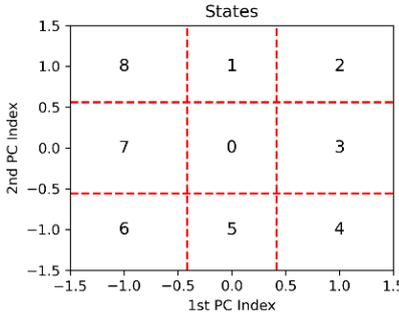
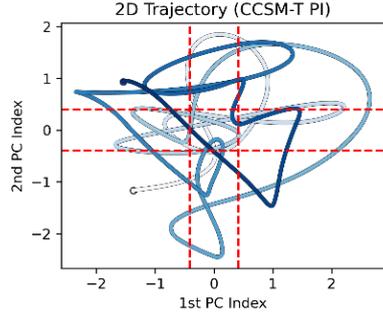
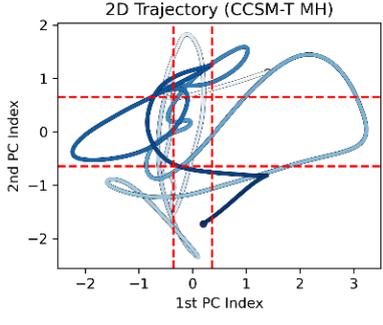
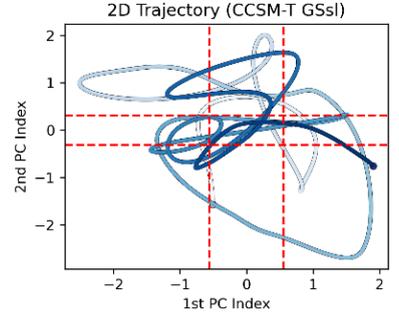
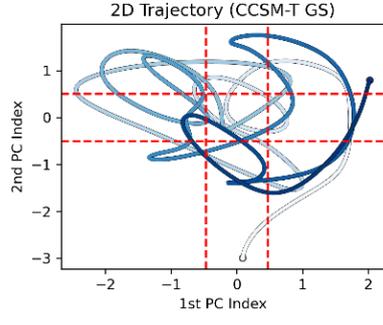
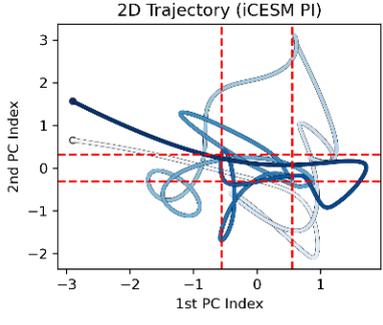
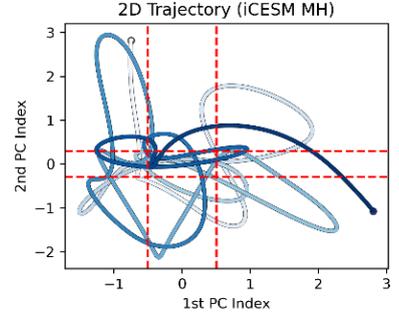
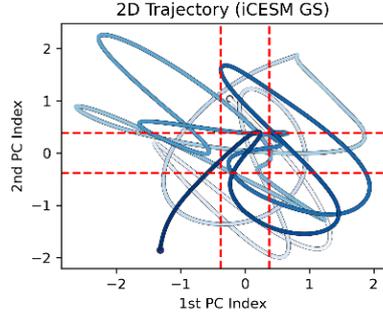
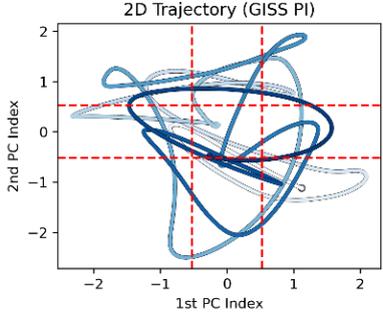
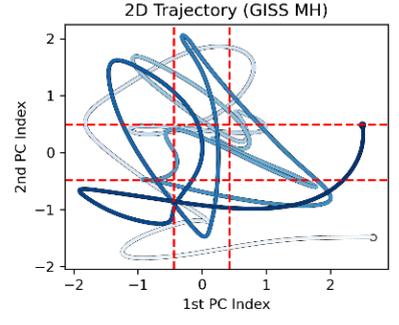
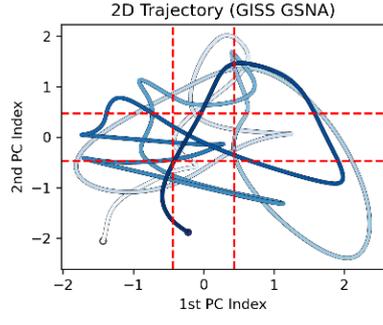
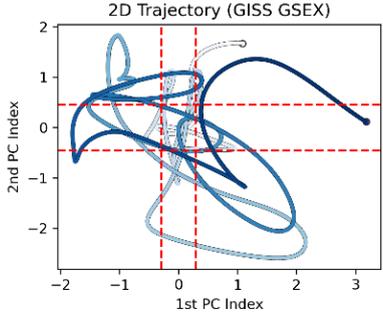
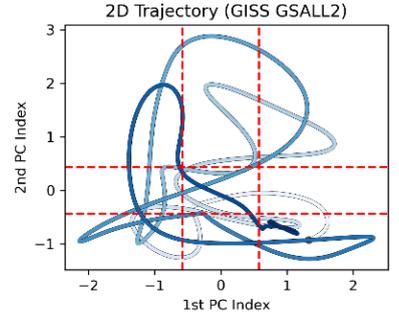
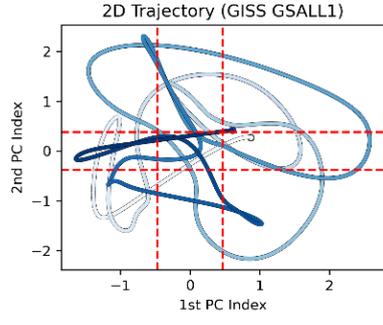
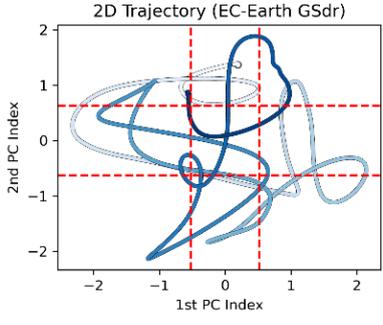
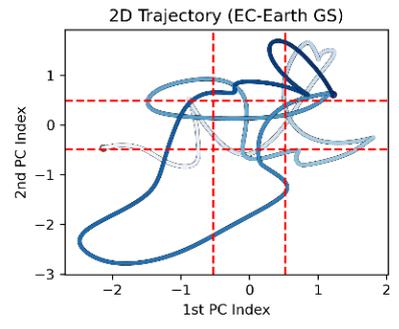
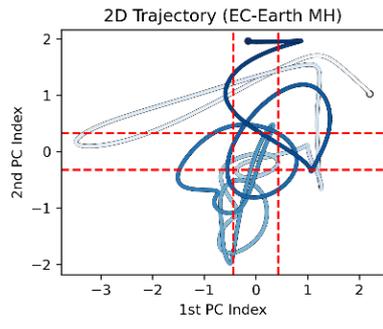
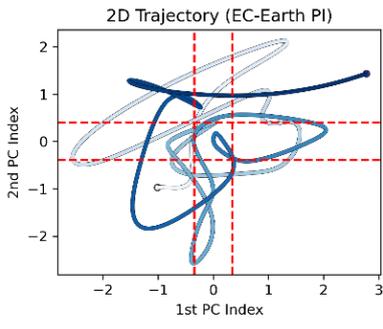


