MS-No.:npg-2024-24

Title: Explaining the high skill of Reservoir Computing method in El Niño prediction

Authors: Francesco Guardamagna, Claudia E. Wieners and Henk A. Dijkstra

Point-by-point reply to reviewer #2

February 9, 2025

We thank the reviewer for their careful reading and for the useful comments on the manuscript.

Overview

Reservoir Computer (RC) is one special version of RNN, which has been applied to build ENSO prediction model including the study in this article. This study aims to explain the high skill of RC in El Nino prediction theoretically. Based on ideal experiments, the author uses various ZC models in different regimes to generate "observation", and uses RC to learn these data so that the ENSO dynamic characteristics of ZC and the Nino trajectory can be learned. From the results, whether in the subcritical or supercritical regime, RC shows high performance. In addition, through CNOP calculation, the sensitivities of RC and ZC to the initial field are explored with meaningful results. However, it seems to be contrary to the main purpose of the study, and it does not seem to fully explain the reason why RC can produce high El Nino forecasting skills. This is my biggest doubt about this work, and of course it is also the most interesting point. I hope the author can have a more elegant explanation. In addition, the following are some thoughts and suggestions on this work or article:

Major comments:

1. *The results in Figure 2 make me think deeply. It tells us that when training a model, it doesn't mean that the richer the data included, the better the results will be. It seems to be related to the inherent dynamic characteristics of the system. Could you please explain why. For example, in the sub-critical state, why the prediction skill is better when wind field is not included in the training period? Although you attribute*

it to the sensitivity to wind noise, this is not specific enough. I think it can be discussed in more detail. By the way, I do not understand the sentence in Line 115.

Author's Reply:

In Machine Learning, adding input features does not always enhance performance, as redundant or irrelevant information can degrade model efficiency. In the subcritical regime, where ENSO variability is primarily noise-driven and the system is linearly stable, including surface wind speed anomalies (τ_c) during training may appear beneficial. However, our results show that the impact of τ_c depends on the forecast horizon. When initialized from ENSO neutral conditions, optimal atmospheric noise patterns can trigger transient growth of perturbations, provided the initial conditions are favorable. Conversely, if a perturbation is already developing, subsequent noise patterns can either reinforce or dampen its evolution. This makes τ_c particularly useful for predicting short-term variability, as it provides critical information about the external forcing that influences early perturbation dynamics. Accordingly, the Reservoir Computer (RC) achieves better accuracy at shorter lead times (3–6 months) when τ_c is included. At longer lead times (9–18 months), improved predictive performance requires the model to rely more on the system's internal dynamics rather than the short-term influence of stochastic noise. Including τ_c during training can lead to overfitting, causing the model to focus excessively on short-term noise patterns instead of learning the internal system dynamics. As a result, model performance deteriorates at longer lead times when τ_c is included. On line 115, we clarify that instead of directly using zonal wind stress anomalies to train the RC and LR model, we use zonal surface wind speed anomalies as a proxy. These two variables are inherently correlated through the bulk formula, conveying similar information. However, a key distinction arises due to how noise is introduced in the Zebiak and Cane (ZC) model: we introduce stochasticity in the form of random zonal wind stress bursts. This results in random local fluctuations in the zonal wind stress signal that are inherently difficult for the RC and LR models to predict and reproduce. In contrast, the surface wind speed anomaly signal is smoother and more predictable, making it easier for the RC and LR models to learn and generalize effectively.

To illustrate this, Fig. 1 below shows the relationship between zonal surface wind speed anomalies and zonal wind stress anomalies, both normalized by their mean and standard deviation, in the supercritical and subcritical regimes.

Changes in the Manuscript:
We will more clearly describe the RC performances in the "RC performances" section on page 8, focusing on the contribution of the variable τ_c in the subcritical and supercritical regime. We will better explain why we choose to use surface zonal wind speed anomalies as a proxy for wind stress in the "Reservoir Computer" section.

