

The Predictability of Tropical Pacific Decadal Variability: Insights from Attractor Reconstruction

NANDINI RAMESH

Department of Earth and Environmental Sciences, Columbia University, New York, New York

MARK A. CANE

Lamont-Doherty Earth Observatory, Palisades, New York

(Manuscript received 13 April 2018, in final form 20 November 2018)

ABSTRACT

Tropical Pacific decadal variability (TPDV), though not the totality of Pacific decadal variability, has wide-ranging climatic impacts. It is currently unclear whether this phenomenon is predictable. In this study, we reconstruct the attractor of the tropical Pacific system in long, unforced simulations from an intermediate-complexity model, two general circulation models (GCMs), and the observations with the aim of assessing the predictability of TPDV in these systems. We find that in the intermediate-complexity model, positive (high variance, El Niño-like) and negative (low variance, La Niña-like) phases of TPDV emerge as a pair of regime-like states. The observed system bears resemblance to this behavior, as does one GCM, while the other GCM does not display this structure. However, these last three time series are too short to confidently characterize the full distribution of interdecadal variability. The intermediate-complexity model is shown to lie in highly predictable parts of its attractor 37% of the time, during which most transitions between TPDV regimes occur. The similarities between the observations and this system suggest that the tropical Pacific may be somewhat predictable on interdecadal time scales.

1. Introduction

The coupled ocean–atmosphere system of the Pacific basin exhibits variability on interdecadal time scales that has impacts on hydroclimate around the world (Herweijer and Seager 2008; Power et al. 1999; Krishnan and Sugi 2003; Mantua and Hare 2002), on global mean surface temperature (Kosaka and Xie 2013; Meehl et al. 2016a; England et al. 2014), and on the ecosystems and fisheries of the Pacific Ocean (Francis et al. 1998). This variability is frequently measured using indices such as the Pacific decadal oscillation (PDO; Mantua and Hare 2002), which is based on temperature anomalies in the North Pacific, or the interdecadal Pacific oscillation (IPO), which additionally incorporates the influence and variability of the southern midlatitudes (Henley et al. 2015). The behavior captured by these indices is composed of contributions from a number of phenomena including the variability of the Aleutian low, air–sea heat flux anomalies in the midlatitudes amplified by

the seasonal cycle, the dynamics of subtropical ocean gyres, internally generated stochastic variability of the atmosphere, and the variability of the tropical Pacific (Newman et al. 2016).

In observations and simulations, the tropical component of this interdecadal variability takes the form of shifts between periods of El Niño-like (positive) states and La Niña-like (negative) states and is among the dominant influences on these indices (Zhang et al. 1997; Newman et al. 2003; Chen and Wallace 2015; Newman et al. 2016; An et al. 2007), accounting for up to 50% of the variance of the PDO. These phases may also be characterized as high- and low-variance states (Atwood et al. 2017; Choi et al. 2009), since interdecadal variance is highly correlated with mean tropical temperature. This variability intrinsic to the tropical Pacific region on decadal-to-interdecadal time scales, which we will refer to as tropical Pacific decadal variability (TPDV), will be the focus of this study.

The pertinence of interdecadal climate variability to the design of policies and infrastructure (Trenberth 2009) makes TPDV a desirable target for prediction. However, the processes responsible for generating this

Corresponding author: Nandini Ramesh, nramesh@ldeo.columbia.edu

phenomenon are not well understood. Some hypotheses place the origins of this variability in the midlatitudes, where anomalies are able to exert an influence on the tropical wind field via air–sea fluxes (Barnett et al. 1999; Vimont 2005). Others have suggested that phenomena external to the Pacific basin, such as variability in the Atlantic Ocean (Dong et al. 2006; Kang et al. 2014) or volcanic aerosols (Adams et al. 2003) are able to induce TPDV as a response. However, multiple studies using varied techniques have demonstrated that the strongest influence on the equatorial thermocline of the Pacific on interdecadal time scales is wind forcing within the tropical belt (McGregor et al. 2007; Karspeck and Cane 2002; Emile-Geay and Cane 2009), implying that the dynamics of the TPDV are governed by the wind stresses and oceanic planetary waves originating within this region; while others have found that the TPDV can be generated by stochastic atmospheric variability intrinsic to the region (Chang et al. 1996; Flügel and Chang 1999; Thompson et al. 2001; Okumura 2013). In addition, a model simulating only the tropical Pacific region has been shown to capture longer-than-interannual variability with a similar level of skill to a general circulation model (GCM; Ramesh et al. 2017), suggesting that factors external to the tropical Pacific region may not be necessary for the generation of TPDV.

Each of the proposed mechanisms has different implications for the predictability characteristics of the system. Based on the existing hypotheses, TPDV could be entirely the product of chance, driven by stochastic processes that render it effectively unpredictable (Kessler 2002; Kleeman 2008; Chang et al. 1996; Wittenberg 2009; Wittenberg et al. 2014), or could arise from low-order chaos in the coupled ocean–atmosphere system of the tropical Pacific causing it to oscillate between different regimes (Timmermann and Jin 2002; Tziperman et al. 1994; Timmermann 2003), allowing for some, if modest, predictability. Hypotheses that place the origins of TPDV outside the tropical Pacific (Dong et al. 2006; Barnett et al. 1999; Kang et al. 2014; Adams et al. 2003) imply that its predictability depends on that of the external factors involved.

The real-world system and climate models lie somewhere on the spectrum ranging from noise driven to chaotic. Recent studies indicate that the interdecadal variability of the Pacific basin possesses some predictability in GCMs (Meehl et al. 2016b; Thoma et al. 2015; Ding et al. 2013), suggesting that it is not entirely stochastic in these models. In an intermediate-complexity model of the tropical Pacific, the 15-yr mean of the Niño-3 index was found to possess potential predictability (Karspeck et al. 2004; Ramesh et al. 2017), and an actual prediction of the future state of the TPDV using

this model in 2004 (Seager et al. 2004) was verified a decade later.

In this study, we further examine the potential predictability of the internally generated TPDV identified in Karspeck et al. (2004) and Ramesh et al. (2017) on the 15-yr time scale by applying an attractor reconstruction scheme to the output from this model and comparing the results with those from two GCMs and the observations. While the real-world and GCM versions of TPDV are influenced by a variety of factors, the TPDV in the intermediate-complexity model is generated purely by the dynamics of the tropical Pacific region, allowing us to study the region as an isolated system.

The reconstruction of the attractor allows us to classify physical states of the models according to their predictability, enabling us to study the characteristics of the states of these systems from which we can hope to make successful predictions of the future. We use the intermediate-complexity Zebiak–Cane (ZC) model, which has been found to possess the properties of a chaotic system (Tziperman et al. 1997); the Geophysical Fluid Dynamics Laboratory (GFDL) Coupled Model, version 2.1, which has been shown to generate a noise-driven TPDV (Wittenberg et al. 2014); and the Community Climate System Model, version 4 (CCSM4), which has neither been proven to be noise driven nor chaotic but is known to possess some predictive skill in the Pacific basin on this time scale (Meehl et al. 2016b). We compare the reconstructed attractor of the observed variability (which is reconstructed with far less confidence than for the models because of the shortness of the observational record) with each of these systems in order to gain insight into whether chaotic or noise-driven processes dominate the observed TPDV.

We first describe the models and data used in section 2, followed by the attractor reconstruction method in section 3. In section 4a, we begin by examining the reconstructed attractor of the ZC model in detail and compare it with those of the other datasets used. We organize the discussion around the ZC model because the long time series available (100 000 years) allows the attractor to be reconstructed with much higher confidence than for any of the other datasets. We then continue the comparison in terms of the distributions of 15-yr standard deviations of the models used in section 4b. Next, in section 4c, we quantify the predictability obtained from the reconstructed attractor of the ZC model. Section 4d confirms that this attractor reconstruction succeeds in placing similar physical states of the system close together, as is expected, and demonstrates the possibility of using the attractor to understand the physical dynamics of the system during transitions between states of the TPDV. The final section summarizes

TABLE 1. List of CMIP5 models used for comparison in this study.

Modeling center (or group)	Institute ID	Model name
Beijing Climate Center, China Meteorological Administration	BCC	BCC_CSM1.1
National Center for Atmospheric Research	NCAR	CCSM4
LASG, Institute of Atmospheric Physics, Chinese Academy of Sciences	LASG/IAP	FGOALS-gl
NASA Goddard Institute for Space Studies	NASA GISS	GISS-E2-R
Met Office Hadley Centre	MOHC	HadCM3
L'Institut Pierre-Simon Laplace	IPSL	IPSL-CM5A-LR
Japan Agency for Marine-Earth Science and Technology, Atmosphere and Ocean Research Institute (The University of Tokyo), and National Institute for Environmental Studies	MIROC	MIROC-ESM
Max-Planck-Institut für Meteorologie Meteorological Research Institute	MPI MRI	MPI-ESM-P MRI-CGCM3

our results, discusses the implications for the predictability of GCMs and the real world, and outlines future research directions building on this work.

2. Models and data

We use the following datasets:

- (i) 100 000 years from an unforced run of the ZC (Zebiak and Cane 1987) model: This model simulates the Pacific as a 1.5-layer ocean from 29°S to 29°N with a global atmosphere based on the Gill model. It produces interdecadal variability (Cane et al. 1995) that is somewhat predictable and may have some utility for real-world predictions (Karspeck et al. 2004; Ramesh et al. 2017; Seager et al. 2004).
- (ii) 4000 years from an unforced run of the GFDL Coupled Model, version 2.1 (Delworth et al. 2006): tropical Pacific variability in this fully coupled GCM has been studied extensively (Karamperidou et al. 2014; Atwood et al. 2017; Wittenberg et al. 2014).
- (iii) 1000 years from a preindustrial control run of the CCSM4, from the CMIP5 archive (Taylor et al. 2012): IPO in this state-of-the-art GCM has been shown to possess some predictability (Meehl et al. 2016b).
- (iv) 160 years of observational data from the Kaplan Extended Sea Surface Temperature, version 2, dataset (Kaplan SST; Kaplan et al. 1998): This dataset assimilates ship and satellite observations from 1856 to the present.
- (v) Preindustrial control runs spanning 1000 years from eight CMIP5 models other than the CCSM4.

These are used briefly in section 4b for comparison with the four datasets described above. The models are listed in Table 1.

3. Attractor reconstruction using simplex projection

A key feature of chaotic systems is the existence of an underlying attractor in phase space whose topology can yield insights into the predictability, stability, and relationships between states of the system. There are several variables in the ocean–atmosphere system of the tropical Pacific, each with time-evolving values at each spatial location. While the strong coupling between these variables means that much of the information contained in this large collection of time series is redundant, it is not a trivial task to reduce the system to the correct subset of time series with which to construct the true phase space.