EC-Earth SST system state identification and evolution during the PI experiment. a) The first EOF; b) The second EOF; (c) and (d): their respective PC series, blue for negative and red for positive, dashed lines indicating the 1 standard deviation value. (e) The cluster identification of the discrete system state, given by the two PC indices above. (f) The system's trajectory in the continuous PC phase space. The line's color represents the time evolution, white for the start (month 0) and dark blue for its end (month 1200). (g) The diagram illustrates the state-by-state evolution of the system over time. Each node represents a specific SST pattern, identified by its cluster number (referenced in panels (e) and (f)) and its spatial pattern displayed below the node. The numeral positioned above each node indicates the duration, in months, that the system remained in that state before transitioning. Darker shading on a node signifies a higher frequency of transitions into that specific cluster throughout the entire time series.

Since the continuous trajectory in a high-dimensional phase space can be challenging to plot, the trajectories of a system in a higher-dimensional phase space can be depicted with directed graphs (Figure 2; 'h' in Figures 3 and 4; and Figures 5-8). These graphs can hold different information regarding the system's trajectory in phase space. Each node in these graphs represents a state of the system at a particular time step, with the node's size indicating the number of months the system remained in that state. Nodes that appear darker have a higher degree, meaning they are connected to more transitions to and from other states. This indicates that the system frequently returns to the same state, which results in a darker color for that node. The distance between nodes is irrelevant; it was adjusted for the graph. The information in a directed graph can be overwhelming to analyze for every simulation run; therefore, we utilize its information to calculate a macro-property that reflects its variability in phase space.

Furthermore, artificial sets were created with two independent random normally distributed indices. However, there was no notable difference between the evolution of these systems and their use of a 2-dimensional PC phase space. To satisfy the reviewer's curiosity, we present some of the artificial trajectories created from independent 2-D random PC Gaussian series alongside the 2-dimensional trajectories of the leading SST modes of the Atlantic Ocean in different simulations used in the manuscript. These plots will also be available in the GitHub repository, along with the code used to generate them.





