



Reconstructing the Atlantic Overturning Circulation Using Linear Machine Learning Techniques

Timothy DelSole & Douglas Nedza

To cite this article: Timothy DelSole & Douglas Nedza (2022) Reconstructing the Atlantic Overturning Circulation Using Linear Machine Learning Techniques, Atmosphere-Ocean, 60:5, 541-553, DOI: [10.1080/07055900.2021.1947181](https://doi.org/10.1080/07055900.2021.1947181)

To link to this article: <https://doi.org/10.1080/07055900.2021.1947181>



© 2021 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group



Published online: 12 Aug 2021.



Submit your article to this journal [↗](#)



Article views: 1241





View related articles [↗](#)



View Crossmark data [↗](#)

Reconstructing the Atlantic Overturning Circulation Using Linear Machine Learning Techniques

Timothy DelSole ^{1,2,*} and Douglas Nedza ¹

¹*Department of Atmospheric, Oceanic, and Earth Sciences, George Mason University, Fairfax, Virginia, USA*

²*Center for Ocean-Land-Atmospheric Sciences, Fairfax, Virginia, USA*

[Original manuscript received 15 December 2020; accepted 24 May 2021]

ABSTRACT *This paper examines the potential of reconstructing the Atlantic Meridional Overturning Circulation (AMOC) using surface data and linear machine learning algorithms. The algorithms are trained on pre-industrial control simulations with the aim of finding an algorithm that can reconstruct the AMOC robustly across multiple climate models. Predictors include a combination of surface temperature and surface salinity, as well as a combination of simultaneous and lagged values relative to the AMOC. For most climate models, the correlation skill of the AMOC reconstructions is greater than 0.7. This reconstruction model involves thousands of predictors and is therefore difficult to interpret. To improve interpretability, machine learning algorithms were applied to Laplacian eigenvectors, which are an orthogonal set of spatial patterns that can be ordered from largest to smallest spatial scale. The skill of the new algorithms is comparable to that based on gridded data, but the new algorithms have the advantage that dimension reduction can be more meaningfully interpreted. The most important predictors were simultaneous and lagged time series of area-averaged surface temperature, and a pattern that measures the east–west salinity difference over the basin surface lagged in time. These three predictors could recover a substantial fraction of the total skill from machine learning algorithms for most climate models.*

RÉSUMÉ [Traduit par la rédaction] *Le présent article examine le potentiel de reconstruction de la circulation méridienne de retournement de l'Atlantique (AMOC) en utilisant des données de surface et des algorithmes linéaires d'apprentissage automatique. Les algorithmes sont entraînés sur des simulations de contrôle préindustrielles dans le but de trouver un algorithme capable de reconstruire l'AMOC de façon robuste à travers plusieurs modèles climatiques. Les prédicteurs comprennent une combinaison de la température de surface et de la salinité de surface, et une combinaison de valeurs simultanées et décalées par rapport à l'AMOC. Pour la plupart des modèles climatiques, la compétence de corrélation des reconstructions de l'AMOC est supérieure à 0,7. Ce modèle de reconstruction comprend des milliers de prédicteurs et est donc difficile à interpréter. Pour améliorer l'interprétabilité, des algorithmes d'apprentissage automatique ont été appliqués aux vecteurs propres du laplacien, qui constituent un ensemble orthogonal de motifs spatiaux pouvant être classés de la plus grande à la plus petite échelle spatiale. Les compétences des nouveaux algorithmes sont comparables à celles basées sur des données rectangulaires, mais présentent l'avantage que la réduction de la dimension peut être interprétée de manière plus significative. Les prédicteurs les plus importants étaient des séries chronologiques simultanées et décalées de la température de surface moyenne de la zone, et un modèle qui mesure la différence de salinité est-ouest à la surface du bassin, décalé dans le temps. Ces trois prédicteurs pourraient récupérer une fraction substantielle de la compétence totale des algorithmes d'apprentissage automatique pour la plupart des modèles climatiques.*

KEYWORDS machine learning; AMOC reconstruction; decadal; CMIP

1 Introduction

The Atlantic Meridional Overturning Circulation (AMOC) is a component of oceanic circulation that produces basin-wide temperature variations in the North Atlantic (Delworth et al.,

1993; Zhang et al., 2019). Coupled atmosphere–ocean models indicate that this circulation has significant variability on multi-decadal time scales (Zhang & Wang, 2013); hence, it is likely a key factor in monitoring and predicting multi-

*Corresponding author's email: tdelsole@gmu.edu

© 2021 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group

This is an Open Access article distributed under the terms of the Creative Commons Attribution-NonCommercial-NoDerivatives License (<http://creativecommons.org/licenses/by-nc-nd/4.0/>), which permits non-commercial re-use, distribution, and reproduction in any medium, provided the original work is properly cited, and is not altered, transformed, or built upon in any way.

TABLE 1. CMIP5 models and variables analyzed in this study.

Model	SST	SSS
CanESM2	✓	✓
CCSM4	✓	✓
CESM1-BGC	✓	✓
CNRM-CM5	✓	✓
inmcm4	✓	✓
MPI-ESM-LR	✓	✓
MPI-ESM-MR	✓	✓
MPI-ESM-P	✓	✓
MRI-CGCM3	✓	no
NorESM1-M	✓	no

decadal climate variations. Unfortunately, direct observations are either too short (e.g., RAPID) or too infrequent (e.g., hydrographic sections) to constrain the variability of the AMOC on decadal time scales (Frajka-Williams et al., 2019). Accordingly, various attempts have been made to reconstruct the AMOC using other observations, especially sea surface temperature (SST). Most of these reconstructions are based on a linear function of a small number of predictors, such as the area-average SST in the North Atlantic, typically between 40° and 60°N (Latif et al., 2004; Rahmstorf et al., 2015), or a few empirical orthogonal functions of subsurface temperature and sea surface height (Mahajan et al., 2011; Zhang, 2008). Some non-linear reconstruction models based on a small number of predictors have also been explored (Ayala-Solares et al., 2018).