The simplex projection method (Sugihara et al. 2012; Deyle and Sugihara 2011) addresses this problem by applying Takens's embedding theorem (Takens 1981), according to which a *shadow manifold* preserving the topology of the true attractor can be reconstructed using any single time series from the system. The shadow manifold is n -dimensional, and each dimension is the chosen time series lagged by an integral multiple of a time interval τ . Thus the n dimensions of a shadow manifold reconstructed from a time series $x(t)$ are $x(t)$, $x(t - \tau)$, $x(t - 2\tau)$, \dots , $x[t - (n - 1)\tau]$. The algorithm seeks, and provides as output, the pair (n, τ) such that n is the dimension of the true attractor.

The shadow manifold preserves relative distances between points on the attractor, where each point

uniquely represents some state of the system. The simplex projection method leverages this property to select appropriate values of n and τ for the system as follows:

- 1) For each (n, τ) pair, a version of the shadow manifold is constructed by producing a “map” of the system’s trajectory in an n -dimensional space. For example, for the pair $n = 2$, $\tau = 3$ months, this would be generated by “plotting” the values from the time series in an x - y plane such that the points (x, y) are the values of the time series at a given time t and $t - 3$ (since the lag τ is 3 months) for all values of t available in the dataset. Although this cannot be visualized for values of n larger than 3, the concept is generalized to these dimensions by storing the time series in a matrix where each column of the matrix corresponds to a dimension. The first column contains the time series, and each of the remaining columns contain the time series lagged by τ months with respect to the column preceding it. Therefore, each row in the matrix represents a set of coordinates in an n -dimensional space. This type of matrix is generated for each possible (n, τ) pair.
- 2) A “prediction” is made for each of a number of randomly selected points in time using each of the matrices generated in step 1. This is done by tracing the trajectories of the selected point’s nearest neighbors (i.e., the points along the time series whose coordinates in the n -dimensional space place it closest to the coordinates of the selected point) into the future. The correct (n, τ) pair should, in theory, place points in time where similar behaviors occurred near each other (as would be the case on the system’s true attractor), meaning that the future trajectories of neighbors in the correct shadow manifold should be good predictors of the future behavior of points near them. The weighted average of these trajectories is calculated using weights that are inversely proportional to the distance between the neighbor and the selected point to produce the prediction for that point.
- 3) Then the error (absolute difference) between the predictions thus obtained and the actual trajectories following the selected points being predicted is measured for each of the (n, τ) pairs. The (n, τ) pair with the minimum average error over all predicted points is returned by the algorithm as the correct dimension and lag for reconstructing the shadow manifold of the system. Choosing the (n, τ) pair with the minimum average error over all predicted points thereby selects the shadow manifold with maximal predictive power.

In this study, we use the Niño-3 index [the monthly mean sea surface temperature (SST) anomaly in the region spanning 5°S–5°N, 150°W–90°E] to reconstruct the shadow manifold after smoothing using a low-pass second-order Butterworth filter with a 4-month cutoff. To capture interdecadal variability, we use the 15-yr standard deviation of the filtered Niño-3 index as the target for prediction in the simplex projection procedure. The 15-yr period is based on earlier work that found this to be an appropriate time scale for the identification of interdecadal shifts in the behavior of the tropical Pacific in the observations (Karspeck et al. 2004; Seager et al. 2004). (Using a 20-yr period made little difference for the results shown below.) These studies concluded that this behavior is potentially predictable in the ZC model. Another study (Ramesh et al. 2017) also showed that the Niño-3 index, filtered in an identical manner, contained sufficient information to make predictions of the 15-yr mean state.

Our earlier studies of TPDV (Karspeck et al. 2004; Seager et al. 2004; Ramesh et al. 2017) used the 15-yr mean of Niño-3 as an index, but here we abandon it in favor of the standard deviation over 15-yr periods. Prior work by others has used the standard deviation or variance over similar lengths of time to identify the phases of interdecadal variability in this region (e.g., Choi et al. 2009; Atwood et al. 2017). High-variance states correspond to warm, El Niño-like mean states, and low-variance states correspond to cool, La Niña-like mean states. The correlation of 15-yr mean and variance is so high (0.85 in the ZC model and 0.47 in observations, both significant at the 99% level) that the two must be regarded as different indices of the same phenomena. This correspondence arises in good measure from the asymmetry between El Niño (warm, high amplitude) and La Niña events (cold, lower amplitude; e.g., Okumura et al. 2017). While either the mean or the standard deviation are satisfactory indices of TPDV, we find that the standard deviation index more effectively distinguishes the extreme phases of TPDV. We discuss this further in the context of the structure of the attractor.

We focus on the ZC model because the substantially longer time series available yields greater certainty of the results than from observations or GCMs. In addition to its proven ability to simulate and predict TPDV, this model also has the advantage of isolating the physics internal to the tropical Pacific system meaning that all results relating to its predictability can be attributed to processes arising within the region. At every stage of this study, we compare our results with observations and the GFDL model and CCSM4, all of which incorporate the various external factors that can interact with the

tropical Pacific system. The differences between these datasets and the ZC model could result from these external factors, differences in the simulation of processes internal to the tropical Pacific, or the uncertainty associated with the shorter record lengths that may not fully characterize all possible behaviors of the system.

4. Results

a. Attractor reconstruction

For the ZC model, the optimal dimension and lag were found to be $n = 3$ and $\tau = 6$ months. This means that the system is reducible to three dimensions (i.e., three time series determine its state) and that information from a span of 13 months (months t , $t - 6$, and $t - 12$) of the Niño-3 time series contain sufficient information to make predictions of its 15-yr standard deviation. The mean error obtained for this combination of dimension and lag obtained from the simplex projection method was 0.03°C , which is 1.6% of the range 0.04° – 1.83°C of the 15-yr standard deviations. We begin our interpretation of the results by examining a short segment of the Niño-3 index in the state space as the model state evolves through time. Statistics regarding the full dataset are presented after this short example.

Figure 1 depicts the trajectory of the system along the reconstructed shadow manifold through a 60-yr time interval, displayed every 5 years. The state space within which the system traverses its trajectory is, based on the results of the simplex projection procedure, defined by three dimensions: the Niño-3 time series on the x axis and the same time series lagged by 6 months on the y axis (as the optimal lag was determined to be 6 months) and lagged by 12 months on the z axis. Each location in this abstract space corresponds to a specific physical state of the system, and the full attractor encompasses all possible states that the system is able to occupy, barring any additional forcing. The location of the system in this space at a given point in time t is, by definition, the Niño-3 index value at time t plotted along the x dimension, the value of the same index at time $t - 6$ plotted along the y dimension, and at time $t - 12$ along the z dimension. The system thus moves along a trajectory in this space as the model state evolves over time, and this trajectory defines its shadow manifold. The pink line in Fig. 1, which traces the Niño-3 index with its lagged counterparts to give the shadow manifold, connects these states as they occur in chronological order over a 60-yr interval. The panels of Fig. 1 sequentially depict the location of the model state (blue dot) in its state space as determined by the Niño-3 index in order to show the temporal evolution of the system.

The first 10 years of the trajectory segment shown are spent tracing a large “orbit” in the three-dimensional state space (Figs. 1a,b). The system (whose current position is marked by the blue dot) moves clockwise along this orbit. By the year 15 (Fig. 1c), the system has entered a smaller, inner orbit. This orbit is also traversed clockwise, until the year 40 (Fig. 1h). By the year 45 (Fig. 1i), having spent approximately 30 years in the inner orbit, the system exits the inner orbit and reenters the outer orbit.

The full shadow manifold of the ZC model (a representative segment of which is depicted in Fig. 2a) possesses densely populated inner and outer orbits, separated by a sparsely populated region. The median time interval between transitions¹ from the outer to inner orbits is 62 years. This is similar to the length of the interval between the known shifts from a positive to a negative PDO—that is, in 1943 and 1999—in the observations, suggesting that the two orbits of the shadow manifold correspond to the regimes of behavior that comprise the model’s interdecadal variability. The size of an orbit is smaller for periods of time when the Niño-3 index displays low variance, and larger orbits correspond to periods of high variance. The structure of the attractor suggests that the standard deviation is a more discriminating measure of TPDV, as the inner and outer orbits possess clearly distinct standard deviations (as can be inferred from the large and small radii of the orbits, which correspond approximately to the standard deviation of the Niño-3 index) but similar means (which correspond approximately to the centers of the orbits). This is not simply an artifact of the 15-yr standard deviation being used as the prediction target in the attractor reconstruction procedure; when the same procedure was performed using the 15-yr mean as the target for prediction instead (result not shown), a similar three-dimensional shadow manifold was obtained. The error associated with the predictions from the simplex projection was higher (mean error = 0.26), and the lag obtained was longer (9 months) in this case, suggesting that the mean is both less predictable and requires more information in order to be predicted.

Another noteworthy feature of the shadow manifold is that the transitions between the inner and outer orbits appear confined to a region in the lower-left corner of the figure. Since proximity of the points on the attractor is indicative of similarity in the physical states represented by those points, this implies that they occur

¹The procedure used to identify the points in time at which transitions between the orbits occur is described in section 4c.




FIG. 1. The trajectory of the ZC model system through the three-dimensional state space of its shadow manifold over time. The axes represent the Niño-3 index ($^{\circ}\text{C}$; x axis) and the same index lagged by 6 months (y axis) and 12 months (z axis). The yellow \times marks the starting point (at time = 0), and the blue dot represents the position of the system at the time indicated in the panel title, after having traced the trajectory indicated by the pink line. The orbits are traced clockwise.