R1.3 *Your third innovation involves the use of Shannon's entropy. You provide an equation for this and then describe some of its properties in the methods section. However, when showing the results, all the values are approximately 3. I didn't get why that should be so. I guess it's related to either the fact there are 3 EOFs or the fact that each PC is divided into 3 states, but I don't know which and you didn't explain.*

The entropy values are close to 3. This is due to the structure of our phase space: with three defined phases (positive, neutral, and negative) for each of the three principal components, the system has $3^3 = 27$ possible states. The maximum entropy (defined by the manuscript's equation 1) of a discrete space such as this would result in $\ln(27) \approx 3.296$. Each simulation's threshold is tuned to yield its maximum possible entropy, which naturally approaches $\ln(27)$. The following paragraph was added to the manuscript:

(line 289) "As seen in Figure 9 and Table 1, the entropy values are approximately 3. This is due to the structure of our phase space: with three defined phases (positive, neutral, and negative) for each of the three principal components, the system has $3^3 = 27$ possible states. The maximum entropy (equation 1) of a discrete space such as this would result in $\ln(27) \approx 3.296$. Since the simulation's thresholds in phase space are individually tuned to yield its maximum possible entropy, it naturally approaches $\ln(27)$."

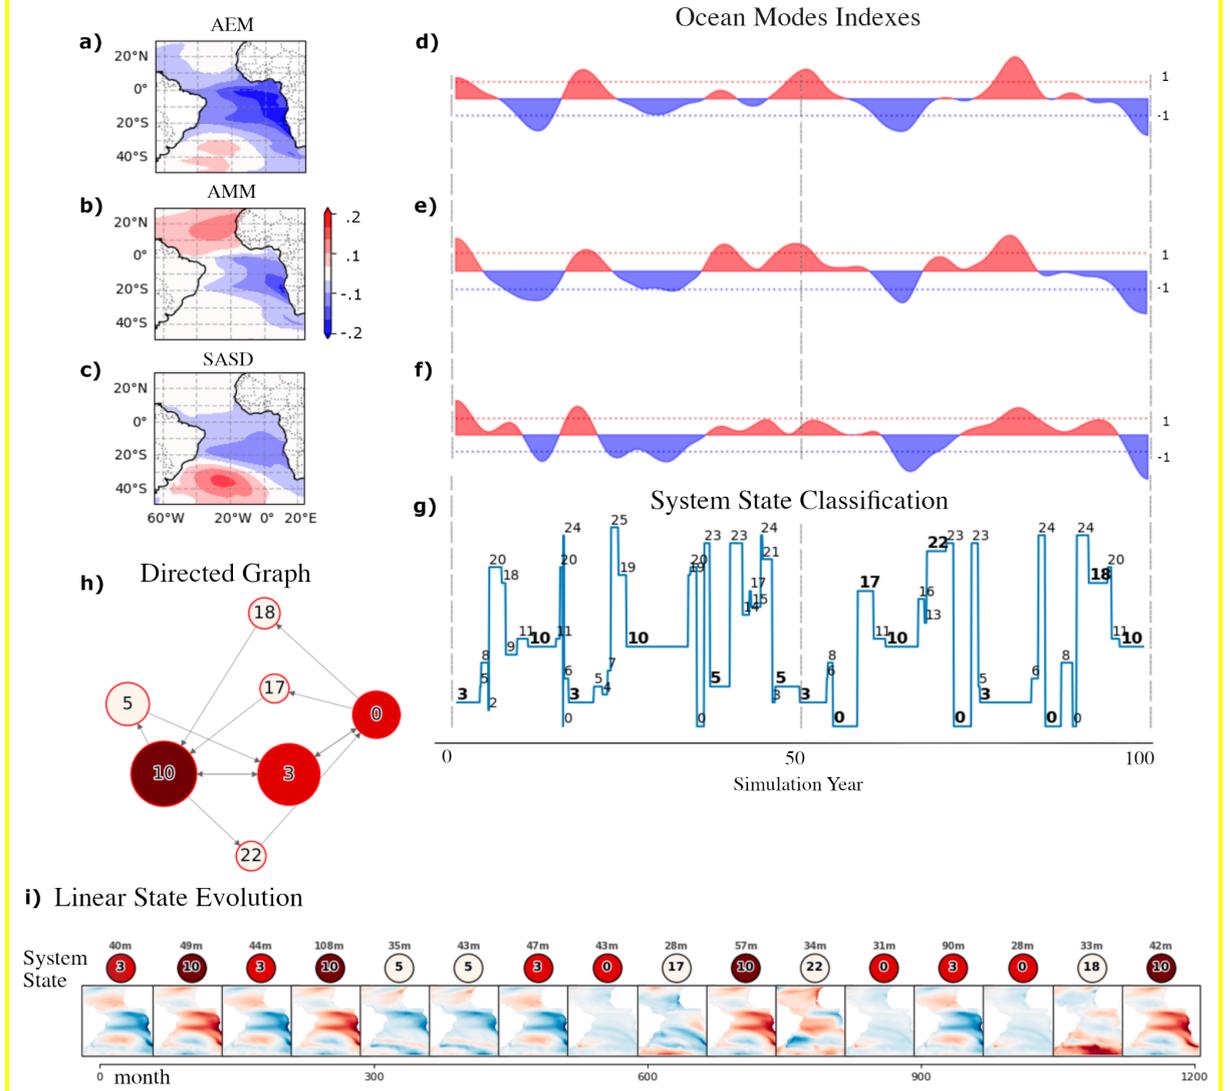
I have 3 other substantive comments:

R1.4 *Both the AMM and AZM/AEM are more often defined by SST indices, rather than through principal components. It would be informative to compare your new method against this style of definition.*

This revised manuscript's Result section uses the traditional AMM, SASD, and AEM modes to define an additional 3-dimensional phase space, yielding different comparable Entropy results. A new discussion of the PCs and Atlantic Ocean Modes approaches to defining the phase space has been added. To highlight some of the additions made in the rewritten Results section, the new figure in Results follows:

(line 315) "... An alternative approach to reducing Atlantic SST dimensions into a 3D phase space uses regional SST indices. These averages provide simple, physically interpretable metrics of climate mode strength. To test how phase space definitions affect Shannon's Entropy, we repeat our calculations using indices for the AMM, AEM, and SASD."

EC-Earth SST Modes - PI experiment



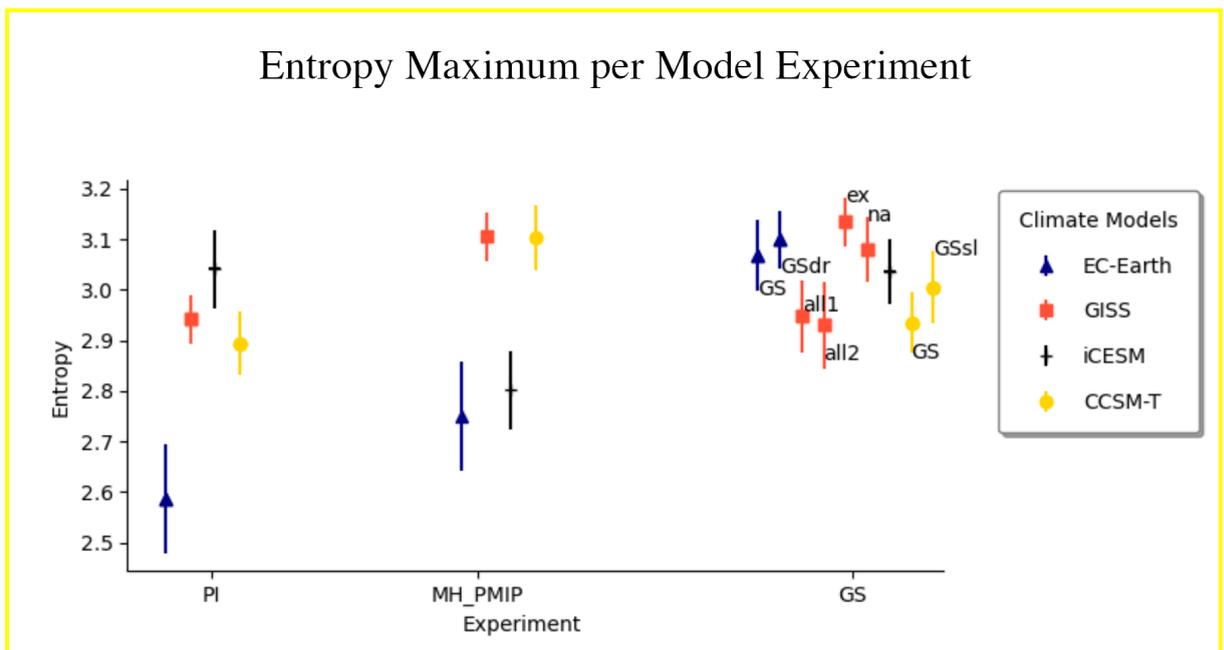
EC-Earth SST system state identification and evolution during the PI experiment. (a),(b), and (c): The merged ensemble AMM, AEM, and SASD, respectively; (d),(e), and (f): their respective index series, blue for negative and red for positive, dashed lines indicating the 1 standard deviation value. (g) The cluster identification of the discrete system state, given the three above Ocean Mode indices, bold font used in the most persistent states (lasting more than 24 months); (h) illustrates state evolution of the most persistent states. (i) tracks the state evolution in linear form. Each node displays its cluster number, spatial pattern, and duration in months before transitioning.

In the merged dataset, the first PC correlates strongly with the AEM (70%) and SASD (64%), while the second PC correlates with the AMM (71%). However, while the SST indices are significantly intercorrelated (39–43% in the merged dataset), the PCs maintain negligible correlation (below 3%). As an orthonormal basis, PCs inherently ensure minimal correlation between dimensions, unlike regional indices.

This index-based phase space yields generally lower entropy values than the PC-based space. While Shannon Entropy does not directly measure correlation, the dependencies between the regional indices significantly impact the system's total entropy.