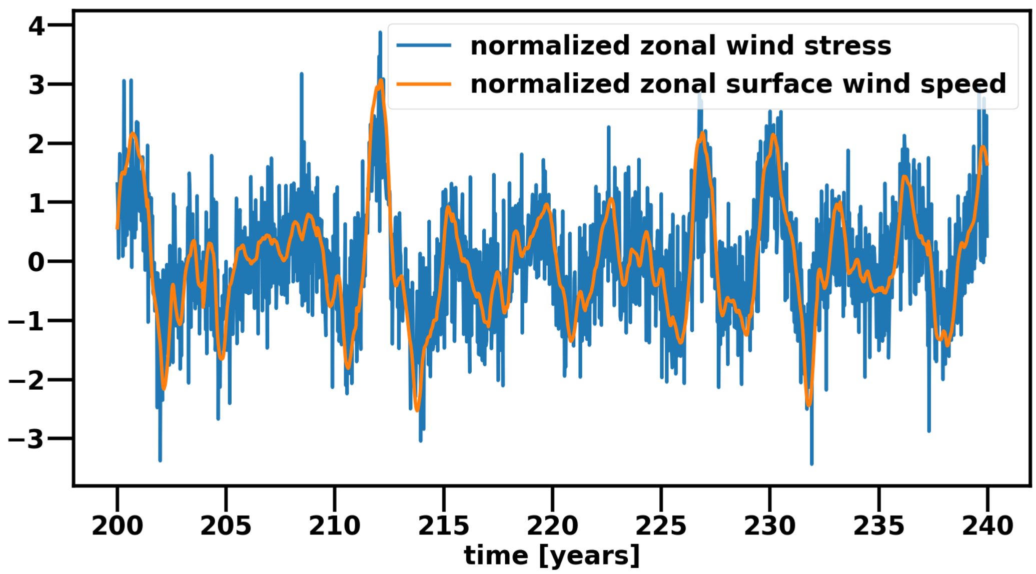

(a)

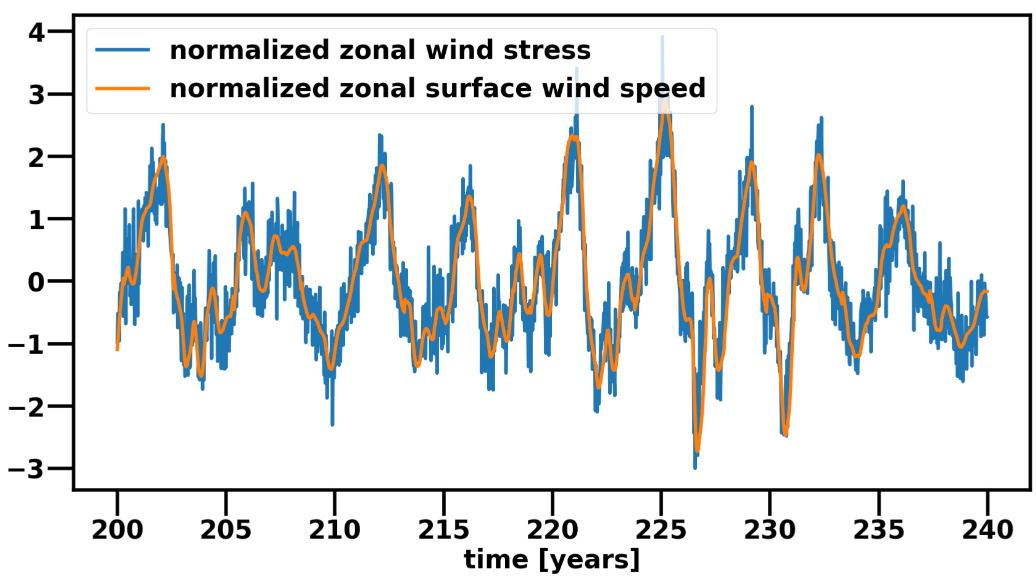

(b)

Figure 1: Relationship between normalized zonal surface wind speed anomalies and zonal wind stress anomalies in the subcritical (a) and supercritical regimes (b).