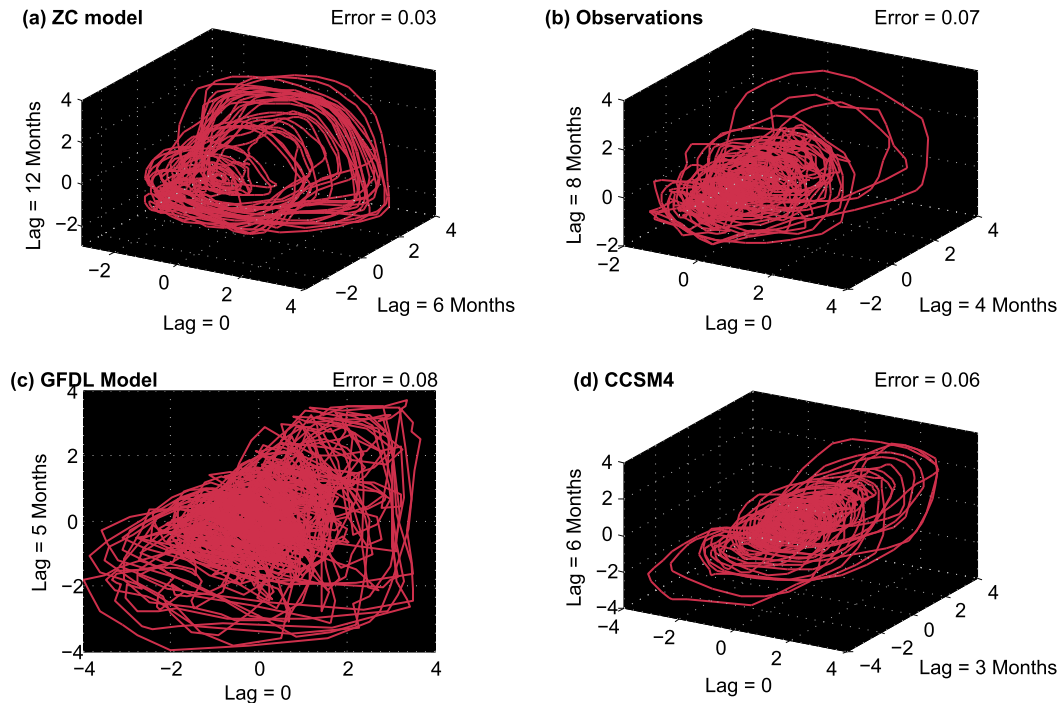


FIG. 2. The shadow manifolds constructed using the simplex projection method, which determines the number of dimensions and length of the lag, for each of the datasets used. For each dataset, 160 years of the time series are shown. The average error in prediction, normalized by dividing by the standard deviation of the target variable being predicted, obtained for each shadow manifold is indicated in the top-right corner of each panel.

during a specific subset of similar physical states of the system and therefore possess potential predictability.

The shadow manifolds obtained from all four datasets are visualized in Fig. 2. The time series for the ZC model is approximately two orders of magnitude longer than the GCM time series and three orders longer than the observed series, yielding more certainty in its manifold reconstruction. Also indicated are the minimum error values obtained during the reconstruction procedure, which can be interpreted as a measure of the overall predictability of TPDV in the system. While the lower error in the ZC model is no doubt in part attributable to the longer time series available, we note that the observations, with the shortest time series, are nonetheless more predictable than the GFDL model but less than CCSM4 by this measure.

The shadow manifold obtained for the observations from the simplex projection procedure is, like that of the ZC model, three-dimensional. A sparsely populated region is seen between a dense inner region and a few trajectories that form a wide outer orbit, suggesting the presence of a pair of orbits. Transitions between the orbits are confined to one side of the manifold. This suggests that, like the ZC model, the real world may possess some predictability associated with regime-like behavior. The short observational record means that

confidence in the accuracy of this shadow manifold is low, and the shape of the true attractor for this dataset must remain uncertain.

The GFDL model, unlike the other datasets examined, possesses a two-dimensional shadow manifold and lacks well-defined orbits. This, in combination with its relatively low predictive power from the simplex projection procedure, suggests that interdecadal variability in this model may be driven by stochastic processes rather than low-order chaos, consistent with earlier studies of the same dataset (Wittenberg 2009; Wittenberg et al. 2014).

CCSM4, like the ZC model and observations, produces a three-dimensional shadow manifold but does not display as distinct a separation between inner and outer orbits.

The lengths of the time series available from the last three of these datasets are orders of magnitude smaller than that from the ZC model. To assess the level of confidence in the shadow manifolds obtained from these shorter time series, we perform the same attractor reconstruction procedure on all subsamples of lengths 160, 4000, and 1000 years (the lengths of the time series from the observations, GFDL model, and CCSM4, respectively) from the ZC model, in order to estimate the probability of the shorter samples arriving at the same

number of dimensions (3) and lag (6 months) obtained from the full ZC model time series. The results are reported in Table 2. We find that the estimate of the dimension is fairly stable, as 72% of the 160-yr samples were able to arrive at the correct value despite the shortness of the record. The longer time series perform better, as expected. The dimension is far more likely to be correctly estimated than the lag: for the 160-yr segments, only 16% of the samples estimate the correct lag of 6 months in combination with the correct dimension. While there is substantial improvement in the estimation of the lag with the length of the time series, with the correct lag estimated for 44% of the 4000-yr segments, the confidence in the lags selected by the algorithm for the shorter datasets is still low. We also examine the probability of the lag being estimated correctly to within 3 months. In all three cases, there is a preference for values close to the correct lag of 6 months, as the majority of segments predict the lag to within 3 months of the correct value. When the shadow manifold of the ZC model is reconstructed using three dimensions and lags ranging from 3 to 9 months, we find that the “double orbit” structure is preserved, suggesting that even though the confidence in the lags selected by the shorter datasets is low, the resulting shadow manifolds retain structural similarities to the correct shadow manifolds provided that the dimension was determined correctly.

The simplex projection method requires the number of nearest neighbors used to be specified. This number is sensitive to the length of the record used and needs to be chosen so as to not include too much or too little information. We vary this number in order to test the stability of the dimension and lag obtained for the shorter datasets. The reported dimension and lag for each dataset were stable within a range of values: 9–15 neighbors for the observations, 21–40 neighbors for CCSM4, and 90–140 neighbors for the GFDL model. Thus the estimates of dimension and lag for the segments available from each system are robust, but this does not prove that they would be unchanged if much longer samples were available.

In the following section, we further examine the bimodality seen in the ZC model’s shadow manifold and compare the shorter datasets’ tendencies to display this behavior.

b. Comparison of the distributions of interdecadal standard deviations

The radius of an orbit in the constructed shadow manifolds corresponds approximately to the standard deviation of the time series over the duration of the orbit. Thus, a well-separated inner and outer orbit

TABLE 2. The probabilities of segments of different lengths making correct estimates of the properties of the full shadow manifold, as estimated using short segments from the ZC model. The correct dimension and lag in this case are 3 and 6 months, respectively. From left to right, the columns show the percentage of segments from the ZC model for which correct estimates were made of the number of dimensions, the lag, the number of dimensions and the lag in combination, the number of dimensions along with a lag within the range from 3 to 9 months.

Segment length (years)	Probability of correct estimate (%)			
	n	τ (months)	(n, τ)	$(n, \tau \pm 3$ months)
160	72	18	16	55
1000	85	26	23	68
4000	92	44	44	88

should yield a bimodal distribution of standard deviation values (computed over the relevant time scale). We verify this in Fig. 3, which shows the distributions of 15-yr standard deviations evaluated for all continuous 15-yr segments, including overlapping segments, in each of the datasets. The ZC model shows a clearly bimodal distribution, indicating two distinct regimes. Bimodality can also be seen to a lesser degree (and with less confidence) in the corresponding distributions for the observations and CCSM4, while the GFDL model displays the least bimodal tendency.

The differing lengths of the time series from each dataset complicates the comparison, as these lengths of time may not be sufficient to characterize the full distribution of interdecadal behaviors of each system. We make use of the longest time series, the 100 000 years from the ZC model, to estimate the length of time required to reliably capture this distribution.

Assuming that the true distribution of 15-yr standard deviations for this model is captured by the 100 000-yr time series, we compute the corresponding distribution for each continuous segment of a given length from the same model and calculate its Kolmogorov distance (the statistic used in the Kolmogorov–Smirnov test) to this true distribution. This is a measure of the similarity between two distributions, defined as the maximum difference between their empirical cumulative probability density functions. It can take values from 0 to 1, with 0 indicating identical distributions and increasing values indicating increasing differences. Thus, for each segment length, we arrive at a range of possible Kolmogorov distances between the distributions derived from these shorter segments and the true distribution. The Kolmogorov distances of the segments should converge to a small value at some time scale that captures the characteristics of the full distribution of interdecadal variability.

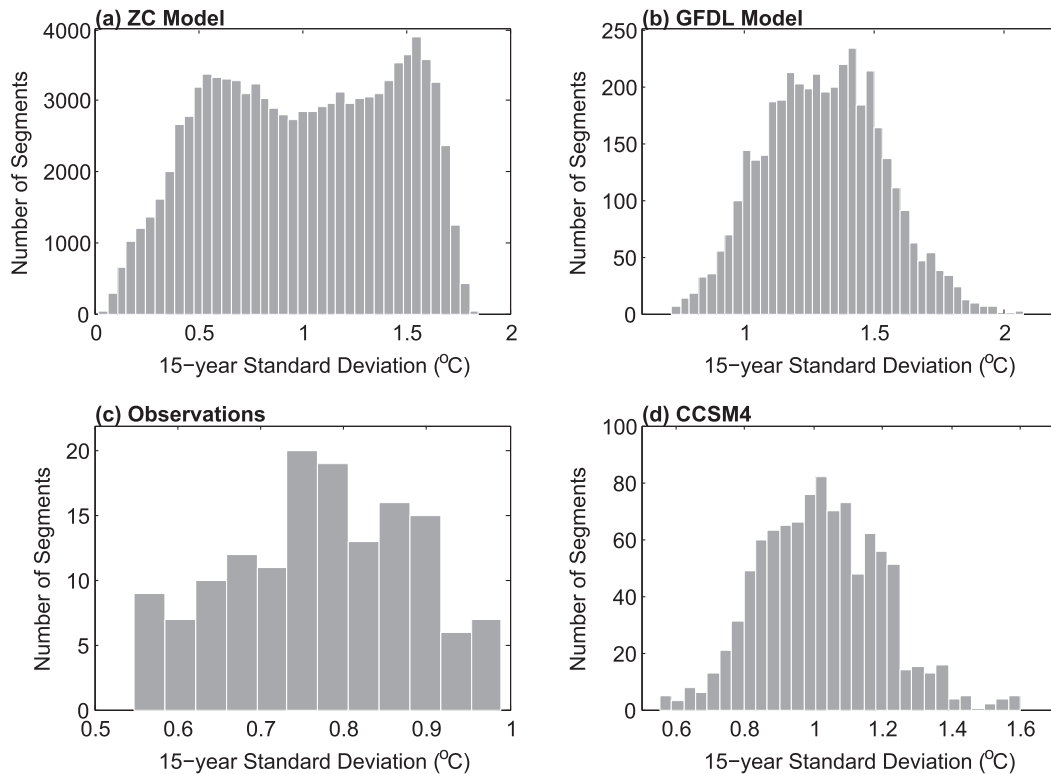


FIG. 3. Histograms of the standard deviations of all possible continuous 15-yr segments of the filtered Niño-3 index for each of the datasets used in this study.