Furthermore, this exercise helped debate the sensitivity of Shannon's Entropy to the phase space. We show the following paragraph from the revised manuscript to highlight some of the additions made:

(line 327) “



Maximum Entropy values and uncertainty for each model experiment for SST using the Atlantic Ocean modes phase space. From left to right: Pre-Industrial PI; mid-Holocene only orbital forcing (MH_PMIP); and mid-Holocene with Green Sahara boundary conditions (MH_GS) experiments. All models include PI, MH_PMIP, and MH_GS runs. Under Green Sahara conditions, additional experiments include: EC-Earth (ECE) with northern African vegetation and dust reduction (dr); GISS with both extratropical and northern African vegetation (GISSall1 and GISSall2), with only extratropical vegetation (GISSex), and with only north African vegetation (GISSna); CCSM-T with soil and lakes (sl).

EC-Earth

The EC-Earth PI experiment (Figure 10 - 'd', 'e' and 'f' and first line of Table 2) best illustrates how high correlation results in low entropy, showing the lowest entropy of all simulations within the Atlantic Ocean Modes phase space. In this specific setup, the AEM and AMM are correlated at 88%, while the AMM and SASD show a 76% correlation. Conversely, all MH scenarios exhibit higher entropy than the PI (6% for MH_PMIP and 19% for GS with/without dust) as these Atlantic modes decouple.

GISS

Compared with the PC phase space, GISS \$PI\$ entropy reduces using the Atlantic ocean modes indices (Table 2). The entropy from the remaining MH scenarios mostly decreases; however, they show no significant changes to the entropies calculated in the PCs phase space.

iCESM

For this model, the SST entropies show a general decrease, while holding the same internal biases seen in the PC phase space entropies (Table 2).

CCSM-Toronto

The only significant difference is still between PI and MH_PMIP; however, in this phase space, MH_PMIP holds the largest entropy amongst this model's different experiments (7% higher than PI)."

R1.5 There is no comparison to observations contained within your analysis. Can you project the resultant EOFs onto a reanalysis dataset and look at the real world?

This revised manuscript's Result section uses observational data to compare the model's PCs and SST Atlantic Ocean modes to the satellite era precipitation and SST patterns.

(line 242) "We test our methodology in three steps. First, we analyze the outputs of four different numerical climate models in PI and various MH climate scenarios (see Methods). We extract the tropical and South Atlantic SST and precipitation three leading Principal Components (PCs), calculate their trajectory in this 3-dimensional phase space, and compute their respective Shannon's Entropy. In the second, we explore the classical Oceanography SST-based indexes to create the Tropical and South Atlantic Ocean's SST phase space (see Methods), calculate its trajectory, and Shannon's Entropy once again. The third step is a comparison between the calculated Entropy values from the model simulations and observational data from satellites using the previously discussed phase spaces."

(line 343) "A primary goal of numerical climate models is the faithful reproduction of low-frequency decadal variability. Notable differences exist between the PCs extracted from observational data and model data, nor do they represent the same variance as the principal components from the leading components seen in the observational satellite era (Figure S1). However, the EOFs extracted from the merged simulations dataset can be projected onto the observational data time series. As long as they have the same length, the use of maximum Entropy enables a comparison between the observational data and the simulations' trajectories in a common phase space. Since the observational datasets range from 1980 to 2015, the results shown in the Table ahead are the Entropy values of the first 420 months (35 years) of each simulation. Within these time windows, the resulting entropy values are lower, indicating that shorter intervals do not fully capture the decadal variability envelope of the tropical and South Atlantic system."

Table 3: Entropy mean and Standard Deviation from the 420-month-long series model runs and the observation data set

Model	Scenario	Entropies in PCs' phase space		Entropies in Modes' phase space
		H_{sst}	H_{ppt}	H_{sst}
ECE	<i>PI</i>	2.60 ± 0.11	2.70 ± 0.08	2.26 ± 0.13
ECE	<i>MH_{PMIP}</i>	2.79 ± 0.09	2.91 ± 0.13	2.45 ± 0.10
ECE	<i>MH_{GS}</i>	2.63 ± 0.07	2.58 ± 0.09	2.80 ± 0.11
ECE	<i>MH_{GSdr}</i>	2.62 ± 0.09	2.80 ± 0.08	2.67 ± 0.09
GISS	<i>PI</i>	2.73 ± 0.08	2.95 ± 0.08	2.76 ± 0.09
GISS	<i>MH_{PMIP}</i>	2.75 ± 0.08	2.63 ± 0.10	2.81 ± 0.09
GISS	<i>MH_{GSall1}</i>	2.83 ± 0.10	2.78 ± 0.14	2.64 ± 0.10
GISS	<i>MH_{GSall2}</i>	2.85 ± 0.09	2.68 ± 0.10	2.62 ± 0.11
GISS	<i>MH_{GSex}</i>	2.95 ± 0.12	2.74 ± 0.10	3.04 ± 0.09
GISS	<i>MH_{GSna}</i>	2.80 ± 0.10	2.81 ± 0.11	2.72 ± 0.10
iCESM	<i>PI</i>	2.71 ± 0.10	2.72 ± 0.10	2.77 ± 0.09
iCESM	<i>MH_{PMIP}</i>	2.85 ± 0.09	2.79 ± 0.08	2.74 ± 0.12
iCESM	<i>MH_{GS}</i>	2.90 ± 0.11	2.71 ± 0.10	2.84 ± 0.09
CCSM-T	<i>PI</i>	2.90 ± 0.09	2.56 ± 0.17	2.72 ± 0.11
CCSM-T	<i>MH_{PMIP}</i>	2.57 ± 0.10	2.70 ± 0.07	2.57 ± 0.10
CCSM-T	<i>MH_{GS}</i>	2.57 ± 0.07	2.92 ± 0.13	2.73 ± 0.12
CCSM-T	<i>MH_{GSsl}</i>	2.65 ± 0.09	2.76 ± 0.10	2.83 ± 0.08
Observation	1980 – 2015	2.83 ± 0.09	2.78 ± 0.09	2.40 ± 0.13