2. *For the CNOP part, "lead time" in ms is optimization time, right?*

   Author's Reply:
   In Fig. 4 and 5 in the "CNOP Analysis" section, the lead time in months on the x-axis corresponds to the optimization time considered during CNOP computation.

   Changes in the Manuscript:
   We will clarify that the lead time on the x-axis of Fig. 4 and 5 represents the optimization time considered during CNOP computation.

3. *Some pictures need to be refined. For example, it is recommended that the abscissa and ordinate in Figure A1 should be changed into the format of latitude and longitude coordinates.*

   Author's Reply:
   Suggestions will be followed.

   Changes in the Manuscript:
   In the revised manuscript, we will refine the figures, including modifying the axes in Fig. A1.

4. *The calculation of CNOP in complicated climate models has always been a major problem. How did you use the gradient-free Cobyla optimization algorithm to solve it? In addition, it is necessary to further verify whether the obtained CNOP is truly the CNOP. It is recommended to add random small perturbations to the obtained CNOP and project it onto the constraint conditions to compare the development of errors, so as to prove that the solution of CNOP is optimal.*

   Author's Reply:
   The Constrained Optimization BY Linear Approximation (COBYLA) algorithm is a gradient-free optimization method designed for solving

nonlinear optimization problems. At each iteration, the algorithm constructs linear approximations of the objective function and constraints using linear interpolation at n + 1 points in the space of the optimization variables. The worst-performing point is identified based on the original, not approximated objective function. Using the linear approximations, the algorithm then formulates and solves a linear optimization problem within a small radius around this point to update its value. The COBYLA algorithm is particularly well-suited for nonlinear optimization problems with a relatively small number of variables, especially in cases where computing derivatives is challenging or infeasible. These features make COBYLA an ideal choice for our analysis. Our study compares how the error due to initial uncertainties evolves over time for two different models, the RC and the ZC model. The RC is trained on one-dimensional indices, including the NINO3 index, the mean thermocline depth anomalies in two regions (5°N–5°S, 120°E–180°E in the western Pacific and 5°N–5°S, 180°E–290°E in the eastern Pacific), and the zonal surface wind speed anomalies (τ_c) over the area 5°N–5°S, 145°E–190°E. In contrast, the state vector of the ZC model consists of 2-dimensional fields of sea surface temperature, thermocline depth, oceanic and atmospheric velocities, and atmospheric geopotential. To address these differences in state vector dimensionality and ensure a fair comparison between the RC and ZC models, we have applied a distinct uniform constant perturbation to all the ZC model's SST fields in the NINO3 area, all the thermocline depth fields over the area 5°N-5°S 120°E-180°E and all the thermocline fields over the area 5°N-5°S 180°E-290°E. This has been done to change the mean values over these three areas of a specific quantity. By doing so, we also reduced the number of variables of our optimization problem to three, making the COBYLA algorithm an appropriate choice for our analysis. To validate the estimated Conditional Optimal Nonlinear Perturbations (CNOPs), we didn't use the methodology suggested by the reviewer. Instead, we evaluate error propagation resulting from applying numerous randomly chosen initial perturbations that satisfy the constraint conditions to determine whether the CNOPs correspond to the largest error growth. Furthermore, we verify whether the estimated CNOPs lie on the boundary of the sphere defined by the constraints.

Changes in the Manuscript:
In the revised manuscript, we will include a new appendix section providing a detailed description of the COBYLA algorithm. Additionally, we will provide a more detailed description on the validation of the CNOP.