Figure 4a shows the resulting distributions of Kolmogorov distances in the ZC model for a range of time scales from 100 to 90 000 years. The median Kolmogorov distance to the full distribution decreases rapidly as the length of the segments increases. The envelope of possible Kolmogorov distances narrows significantly between 15 000 and 20 000 years, with the standard deviation falling below 0.01 at 20 000 years. This suggests that at least 20 000 years of data are needed to fully capture the features of the distribution. At the length of the longest other dataset (i.e., 4000 years), the spread of 95% of the distribution of Kolmogorov distances (indicated by the gray shading in Fig. 4) falls below 0.1 but continues to decrease to a value of 0.04 at 20 000 years, indicating that the shape of the distribution of 15-yr standard deviations has not yet stabilized at 4000 years. It is clear from Fig. 4a that the distribution of Kolmogorov distances has a large spread at both 1000 and 160 years, implying that the samples from CCSM4 and the observations are unlikely to be representative of the true distributions of 15-yr standard deviations in these systems.

We indicate the Kolmogorov distance between the full distribution of 15-yr standard deviations for each dataset (i.e., each CMIP model and observations) and

that of the ZC model in Fig. 4a. The observations are found to be at a Kolmogorov distance of 0.52 from the full distribution of the ZC model (indicated by the black triangle). This lies within the range of the possible distances between 160-yr segments from the ZC model and its own full distribution. The distribution from the observations is further from the full distribution than for 97.5% of the 160-yr samples from the ZC model, with 2296 segments at a distance greater than this value.

The distance between the 1000-yr segment of the CCSM4 (black circle) and the full distribution is slightly below the maximum possible distance between segments from the ZC model and its own full distribution, with 39 segments at a greater distance, indicating marginal similarity between the datasets. The 4000-yr GFDL model segment (black square) is markedly different from the ZC model and lies far outside the distribution.

We also perform this calculation for 1000-yr pre-industrial runs of nine other models from the CMIP5 archive (colored circles). The segment from CCSM4 is closer to the full ZC model distribution than the segments from the other CMIP5 models, all of which lie outside the distribution.

When the same calculation is performed using the 4000-yr time series of the GFDL model as the reference

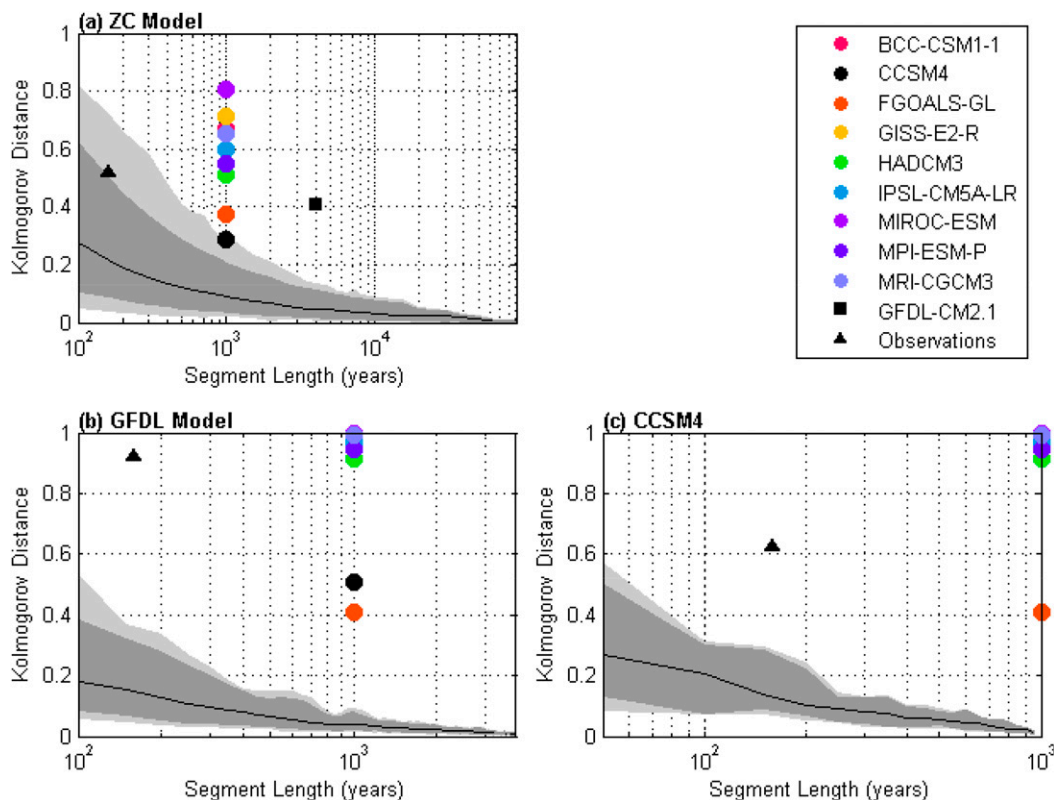


FIG. 4. (a)–(c) The range of Kolmogorov distances between the distributions in Figs. 3a, 3b, and 3d, respectively, and similarly derived distributions from shorter segments of the same dataset are shown. The black line indicates the median of the Kolmogorov distances, the dark shading fills between the 2.5th and 97.5th percentiles, and the light shading extends to the extreme values. The corresponding distances calculated from the full time series of observations, from 1000-yr simulations from nine CMIP5 models (including CCSM4), and the GFDL model are also shown.

distribution (Fig. 4b), we find the Kolmogorov distance from the observations to be far outside the distribution of distances for 160-yr segments from the GFDL model, indicating that the interdecadal behavior in the observed period is not captured by the 4000 years from the GFDL model used in this study. This is consistent with an earlier analysis of this dataset (Atwood et al. 2017). Most of the CMIP5 models cluster far from the GFDL model, showing that this model is distinctly different from all other CMIP5 models with available 1000-yr-long simulations.

Using the 1000-yr run of CCSM4 as the reference distribution (Fig. 4c), we find that the observations are closer to this model than the GFDL model but still lie outside the distribution. The other CMIP5 models are also clearly different from this model. This indicates that the representation of the TPDV differs widely between the available long runs of GCMs. This could be either a result of the shortness of these datasets or because the mechanisms that generate the TPDV vary across the GCMs.

Of the three models examined in our study, the observations are closest to the ZC model by this metric. This analysis indicates that the interdecadal variability of the real-world tropical Pacific system could lie within the range of behaviors captured by the ZC model, possessing distinct regimes, and that it certainly lies outside the range of behaviors found in the 4000-yr run of the GFDL model, which does not possess such regimes and has been previously identified as noise driven. Based on these results, we further investigate the predictability of the ZC model within the framework of its shadow manifold, encouraged by the possibility of its relevance to the real-world system.

c. Locating regions of predictability on the shadow manifold

A valuable property of the attractor of a chaotic system is that it can be used to classify states of the system based on their predictability: regions of the attractor where points lying close together (and therefore possessing similar physical characteristics) separate into

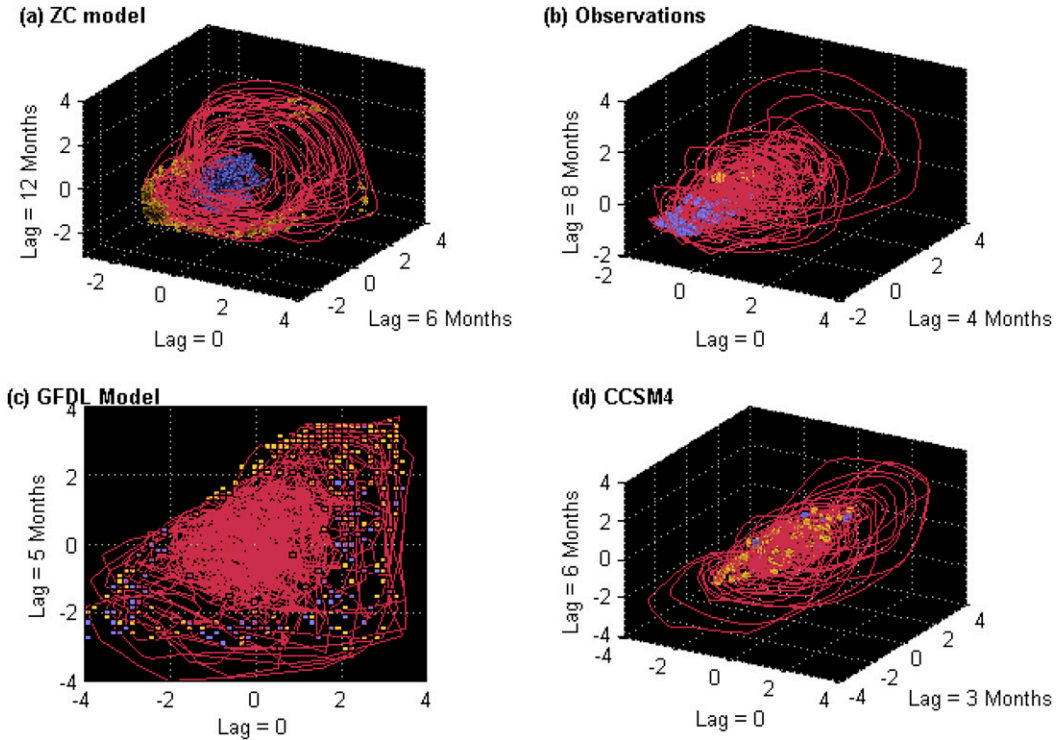


FIG. 5. The shadow manifold of each of the systems studied, with regions of high predictability indicated by colored grid cells. Blue (yellow) grid cells indicate high-confidence predictions of a low-variance (high variance) state 15 years into the future.

widely differing trajectories can be considered highly unpredictable or sensitive, as a small perturbation (which would amount to a small shift in position on the attractor) can lead to a drastically different future. On the other hand, regions where points lying close together follow similar trajectories can be considered predictable, as similar states lead to similar outcomes and small perturbations do not alter this future in a significant way. By preserving the topology of the attractor, the shadow manifold preserves this property.

In this section, we continue to use the 15-yr standard deviation of the filtered Niño-3 index as the target for prediction. Based on the bimodality of the ZC model’s distribution in Fig. 3, we reduce the problem of classifying predictions to a simple binary: “positive” or “negative.” For the ZC model, this means predicting whether the 15-yr standard deviation of the 15 years following a given point is greater than (high variance, referred to as positive) or less than (low variance, referred to as negative) a threshold of 1.1°C, the midpoint between the two peaks of the distribution. To classify regions of the shadow manifold on the basis of predictability, we divide the shadow manifold into grid cells of volume $(0.1^\circ\text{C})^3$ and consider only cells containing 50 or more points from the time series.

The prediction (positive or negative) for each grid cell is said to be that made by the majority of trajectories passing through the grid cell. We measure the predictability of each cell as the percentage of predictions from the cell that agree with the majority prediction, yielding a metric taking values from 50% to 100%.