”

Spatial Patterns and Variance Explained

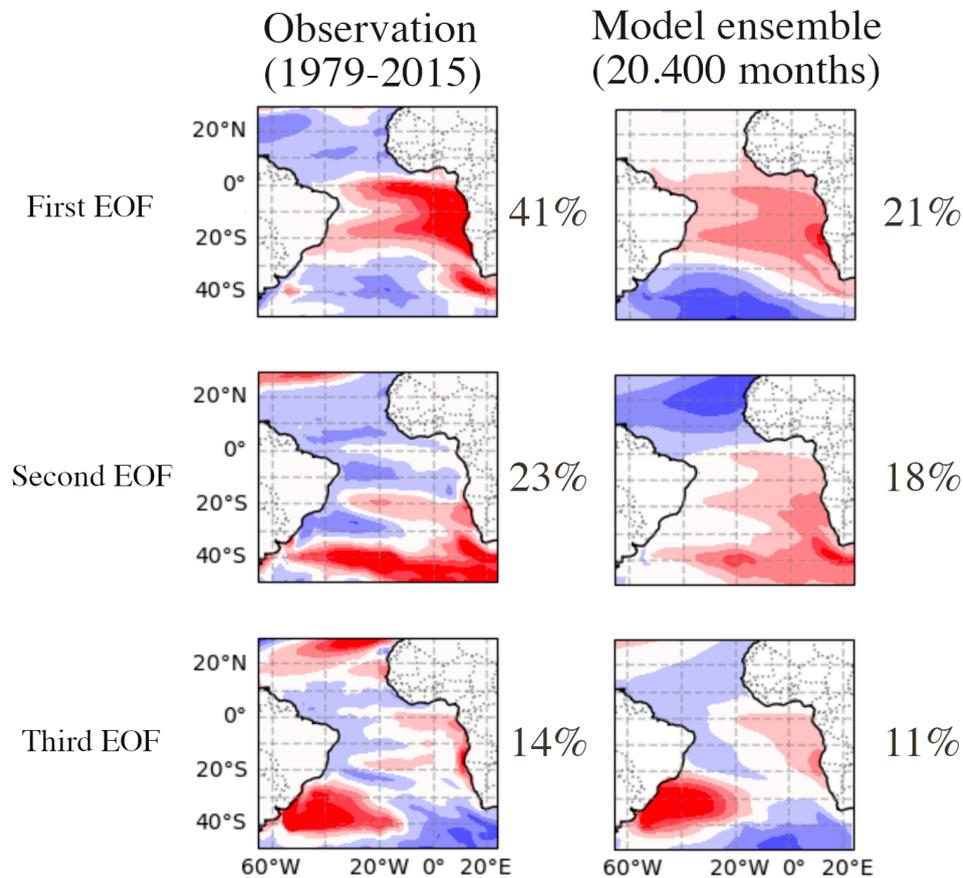


Figure S1. The three leading empirical orthogonal functions (EOFs) from the principal component analysis. From left to right, the observational EOFs, the merged model ensemble EOFs, and the percentage of explained variance in their respective dataset. From top to bottom, the first, second, and third leading modes.”

Furthermore, we do not wish to overstate results about the mid-Holocene experiments using the observational dataset alone; when using numerical climate models, it is wisest to compare satellite era data with historical experiments. However, we try to use these results to compare the mean Entropy from the multi-model ensemble in each phase space, discussing what the observation dataset's entropy indicates is the most reliable procedure (the PC or the regional SST boxes). The discussion made in the manuscript follows:

(line 352) “Typically, satellite-era observations are compared against historical model runs. Previous research categorized the tropical and South Atlantic systems into a discrete evolution of patterns, tracking the SST and precipitation cycles using reanalysis data from 1890-2015 with a directed graph format (Gorenstein et al., 2023). However, as an attempt to quantitatively compare the values emerging from the numerical climate models' entropy in these MH experiments, we calculate the entropy from two observational datasets (the HadISST and the GPCP datasets - see Methods).

In the PC phase space, the ensemble model mean (2.76 for SST and 2.75 for precipitation) falls within the observational entropy uncertainty interval. However, in the Atlantic Ocean Modes phase space, the ensemble model mean (2.70) is significantly higher than the observational entropy value (2.40 ± 0.13). While this could suggest that historical variability is lower than most of the PI and MH experiments presented, such a direct comparison assumes the models represent these modes accurately. Instead, this discrepancy may indicate that, for this methodology, studying Atlantic decadal variability in simulations is more reliable when using principal components rather than traditional SST indices.”