5. *Compared with linear regression, it seems that the advantages of RC are not particularly significant either. What's your view on this issue?*

Author's Reply:

While we acknowledge that the RC does not drastically outperform the LR, our results demonstrate a clear advantage in adopting the RC, as its ability to capture nonlinear relationships between input variables, made possible by the use of a nonlinear activation function (the hyperbolic tangent in our study), leads to a consistent performance improvement, particularly in the supercritical regime, where non-linearities play a more prominent role. This is further supported by the fact that in this regime, model performance improves when τ_c is included during training, highlighting the importance of the nonlinear effects introduced by this variable [1]. These effects are better captured by the RC, whereas the LR can only provide a linear approximation. Furthermore, the relatively small performance gap found in this study can be attributed to the ZC model being a model of intermediate complexity in which ENSO is a weakly nonlinear phenomena (e.g. all wave dynamics in ocean and atmosphere is linear in the model). The data generated from the ZC model exhibit simpler dynamics compared to real-world observations or data from simulations with more complex General Circulation Models (GCMs). In such cases, the performance advantage of the RC over the LR is expected to be more pronounced.

Changes in the Manuscript:
In the revised manuscript, we will clarify the difference in performances between the RC and the LR in the "RC performances" section, explaining why the increase in performances is moderate.

6. *In RC, CNOP is not sensitive to the forecast duration, while the opposite is true in ZC (Fig. 5). Why is this the case and what does it imply?*

   Author's Reply:

   Figure 5 shows that the Zebiak-Cane (ZC) model is more sensitive to initial sea surface temperature (SST) perturbations at shorter lead times (3 months) and more sensitive to initial thermocline depth perturbations at longer lead times (6 and 9 months). In contrast, the RC appears to be more sensitive to SST perturbations across all lead times (3 to 9 months). As discussed in the conclusion, thermocline anomalies play a crucial role in error propagation, particularly in the ZC model, where ENSO variability is strongly influenced by the thermocline feedback. The Reservoir, however, effectively reduces sensitivity to initial thermocline perturbations, reducing error propagation.

   Changes in the Manuscript:
   We will include a better clarification of the results in the "CNOP analysis" section on page 11 of the revised manuscript.

7. *Personally, to explain the advantages of RC in ENSO prediction, the key is to focus on the extent to which RC has learned the ENSO dynamics or nonlinear behaviors.*

   Author's Reply:
   In our view, performing short-term forecasts and developing a perfect model emulator capable of capturing the long-term dynamics and statistical properties of a system are fundamentally different tasks, each requiring distinct model architectures, hyperparameter configurations, and evaluation criteria. Assessing how well a Machine Learning model captures a system's dynamics and nonlinear behaviors is more relevant when analyzing its ability to replicate the long-term characteristics and statistics of the system (in the case of ENSO, on a decadal timescale). In recent years, various Machine Learning models have demonstrated strong ENSO forecasting skills on relatively short timescales (up to 21 months) without necessarily capturing all the underlying physical processes. This suggests that their ability to achieve high short-term

predictive skill relies on different factors. Here, we define forecasts spanning a couple of years as ”short-term” compared to the decadal timescales required to assess ENSO’s long-term behavior. In this study, we focus on the RC model, specifically investigating the hypothesis that its superior predictive performance, particularly its ability to overcome the spring predictability barrier, stems from its capacity to reduce error propagation caused by initial uncertainties. This aligns with the perspective proposed in previous studies [2], where the spring predictability barrier in the ZC model was quantified in terms of sensitivity to initial perturbations.

Changes in the Manuscript:
We will make a remark on this in the revised “Summary and Discussion” section.

References

[1] Wansuo Duan, Yanshan Yu, Hui Xu, and Peng Zhao. Behaviors of nonlinearities modulating the el niño events induced by optimal precursory disturbances. *Climate Dynamics*, 40(5):1399–1413, March 2013.

[2] Mu Mu, Hui Xu, and Wansuo Duan. A kind of initial errors related to “spring predictability barrier” for el niño events in zebiak-cane model. *Geophysical Research Letters*, 34(3), 2007.