In Fig. 5, we highlight the regions of high predictability on the shadow manifolds of Fig. 2 based on this metric. For the ZC model, we highlight the grid cells with 85% or more agreement and indicate positive predictions in yellow and negative predictions in blue. The high-predictability cells form distinct clusters on the shadow manifold, with the positive and negative predictions being well separated. This has the following implications:

- 1) 25.4% of grid cells, enclosing 37% of the points on the time series, possess a predictability of 85% or higher. Both positive and negative states can be predicted with high confidence from these points.
- 2) The clustering of these cells means that the high predictability is not a result of pure chance: since physically similar states lie close together on the shadow manifold, a cluster of grid cells corresponds to a small range of physically similar states. These are

the physical states of the system that are highly predictable on interdecadal time scales.

- 3) The separation of the positive and negative clusters shows that the states making each type of prediction are physically distinct and relatively stable: a small perturbation is unlikely to convert a point from making a high-confidence negative to positive prediction (or vice versa).

The frequency with which the model passes through the predictable grid cells is also of interest, as it indicates how often useful predictions can be made. We find that the frequency is highly irregular but that the model does not spend long durations of time without passing through these clusters, as the mean length of the intervals spent outside of the highly predictable clusters is 16 months, with 95% of such intervals being shorter than 6 years.

We conduct a similar analysis for the other shadow manifolds, using a lower threshold of predictability (70%) and an appropriate minimum number of trajectories per grid cell (5 for observations, 10 for GCMs) to identify regions of relatively high predictability in these systems. We use thresholds of 0.82° , 1.16° , and 1.31°C for the observations, CCSM4, and the GFDL model, respectively. For the first two of these datasets, these thresholds are the midpoint between the two highest peaks of the distributions in Fig. 3, as for the ZC model. For the GFDL model, we use the median value from the distribution since bimodality is not apparent in this case. For the sparsely populated shadow manifold from the observations, we also use a larger cell volume of $(0.2^\circ\text{C})^3$. The observations show a clustering of the high-confidence negative predictions, with the grid cells making high-confidence positive predictions in a separate part of the space. In the shadow manifold of the CCSM4, there is a cluster of high-confidence positive predictions, but high-confidence negative predictions are rare and do not cluster, suggesting a stark asymmetry between the predictability of positive and negative states in this model. The uncertainty in the lag of these two shadow manifolds means that there is low confidence in the location of these clusters.

The shadow manifold of the GFDL model possesses many cells passing the predictability threshold for both types of predictions. The cells making positive and negative predictions are well mixed and situated close together, indicating a highly unpredictable system: small perturbations can shift the system from a state that leads to a high-variance future to one that leads to a low-variance future. Such a system can be expected to respond to noise or external forcing with changes in interdecadal behavior, whereas predictable states in

systems like the ZC model and perhaps the observations would remain predictable even if perturbed.

We examine the predictability of the ZC model further in Fig. 6. Classifying the grid cells by predictability (Fig. 6a) shows that while predictability varies widely over the full shadow manifold, more than half of the cells exhibit at least 75% predictability.

The transitions between orbits in the ZC model's shadow manifold appear to be confined to a region in the lower left of Fig. 5a. To explore the possibility of predicting transitions between states, we locate the points of the Niño-3 time series when these transitions occur in the three-dimensional space of the shadow manifold and determine the predictability of each transition based on the grid cell in which it is situated.

To select the points in time that represent transitions, we compute the standard deviation of all 15-yr segments of the filtered Niño-3 time series beginning in January (therefore, the 100 000-yr time series yields 99 986 standard deviation values). The standard deviations are then classified into high or low variance based on the cutoff of 1.1°C . Then we select the points in time where a high-variance (low variance) 15-yr segment starting in a given year is followed by a low-variance (high variance) 15-yr segment starting in the subsequent year. Note that consecutive segments overlap, with start points separated by one year (i.e., each January).

We assess the predictability of each point within a 2-yr window centered on the selected point based on the grid cell of the shadow manifold in which it is located and choose the point that possesses the maximum predictability. This allows us to select the points in time that represent the maximum predictive power within the time frame during which each transition takes place and eliminate the arbitrary selection of the month of January.

Figure 6b shows the distribution of the transitions based on their predictability assessed in this manner. Comparing this to the distribution in Fig. 6a indicates that transitions are more predictable than the average state of the system, with 80.2% exhibiting a predictability of 80% or higher; 66.4% of the transitions have a predictability of 85% or higher and therefore lie within the predictable regions highlighted in Fig. 5a. Examining the negative and positive transitions separately (Figs. 6c,d) reveals that most transitions of both types have high predictability, with transitions into high-variance states slightly more likely to be highly predictable than transitions into low-variance states. This is consistent with earlier studies of interdecadal predictability in this model that found positive states to be more predictable than negative states (Karspeck et al. 2004; Ramesh et al. 2017).

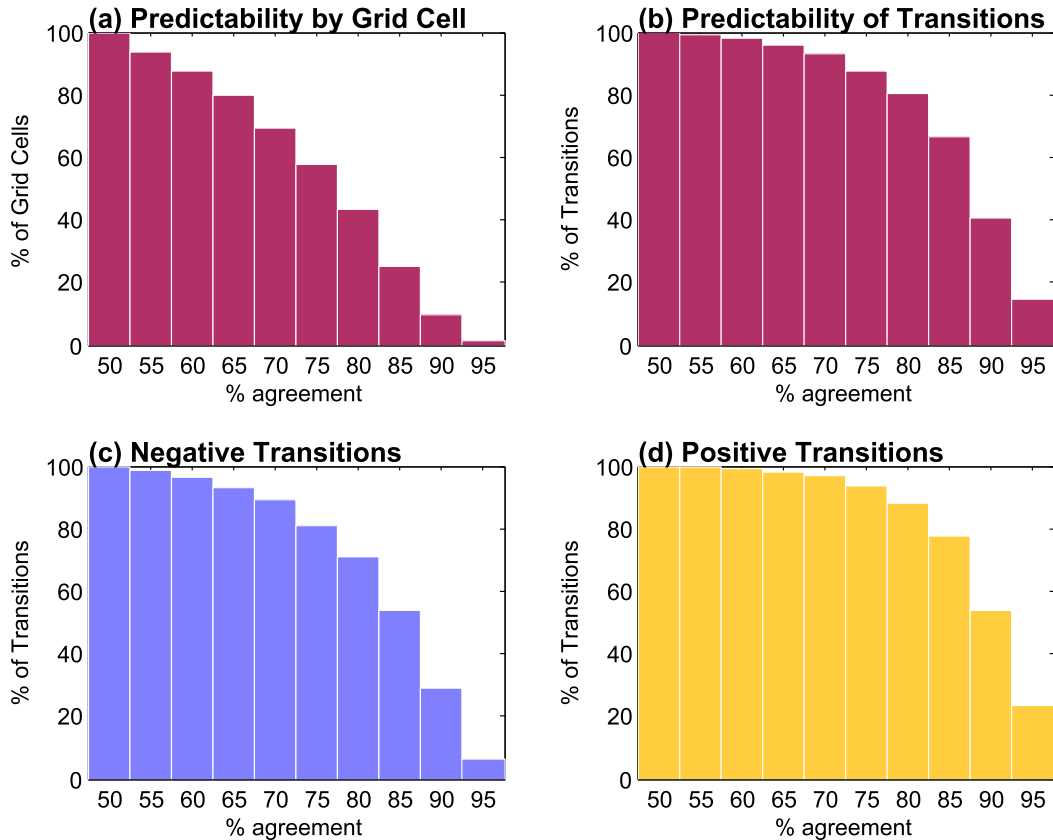


FIG. 6. (a) Cumulative histogram of grid cells on the shadow manifold of the ZC model classified by their predictability and (b) cumulative histogram of transitions between regimes classified based on the predictability of the grid cells in which they are located on the ZC model’s shadow manifold. (c),(d) As in (b), but for transitions to low- and high-variance states, respectively. The values in all panels are cumulatively summed from right to left, such that each bar indicates the number of (a) grid cells or (b)–(d) transitions that have agreement greater than or equal to the percentage on the x axis. The trajectories passing through each grid cell are traced for the 15-yr period after they exit the grid cell, and the sign of the prediction associated with the grid cell is that of the majority of these trajectories. The predictability depicted in this figure is the percentage of trajectories passing through the grid cell that are in agreement with the majority prediction.

d. Toward characterizing predictable states of the attractor in terms of physical variables

Each location on the attractor corresponds to a particular physical state of the model system. As the model’s trajectory passes through these locations, each point on the trajectory corresponds to a point in time of the model’s evolution, during which it occupies the particular physical state corresponding to these locations, in terms of variables such as sea surface temperatures, wind velocities, and thermocline depth. The same should be true of the shadow manifold if the attractor reconstruction has been performed correctly, as it preserves the topology of the attractor. In this section, we confirm this to be the case by examining the properties of one small region of the shadow manifold.

Figure 5a shows that the region where the model predicts a negative state with high confidence forms a

single small cluster of points lying close together (indicated in blue). If the reconstruction procedure has correctly preserved the structure of the system’s true attractor, the states within this cluster should show more similarity to each other in terms of the important physical variables of the system than to regions outside of the cluster. To test whether this is in fact the case, we locate all the points in time that lie within this high-predictability cluster and extract the model’s physical states at these times. We then compute the standard deviation of a few physical variables over these points in time, which we expect should be significantly smaller for this subset of points than for the full time series.

We report the following in Table 3:

- 1) The standard deviation for some physical variables for the model states that lie within the volume of the cluster, averaged over all spatial points. In all cases,

TABLE 3. Spatial averages of the standard deviations of the sea surface temperature anomaly, upper-ocean layer thickness, surface zonal wind speed, and divergence of surface wind speed compared between (column 2) the full range of model behaviors and (column 1) the range of behaviors constrained by the cluster shown in Fig. 5a. (column 3) The ratio between the values in columns 1 and 2. (column 4) The minimum value over the spatial domain when the ratio in column 3 is calculated individually for each spatial point.

Variable (anomaly from annual cycle)	Standard deviation within cluster	Standard deviation over full time series	Standard deviation ratio	Minimum standard deviation ratio
SST	0.22°C	0.47°C	0.47	0.32
Upper-ocean layer thickness	7.17 m	13.14 m	0.54	0.29
Surface zonal wind speed	0.22 m s ⁻¹	0.58 m s ⁻¹	0.38	0.32
Surface wind speed divergence	$2.37 \times 10^{-7} \text{ s}^{-1}$	$4.54 \times 10^{-7} \text{ s}^{-1}$	0.52	0.28

this value is found to be considerably smaller than the full range of variability expected for these variables.

- 2) The standard deviation for each variable, averaged over all spatial points, for all model states from the full time series.
- 3) The ratio between the standard deviations reported in columns 1 and 2: the smaller this ratio, the more tightly the spatial field of the variable is constrained by the cluster in phase space.
- 4) The minimum value of the standard deviation ratio in the spatial domain of the model: for each variable, the ratio of the standard deviations is computed for each spatial point, and the minimum value over the spatial domain is reported. Since the three dimensions of the shadow manifold correspond to three unknown physical variables whose spatial extent is unknown, there may be regions in space that are better constrained by the shadow manifold than others. This ratio is found to be approximately 0.3 for all the variables examined, showing that there is some spatial region where each of these variables is highly constrained.