Finally, these changes are also reflected in the conclusion section as follows:

(line 417) “Across four models, the results show that mid-Holocene forcings can significantly alter the degree of organization of Atlantic decadal variability, with model-dependent entropy changes, and different SST and precipitation responses. Comparisons with observational datasets indicate that PC-based phase spaces provide a more consistent and robust basis for evaluating low-frequency Atlantic variability than traditional SST indices, highlighting entropy as a powerful metric to diagnose how external forcings reshape the structure, persistence, and transitions of dominant climate modes.”

(line 437) “This methodology characterizes system dynamics and variability in terms of state occupancy and transitions rather than absolute anomaly magnitudes. By discretizing the climate system in a unified low-dimensional phase space and computing Shannon’s entropy, the underlying climate variability can be directly compared between different models, experiments, and observations.”

R1.6 You have chosen a case study (idealised green Sahara experiments during the mid-Holocene) that is not readily accessible to the general reader, and therefore requires a substantial amount of additional introduction and methods that get in the way of the methodological innovations you are trying to explain.

We agree that other types of data or model experiments could be equally interesting to investigate using our methodology. In the previous version of the manuscript, the discussion focused too heavily on the results derived from the chosen dataset, which may have obscured a more direct presentation of the proposed methodology. To address the reviewers’ concerns, the revised manuscript has been substantially rewritten to emphasize the methodological framework rather than the specific results of the Green Sahara. In particular, the Introduction and Discussion sections were revised to remove mid-Holocene-specific details that may not be of broad interest, while retaining this period as a motivating example for applying a novel methodology to study climate variability in the tropical and South Atlantic region. The revised paragraphs in the Introduction section now read as follows:

(line 35) “The mid-Holocene (MH) period, approximately 5000-7000 years Before Present, provides a natural framework to investigate how external forcings reshape climate variability. Orbital changes during this period altered the seasonal and latitudinal distribution of insolation, leading to profound hydroclimatic changes, most notably the African Humid Period and the expansion of vegetation over the Sahara. Proxy studies and climate model simulations suggest that these boundary-condition changes affected not only mean climate states but also the dynamics of tropical and Atlantic variability through

land–atmosphere–ocean feedbacks. Rather than simply amplifying or damping anomalies, quantifying how such forcings reorganize decadal variability remains an open problem, particularly when comparing across models with differing physics and parameterizations.

The geophysical mechanisms behind the Atlantic Ocean modes coupling with pressure and wind driving decadal precipitation anomalies were unraveled using observational data in Gorenstein et al. (2023). When examining the dynamics of this region in climate simulations, a more fundamental question arose regarding how to assess its decadal variability. Since numerical models present biased climate representations compared to observational data and among themselves, their climate variability is not typically measured in relation to large-scale climate patterns, such as ocean modes. This motivated us to create a new method using ocean modes and their precipitation counterparts to quantify decadal variability in numerical climate models.

Furthermore, the Discussion section is used to explore the broader applicability of the proposed methodology. There, we draw on previous proxy and numerical modeling studies to discuss how our information-theoretic definition of climate variability can be compared with traditional variance-based metrics used in the literature. In this context, it is necessary to address mid-Holocene-specific features and to assess whether our results are consistent with earlier findings for this period. While we have aimed to keep the manuscript direct and concise, removing this discussion would limit the paper's ability to evaluate the methodology in relation to existing studies and to illustrate its potential applications.

Other comments:

R1.7 You need to label each panel of a figure rather than collect them together (so Fig 1 has panel a-h, not just a-b).

The figures have been changed according to the reviewer's suggestions.

R1.8 Can you make the system-state classification centre around the neutral state of the 3 modes? This is the most intuitive way of thinking about them, but I don't see where that state is - presumably it is the most common mode in the directed-graph (but not necessarily).