Therefore, the points lying within the cluster and close together in phase space represent some subset of physically similar model states, as expected if the attractor topology has been preserved. In Fig. 7, we show statistics related to the spatial fields of the sea surface temperature anomalies (SSTAs), thermocline depth anomalies, and wind velocities extracted at the points in time lying within the region of the shadow manifold that predict a negative state with high confidence (shown in blue in Fig. 5a).

The left column of Fig. 7 displays the spatial fields of the standard deviation ratios computed for Table 3, that is, the standard deviation of the variable of interest at each spatial point over the states lying within the cluster divided by the standard deviation of the same variable at each spatial point over the entire time series. Note that these are not the 15-yr standard deviations but the

standard deviations calculated over the selected subset of points in time. We use this value to assess the relative importance of each variable at each location in determining the predictable state represented by the cluster on the shadow manifold. We rely on this ratio rather than the standard deviation within the cluster alone because a variable with a low standard deviation in a region of generally low variability is less informative than a variable with a low standard deviation in a region of large variability.

To confirm that the mean values of these variables within the set of states lying within the cluster are distinct from the mean of the values outside the cluster, we perform Welch's (unequal variance) *t* test at each spatial point, comparing the distribution of values within the set of states encompassed by the region predicting a negative state with high confidence to the distribution of values lying outside this cluster. We find that the hypothesis that these two sets of values come from distributions with equal means can be rejected at the 5% significance level at nearly all of the spatial points in the domain for the three variables examined (i.e., anomalies of SST, thermocline depth, and wind stress divergence). The regions where this is not the case are stippled in Fig. 7. These are found to be only two to four grid cells for each of the variables, with spatial averages of the *p* values over the domain being 0.0029, 0.0042, and 0.0026 for SST, thermocline depth, and wind stress divergence anomalies, respectively.

The right column of Fig. 7 shows the average spatial fields of SSTAs, thermocline depth anomalies, and wind velocity anomalies corresponding to the states within the cluster. These are averaged over the subset of points in time lying within the cluster. The black contour line over these maps encloses the regions from the corresponding maps to the left possessing standard deviation ratios smaller than 0.5 and therefore the most informative in defining the state of the system that reliably produces a negative transition. In the case of all three variables shown, the eastern tropical Pacific is decisive in

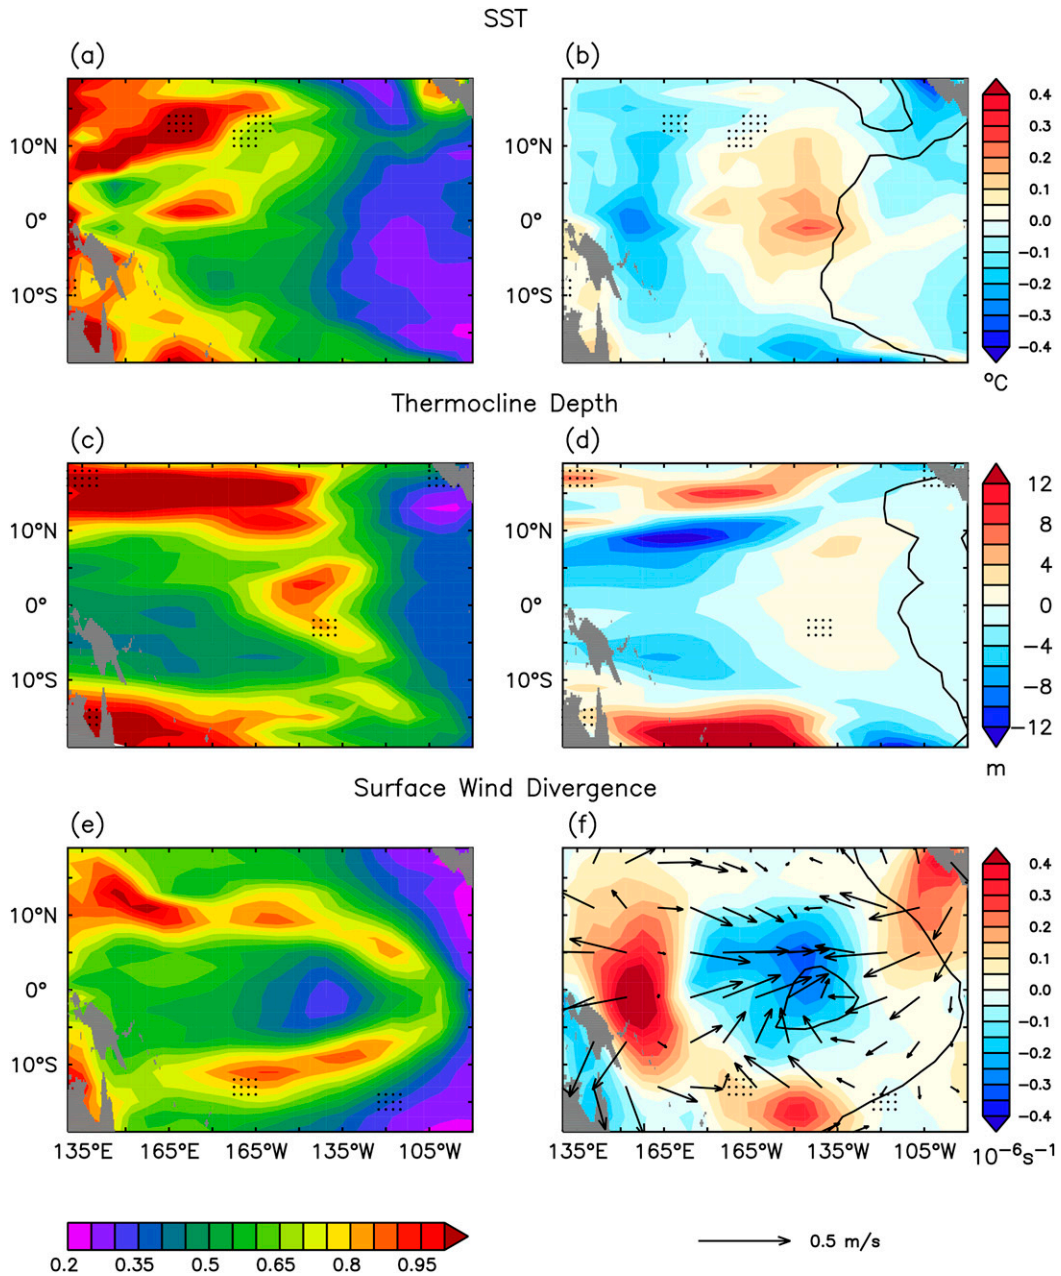


FIG. 7. (a),(c),(e) Ratio between the standard deviation at each spatial point for all temporal points within the high-predictability cluster and the same for all points over the full shadow manifold for (a) sea surface temperature, (c) thermocline depth, and (e) zonal wind velocity. (b),(d),(f) Mean spatial fields of the corresponding physical variables. The black contour encloses regions with a standard deviation ratio less than 0.5. The arrows in (f) indicate the mean wind velocity anomalies in the high-predictability region. Stippling indicates regions where the hypothesis that the values lying within the predictable cluster have the same mean as the values outside the cluster cannot be rejected at the 5% level.

determining this state. The predictable state is associated with small cool SSTAs in the eastern equatorial Pacific, a thermocline in the far eastern tropical Pacific that is slightly shallower than normal, and convergent wind velocity anomalies in the central equatorial Pacific

region accompanied by off-equatorial westerly anomalies in the eastern part of the basin. The small values of the anomalies are not a consequence of averaging over widely differing values, as the standard deviations seen in the left column (as well as in Table 3) for these regions

are small and these variables are close to climatological values in the regions indicated. While the processes in these regions have not been shown to control the evolution into a negative TPDV phase, they are at the very least reliable indicators of the change in phase that is about to occur.

The analysis in this section demonstrates that it is possible to isolate a particular state of interest based on its predictability characteristics using this attractor reconstruction method. Developing an understanding of the physics driving a negative transition will require further analysis of the model variables associated with this state and their evolution through time, which is beyond the scope of this study.

5. Discussion

The shadow manifolds of the four systems in this study were reconstructed based on the predictive power of the Niño-3 time series from each system. These shadow manifolds preserve crucial properties of the systems' true attractors, allowing us both to assess the predictability of the TPDV in these models and to classify the various states of each model by their predictability.

A wide range of behaviors is found in the models examined. When viewed as part of a chaotic system, the TPDV in the ZC model emerges as a pair of separate regimes of behavior corresponding to the two orbits seen in the shadow manifold. The attractor of this model is three-dimensional, implying that there exists some combination of three time series from the model that contains enough information to specify the state of the entire system. For over a third of the time, the model inhabits states where the phase of the TPDV for the following 15 years is predictable with 85% agreement. These predictable states form distinct clusters in the state space that correspond to particular physical states of the model. The predictability of states varies considerably over the shadow manifold, meaning that there are some states of the model from which there is little chance of making a correct prediction of the interdecadal variability but also that there exist states from which confident predictions of interdecadal variability can be made.

In contrast, the GFDL model's representation of the TPDV is one of a nearly unpredictable, noise-driven process, as expected from earlier work (Wittenberg et al. 2014). This model's shadow manifold is two-dimensional with no clear clustering of the predictable states, implying that these predictable states are likely the product of chance and do not correspond to specific physical states of the system that are inherently more predictable than others. This model's distribution of 15-yr standard

deviations is also distinctly different from that of the ZC model and lacks clear bimodality, which is further supported by its large Kolmogorov distance from the ZC model. These differences suggest that the TPDV in this model is likely generated by a fundamentally different set of processes than that in the ZC model.

The GFDL model's higher complexity may be responsible for some of these differences. However, the predictability of TPDV reported in GCMs other than the GFDL model (Meehl et al. 2016b; Ding et al. 2013; Thoma et al. 2015) shows that having higher complexity than the ZC model does not always result in highly unpredictable systems. The shorter time series from the CCSM4 and real world are quite different from the GFDL model and have more in common with the ZC model both in terms of their dimensionality (section 4a) and their distributions of 15-yr standard deviations (section 4b). This suggests that they may have similar predictability characteristics to the ZC model. The number of dimensions in these two cases is determined with some confidence (85% for CCSM4 and 72% for the observations). However, the low confidence in the lag implies that the precise structure of the shadow manifolds may not be correctly determined, so the clustering of predictable states on these shadow manifolds is less certain.