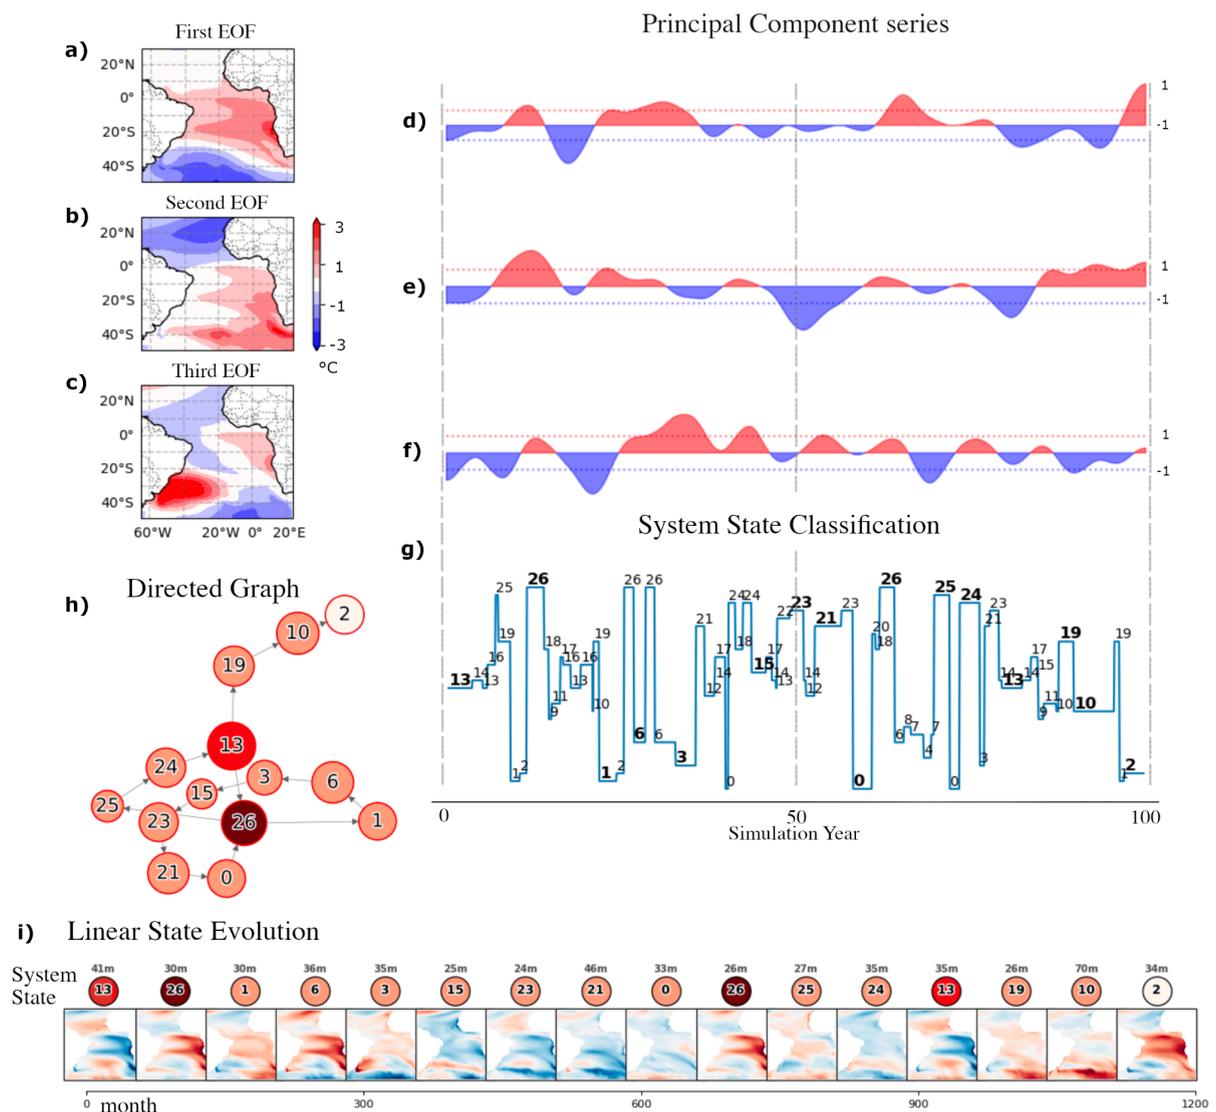
The direct graphs can be controlled until a certain point. For example, it could be interesting to fix the most visited state in the center of the plot and the remaining states accordingly. However, in the revised version of the manuscript, state zero does not seem to be the most visited in the direct graphs, nor is there an evident commonality in the ensemble of simulations. We believe the previous version of the manuscript had Figures and graphs around state zero, as they were plotted using system discretizations with a fixed threshold (thresholds of 1 standard deviation) instead of the maximum entropy threshold.

Furthermore, the networkx library used to plot them is not very adaptable, and forcing a design can compromise the final graph visualization. It is not simple to represent densely connected graphs, maintaining connected nodes close to each other, and neither nodes nor edges overlapping. The graphs presented in the manuscript are a product of extensive testing using Python's library and are found to be feasible representations of the simulation trajectories.

To improve the graph readability, we introduced two new concepts in all the figures that contain graphs (Figures 1, 3-8, and 10). First, the most persistent states concept: the only states used to plot the graphs are those that show a minimum persistence of months; this way, the graphs become more focused and compact. Second, a frame-by-frame depiction of the most persistent system state's evolution, presenting the same information contained in the directed graphs in a linear format. For example, the new Figure 3 follows:

(line 265) “

EC-Earth SST - PI experiment



EC-Earth SST state identification and evolution during the PI experiment. Panels (a)–(c) show the merged ensemble’s first three EOF patterns, with (d)–(f) displaying their corresponding PC series (red/blue indicating positive/negative phases; dashed lines showing 1 standard deviation). (g) identifies discrete system states based on these PC indices, bold font used in the most persistent states (lasting more than 24 months); (h) illustrates state evolution of the most persistent states. (i) tracks the state evolution in linear form. Each node displays its cluster number, spatial pattern, and duration in months before transitioning.”