A key difference between GCMs and the intermediate-complexity ZC model is that the latter computes anomalies from a specified seasonal cycle, whereas GCMs must simulate the seasonal cycle. Representing the seasonal cycle accurately may be crucial in this context, as the interactions between the annual cycle and ENSO are likely to play an important role in generating the system's chaotic behavior (Tziperman et al. 1994; Jin et al. 1994; Stuecker et al. 2015). Accurate simulation of the seasonal cycle in the tropical Pacific remains a challenge for many GCMs (Flato et al. 2013), meaning that if the real-world system is indeed chaotic, they may not be able to correctly capture its predictability characteristics without first producing a realistic seasonal cycle.

Another important difference is that while the GCMs and observations include the effects of variability in other parts of the world on the tropical Pacific region, the ZC model does not. The influence of regions outside the tropical Pacific may alter the interdecadal predictability of the system, so assessing the magnitude of the influence of external factors on this region across different GCMs and in the observations would be valuable.

Our results indicate that the interdecadal characteristics of the real-world system bear most resemblance to the more predictable ZC model in terms of both the

distribution of 15-yr standard deviations and the dimensionality of the attractors obtained. The shortness of the observational record means that the confidence in this resemblance cannot be high. However, the substantial differences between the noise-driven GFDL model and the observations lend support to the hypothesis that TPDV in the real world may be the product of a chaotic oscillator.

If this is the case, the real-world system, like the ZC model's representation of TPDV, passes through states from which the likelihood of making successful interdecadal predictions varies substantially. In the ZC model, the periods when transitions between interdecadal states occur are among the most likely to yield successful predictions, whereas the persistence of either of the TPDV states is more difficult to reliably predict.

Applying the results of this study would require characterizing and identifying the states of the system in nature from which successful predictions can be made. The length of the time series required to confidently characterize interdecadal variability is a significant obstacle. The results of [section 4b](#) suggest on the basis of the Kolmogorov–Smirnov test that 20 000 years of data are needed to fully characterize the distribution of interdecadal variability and that a sample of at least 4000 years is needed to yield a result with fairly high confidence. Even the latter, lesser standard requires a far longer time series than the length of the observations available and of most GCM simulations. Paleoclimate proxies are able to provide us with some knowledge of the variability over time spans of comparable length, but at present, their low temporal resolution and limited precision render them unsuitable for shadow manifold reconstruction.

A more promising way forward would be to characterize the physical variables associated with the predictable states in the ZC model and compare their evolution to that in the observations at the times of known transitions between TPDV shifts. The results of [section 4d](#) show that the points on the shadow manifold lying within the predictable region examined correspond to a specific physical state. The transition to a negative TPDV state is associated with a shallow thermocline and cool SSTAs in the eastern tropical Pacific and convergent wind anomalies in the central equatorial Pacific region. An analysis of this physical state has the potential to shed light on the physical processes that give rise to a negative TPDV state, as the predictable cluster groups together the states that evolve similarly to produce this outcome. Our further studies will build upon this result to trace the temporal evolution of key variables as the system passes through the predictable state and to understand the processes associated with

predictable states and the dynamics underlying transitions between TPDV phases. A comparison of the physical processes occurring during known shifts of the TPDV in observations and the ZC model would shed light on the extent of the similarity between these systems and, therefore, the applicability of our results to real-world prediction.

This study demonstrates the usefulness of data-driven attractor reconstruction methods in studying the tropical Pacific system. The simplex projection method used here not only yields insights into its chaotic nature but also reveals a more complete and nuanced picture of predictability in the ZC model than a conventional perturbed-ensemble approach to the same problem ([Ramesh et al. 2017](#)). The reconstructed shadow manifold shows that the predictive skill of the model varies widely depending upon the state it occupies at the time the prediction is made. Unless the criteria for selecting the initial state of the model in perturbed-ensemble experiments are specifically designed to choose the most predictable states, such experiments can be expected to produce larger errors in their predictions than otherwise. Identifying predictable states would provide the opportunity to make a relatively confident prediction of the state of the TPDV by initializing it at the appropriate time. If we correctly identify the location of the system on the shadow manifold (a trivial task in the case of the ZC model), the likelihood of making a correct prediction at any given time is known. This information can be used to decide which predictions should be given the most weight when making a forecast based on the location of the system in its state space at the time the prediction is made. This is analogous to the seasonal variation of the predictability of ENSO: predictions made at certain times of year, when the system is known to be in a state that is relatively sensitive to perturbations, are known to be less likely to be correct ([Chen and Cane 2008](#)). Similarly, on an interdecadal time scale, the shadow manifold can be used to make the distinction between states that are the most and least likely to yield reliable predictions. Unlike the seasonal cycle, the traversal of predictable states for the interdecadal time scale does not have a regular frequency; however, our results indicate that the model is highly unlikely to spend more than 6 years without passing through a state that has a predictability of 85% or higher. The model's traversal of predictable states is more frequent than the time scale of the predictions that can be made (15 years) and the median length of interdecadal states (31 years). Therefore, a prediction scheme for this model using our results would be to monitor the Niño-3 index as the model evolves, wait until it arrives in a predictable region of the shadow manifold, and issue a forecast of

the standard deviation of the following 15 years at that point in time.

To make use of the results from our study in real-world prediction, recognizing specific physical states, if they exist, from which the natural system is predictable would be a necessary additional step. As there is significant uncertainty associated with the reconstructed shadow manifold of the observations, it is unlikely to be possible to use the Niño-3 index alone to identify whether the system is in a predictable state. Instead, this problem will need to be approached by comparing the physical processes associated with predictable states in the ZC model to the evolution of the system in the real world. If the processes underlying the TPDV in the ZC model and the real world are sufficiently similar, predicting this phenomenon in the real world may well be an attainable goal.

Acknowledgments. N. R. is funded by the National Aeronautics and Space Administration Earth and Space Science Fellowship (Award NNX14AK96H). We thank Dr. Marion Dumas for her comments on the early stages of this study.

REFERENCES

- Adams, J. B., M. E. Mann, and C. M. Ammann, 2003: Proxy evidence for an El Niño-like response to volcanic forcing. *Nature*, **426**, 274–278, <https://doi.org/10.1038/nature02101>.
- An, S.-I., J.-S. Kug, A. Timmermann, I.-S. Kang, and O. Timm, 2007: The influence of ENSO on the generation of decadal variability in the North Pacific. *J. Climate*, **20**, 667–680, <https://doi.org/10.1175/JCLI4017.1>.
- Atwood, A. R., D. S. Battisti, A. T. Wittenberg, W. H. G. Roberts, and D. J. Vimont, 2017: Characterizing unforced multidecadal variability of ENSO: A case study with the GFDL CM2.1 coupled GCM. *Climate Dyn.*, **49**, 2845–2862, <https://doi.org/10.1007/s00382-016-3477-9>.
- Barnett, T. P., D. W. Pierce, M. Latif, D. Dommenges, and R. Saravanan, 1999: Interdecadal interactions between the tropics and midlatitudes in the Pacific basin. *Geophys. Res. Lett.*, **26**, 615–618, <https://doi.org/10.1029/1999GL900042>.
- Cane, M. A., S. E. Zebiak, and Y. Xue, 1995: Model studies of the long term behavior of ENSO. *Natural Climate Variability on Decade-to-Century Time Scales*, D. G. Martinson et al., Eds., National Academy Press, 442–457.
- Chang, P., L. Ji, H. Li, and M. Flügel, 1996: Chaotic dynamics versus stochastic processes in El Niño–Southern Oscillation in coupled ocean–atmosphere models. *Physica D*, **98**, 301–320, [https://doi.org/10.1016/0167-2789\(96\)00116-9](https://doi.org/10.1016/0167-2789(96)00116-9).
- Chen, D., and M. A. Cane, 2008: El Niño prediction and predictability. *J. Comput. Phys.*, **227**, 3625–3640, <https://doi.org/10.1016/j.jcp.2007.05.014>.
- Chen, X., and J. M. Wallace, 2015: ENSO-like variability: 1900–2013. *J. Climate*, **28**, 9623–9641, <https://doi.org/10.1175/JCLI-D-15-0322.1>.
- Choi, J., S.-I. An, B. Dewitte, W. W. Hsieh, J. Choi, S.-I. An, B. Dewitte, and W. W. Hsieh, 2009: Interactive feedback between the tropical Pacific decadal oscillation and ENSO in a coupled general circulation model. *J. Climate*, **22**, 6597–6611, <https://doi.org/10.1175/2009JCLI2782.1>.
- Delworth, T. L., and Coauthors, 2006: GFDL’s CM2 global coupled climate models. Part I: Formulation and simulation characteristics. *J. Climate*, **19**, 643–674, <https://doi.org/10.1175/JCLI3629.1>.
- Deyle, E. R. and G. Sugihara, 2011: Generalized theorems for nonlinear state space reconstruction. *PLOS ONE*, **6**, e18295, <https://doi.org/10.1371/journal.pone.0018295>.
- Ding, H., R. J. Greatbatch, M. Latif, W. Park, and R. Gerdes, 2013: Hindcast of the 1976/77 and 1998/99 climate shifts in the Pacific. *J. Climate*, **26**, 7650–7661, <https://doi.org/10.1175/JCLI-D-12-00626.1>.
- Dong, B., R. T. Sutton, and A. A. Scaife, 2006: Multidecadal modulation of El Niño–Southern Oscillation (ENSO) variance by Atlantic Ocean sea surface temperatures. *Geophys. Res. Lett.*, **33**, L08705, <https://doi.org/10.1029/2006GL025766>.
- Emile-Geay, J., and M. A. Cane, 2009: Pacific decadal variability in the view of linear equatorial wave theory. *J. Phys. Oceanogr.*, **39**, 203–219, <https://doi.org/10.1175/2008JPO3794.1>.
- England, M. H., and Coauthors, 2014: Recent intensification of wind-driven circulation in the Pacific and the ongoing warming hiatus. *Nat. Climate Change*, **4**, 222–227, <https://doi.org/10.1038/nclimate2106>.
- Flato, G., and Coauthors, 2013: Evaluation of climate models. *Climate Change 2013: The Physical Science Basis*, T. F. Stocker et al., Eds., Cambridge University Press 741–866.
- Flügel, M., and P. Chang, 1999: Stochastically induced climate shift of El Niño–Southern Oscillation. *Geophys. Res. Lett.*, **26**, 2473–2476, <https://doi.org/10.1029/1999GL900550>.
- Francis, R. C., S. R. Hare, A. B. Hollowed, and W. S. Wooster, 1998: Effects of interdecadal climate variability on the oceanic ecosystems of the NE Pacific. *Fish. Oceanogr.*, **7**, 1–21, <https://doi.org/10.1046/j.1365-2419.1998.00052.x>.
- Henley, B. J., J. Gergis, D. J. Karoly, S. Power, J. Kennedy, and C. K. Folland, 2015: A tripole index for the interdecadal Pacific oscillation. *Climate Dyn.*, **45**, 3077–3090, <https://doi.org/10.1007/s00382-015-2525-1>.
- Herweijer, C., and R. Seager, 2008: The global footprint of persistent extra-tropical drought in the instrumental era. *Int. J. Climatol.*, **28**, 1761–1774, <https://doi.org/10.1002/joc.1590>.
- Jin, F.-F., D. Neelin, and M. Ghil, 1994: El Niño on the Devil’s Staircase: Annual subharmonic steps to chaos. *Science*, **264**, 70–72, <https://doi.org/10.1126/science.264.5155.70>.
- Kang, I. S., H. H. No, and F. Kucharski, 2014: ENSO amplitude modulation associated with the mean SST changes in the tropical central Pacific induced by Atlantic multidecadal oscillation. *J. Climate*, **27**, 7911–7920, <https://doi.org/10.1175/JCLI-D-14-00018.1>.
- Kaplan, A., M. A. Cane, Y. Kushnir, A. C. Clement, M. B. Blumenthal, and B. Rajagopalan, 1998: Analyses of global sea surface temperature 1856–1991. *J. Geophys. Res.*, **103**, 18 567–18 589, <https://doi.org/10.1029/97JC01736>.
- Karamperidou, C., M. A. Cane, U. Lall, and A. T. Wittenberg, 2014: Intrinsic modulation of ENSO predictability viewed through a local Lyapunov lens. *Climate Dyn.*, **42**, 253–270, <https://doi.org/10.1007/s00382-013-1759-z>.
- Karspeck, A. R., and M. A. Cane, 2002: Tropical Pacific 1976–77 climate shift in a linear, wind-driven model. *J. Phys. Oceanogr.*, **32**, 2350–2360, [https://doi.org/10.1175/1520-0485\(2002\)032<2350:TPCSIA>2.0.CO;2](https://doi.org/10.1175/1520-0485(2002)032<2350:TPCSIA>2.0.CO;2).

- , R. Seager, and M. A. Cane, 2004: Predictability of tropical Pacific decadal variability in an intermediate model. *J. Climate*, **17**, 2842–2850, [https://doi.org/10.1175/1520-0442\(2004\)017<2842:POTPDV>2.0.CO;2](https://doi.org/10.1175/1520-0442(2004)017<2842:POTPDV>2.0.CO;2).
- Kessler, W. S., 2002: Is ENSO a cycle or a series of events? *Geophys. Res. Lett.*, **29**, 2125, <https://doi.org/10.1029/2002GL015924>.
- Kleeman, R., 2008: Stochastic theories for the irregularity of ENSO. *Philos. Trans. Roy. Soc. London*, **366A**, 2511–2526, <https://doi.org/10.1098/rsta.2008.0048>.
- Kosaka, Y., and S.-P. Xie, 2013: Recent global-warming hiatus tied to equatorial Pacific surface cooling. *Nature*, **501**, 403–407, <https://doi.org/10.1038/nature12534>.
- Krishnan, R., and M. Sugi, 2003: Pacific decadal oscillation and variability of the Indian summer monsoon rainfall. *Climate Dyn.*, **21**, 233–242, <https://doi.org/10.1007/s00382-003-0330-8>.
- Mantua, N. J., and S. R. Hare, 2002: The Pacific decadal oscillation. *J. Oceanogr.*, **58**, 35–44, <https://doi.org/10.1023/A:1015820616384>.
- McGregor, S., N. J. Holbrook, and S. B. Power, 2007: Interdecadal sea surface temperature variability in the equatorial Pacific Ocean. Part I: The role of off-equatorial wind stresses and oceanic Rossby waves. *J. Climate*, **20**, 2643–2658, <https://doi.org/10.1175/JCLI4145.1>.
- Meehl, G. A., A. Hu, B. D. Santer, and S.-P. Xie, 2016a: Contribution of the interdecadal Pacific oscillation to twentieth-century global surface temperature trends. *Nat. Climate Change*, **6**, 1005–1008, <https://doi.org/10.1038/nclimate3107>.
- , —, and H. Teng, 2016b: Initialized decadal prediction for transition to positive phase of the interdecadal Pacific oscillation. *Nat. Commun.*, **7**, 11718, <https://doi.org/10.1038/ncomms11718>.
- Newman, M., G. P. Compo, and M. A. Alexander, 2003: ENSO-forced variability of the Pacific decadal oscillation. *J. Climate*, **16**, 3853–3857, [https://doi.org/10.1175/1520-0442\(2003\)016<3853:EVOTPD>2.0.CO;2](https://doi.org/10.1175/1520-0442(2003)016<3853:EVOTPD>2.0.CO;2).
- , and Coauthors, 2016: The Pacific decadal oscillation, revisited. *J. Climate*, **29**, 4399–4427, <https://doi.org/10.1175/JCLI-D-15-0508.1>.
- Okumura, Y. M., 2013: Origins of tropical Pacific decadal variability: Role of stochastic atmospheric forcing from the South Pacific. *J. Climate*, **26**, 9791–9796, <https://doi.org/10.1175/JCLI-D-13-00448.1>.
- , T. Sun, and X. Wu, 2017: Asymmetric modulation of El Niño and La Niña and the linkage to tropical Pacific decadal variability. *J. Climate*, **30**, 4705–4733, <https://doi.org/10.1175/JCLI-D-16-0680.1>.
- Power, S., T. Casey, C. Folland, A. Colman, and V. Mehta, 1999: Interdecadal modulation of the impact of ENSO on Australia. *Climate Dyn.*, **15**, 319–324, <https://doi.org/10.1007/s003820050284>.
- Ramesh, N., M. A. Cane, R. Seager, and D. E. Lee, 2017: Predictability and prediction of persistent cool states of the tropical Pacific Ocean. *Climate Dyn.*, **49**, 2291–2307, <https://doi.org/10.1007/s00382-016-3446-3>.
- Seager, R., A. R. Karspeck, M. A. Cane, Y. Kushnir, A. Giannini, A. Kaplan, B. Kerman, and J. Velez, 2004: Predicting Pacific decadal variability. *Earth's Climate: The Ocean-Atmosphere Interaction*, *Geophys. Monogr.*, Vol. 147, Amer. Geophys. Union, 105–120.
- Stuecker, M. F., F.-F. Jin, and A. Timmermann, 2015: El Niño–Southern Oscillation frequency cascade. *Proc. Natl. Acad. Sci. USA*, **112**, 13 490–13 495, <https://doi.org/10.1073/pnas.1508622112>.
- Sugihara, G., R. May, H. Ye, C. Hsieh, E. Deyle, M. Fogarty, and S. Munch, 2012: Detecting causality in complex ecosystems. *Science*, **338**, 496–500, <https://doi.org/10.1126/science.1227079>.
- Takens, F., 1981: Detecting strange attractors in turbulence. *Dynamical Systems and Turbulence*, D. Rand and L.-S. Young, Eds., Lecture Notes in Mathematics, Vol. 898, Springer-Verlag, 366–381.
- Taylor, K. E., R. J. Stouffer, and G. A. Meehl, 2012: An overview of CMIP5 and the experiment design. *Bull. Amer. Meteor. Soc.*, **93**, 485–498, <https://doi.org/10.1175/BAMS-D-11-00094.1>.
- Thoma, M., R. J. Greatbatch, C. Kadow, and R. Gerdes, 2015: Decadal hindcasts initialized using observed surface wind stress: Evaluation and prediction out to 2024. *Geophys. Res. Lett.*, **42**, 6454–6461, <https://doi.org/10.1002/2015GL064833>.
- Thompson, C. J., D. S. Battisti, C. J. Thompson, and D. S. Battisti, 2001: A linear stochastic dynamical model of ENSO. Part II: Analysis. *J. Climate*, **14**, 445–466, [https://doi.org/10.1175/1520-0442\(2001\)014<0445:ALSDMO>2.0.CO;2](https://doi.org/10.1175/1520-0442(2001)014<0445:ALSDMO>2.0.CO;2).
- Timmermann, A., 2003: Decadal ENSO amplitude modulations: A nonlinear paradigm. *Global Planet. Change*, **37**, 135–156, [https://doi.org/10.1016/S0921-8181\(02\)00194-7](https://doi.org/10.1016/S0921-8181(02)00194-7).
- , and F. Jin, 2002: A nonlinear mechanism for decadal El Niño amplitude changes. *Geophys. Res. Lett.*, **29**, 1003, <https://doi.org/10.1029/2001GL013369>.
- Trenberth, K. E., 2009: An imperative for climate change planning: Tracking Earth's global energy. *Curr. Opin. Environ. Sustainability*, **1**, 19–27, <https://doi.org/10.1016/j.cosust.2009.06.001>.
- Tziperman, E., L. Stone, M. A. Cane, and H. Jarosh, 1994: El Niño chaos: Overlapping of resonances between the seasonal cycle and the Pacific Ocean-atmosphere oscillator. *Science*, **264**, 72–74, <https://doi.org/10.1126/science.264.5155.72>.
- , S. E. Zebiak, and M. A. Cane, 1997: Mechanisms of seasonal–ENSO interaction. *J. Atmos. Sci.*, **54**, 61–71, [https://doi.org/10.1175/1520-0469\(1997\)054<0061:MOSEI>2.0.CO;2](https://doi.org/10.1175/1520-0469(1997)054<0061:MOSEI>2.0.CO;2).
- Vimont, D. J., 2005: The contribution of the interannual ENSO cycle to the spatial pattern of decadal ENSO-like variability. *J. Climate*, **18**, 2080–2092, <https://doi.org/10.1175/JCLI3365.1>.
- Wittenberg, A. T., 2009: Are historical records sufficient to constrain ENSO simulations? *Geophys. Res. Lett.*, **36**, L12702, <https://doi.org/10.1029/2009GL038710>.
- , A. Rosati, T. L. Delworth, G. A. Vecchi, and F. Zeng, 2014: ENSO modulation: Is it decadal predictable? *J. Climate*, **27**, 2667–2681, <https://doi.org/10.1175/JCLI-D-13-00577.1>.
- Zebiak, S. E., and M. A. Cane, 1987: A model El Niño–Southern Oscillation. *Mon. Wea. Rev.*, **115**, 2262–2278, [https://doi.org/10.1175/1520-0493\(1987\)115<2262:AMENO>2.0.CO;2](https://doi.org/10.1175/1520-0493(1987)115<2262:AMENO>2.0.CO;2).
- Zhang, Y., J. M. Wallace, and D. S. Battisti, 1997: ENSO-like interdecadal variability: 1900–93. *J. Climate*, **10**, 1004–1020, [https://doi.org/10.1175/1520-0442\(1997\)010<1004:ELIV>2.0.CO;2](https://doi.org/10.1175/1520-0442(1997)010<1004:ELIV>2.0.CO;2).