A more comprehensive approach to estimating the AMOC circulation is to derive it from ocean data assimilation. Unfortunately, the AMOC inferred in ocean data assimilation depends strongly on the ocean model used in the assimilation (Karspeck et al., 2017). The advent of new, high-quality observations from floats and satellites has led to reconstructions based on multiple datasets but only for the most recent two decades or so (Wunsch & Heimbach, 2006).

Machine learning techniques have emerged as a promising new approach to learning predictive relations from large datasets. It is natural to apply these techniques to the problem of reconstructing the AMOC. However, as mentioned above, observational datasets are problematic because they are relatively short or infrequent. In addition, observations contain variability from both internal variability and external forcing, and their relative importance in the Atlantic is much debated (Booth et al., 2012; Zhang et al., 2013). As a first step, we apply machine learning techniques to reconstruct the AMOC based on thousands of surface variables simulated by coupled atmosphere–ocean models. Time series from model simulations are attractive because they are long (e.g., 500 years), free of observational errors, and free of forced variability (when desired). A disadvantage of model simulations is that models are imperfect and differ from each other. To avoid being biased toward a particular model, we pursue a multi-model approach. Specifically, the quality of a reconstruction is judged by how well it recovers the AMOC in climate models that were not used to train the

machine learning algorithm. By finding an algorithm that skillfully reconstructs the AMOC across multiple climate models, we hope to derive an algorithm that can be trusted to reconstruct the AMOC from observations. Actual application to observations requires a separate study to separate forced and internal variability.

2 Data

The datasets used in this study come from climate model simulations. Climate models predict the behaviour of the planetary atmosphere, ocean, land, and ice based on the laws of physics and chemistry. The partial differential equations derived from these laws are too difficult to solve, so they are discretized on a grid and supplemented by parameterizations to account for processes that operate on scales much shorter than the grid (e.g., clouds and radiation). Even after these simplifications, the resulting equations involve over a million variables, are non-linear, and support chaos, so solutions cannot be obtained analytically and must be obtained by numerical approximations on a computer. The particular simulations used in this study are pre-industrial control runs from phase five of the Coupled Model Intercomparison Project (CMIP5; Taylor et al., 2012). Control simulations provide realizations of internal variability, which refers to variability that occurs naturally in the absence of year-to-year changes in climate forcing. In particular, climate forcings, such as greenhouse gas concentration (e.g., CO₂) and solar irradiance, are held constant from year to year. Different models are initialized at different random initial states, so the resulting simulations are independent, and no correlation exists between control simulations. Because only control simulations are analyzed our reconstructions account only for internal variability. Forced variability (such as from inter-annual changes in greenhouse gas concentration and solar irradiance) will be taken into account in a future study.

Our goal is to reconstruct the AMOC on decadal time scales, which is a critical time scale for near-term climate change prediction (Kirtman et al., 2013). A standard approach is to analyze ten-year running means of the AMOC, but for some models the autocorrelation function at ten years is at or near a negative minimum (not shown), in which case a ten-year mean would average out the corresponding oscillatory signal. Accordingly, we analyze running means based on a five-year window, which is shorter than ten years but still sufficient to filter out seasonal and year-to-year variations. Each five-year running mean time series is defined such that the value at year t equals the mean over years $(t - 2, t - 1, t, t + 1, t + 2)$. When considering lagged predictors, the smallest lag that avoids overlapping windows is five years. Therefore, we consider simultaneous and five-year lagged values of smoothed time series.

Following standard practice, we define an index of AMOC as the maximum meridional overturning streamfunction at a particular latitude. The streamfunction is computed as the northward velocity integrated across the basin and down to the sea

bottom, typically measured in Sverdrups ($10^6 \text{ m}^3 \text{ s}^{-1}$). We chose 40°N because this latitude is often used to study decadal AMOC variability (Buckley & Marshall, 2016; Zhang, 2008).

On decadal time scales, variability in AMOC is related to buoyancy anomalies (Buckley & Marshall, 2016). Because buoyancy is a function of temperature and salinity, we consider predictors based on SST and sea surface salinity (SSS). More precisely, the reconstructions are derived from five-year running means of annual-mean SST and SSS in the Atlantic basin between 0° and 60°N . All surface data were interpolated onto a $2.5^\circ \times 2.5^\circ$ grid, giving 518 cells in the North Atlantic.

We consider models that have pre-industrial control simulations spanning at least 500 years and include monthly mean SST and meridional overturning circulation as output variables. Only ten models in the CMIP5 dataset satisfy these criteria. Of these, eight models also contain SSS. Only the last 500 years of each pre-industrial control run were analyzed. A list of models and available variables is given in Table 1.

To mitigate model differences, the AMOC index from each model is standardized to zero mean and to unit variance. Also, the SST time series at each grid point is centred and then multiplied by the same model-dependent constant at each grid point. The constant for each model is chosen so that the area-averaged variance in that model equals one. Thus, if $T'(x, y)$ denotes a temperature anomaly, then we define a re-scaled variable $T^* = \alpha T'$ such that $(T^*)^2 = 1$, where the bar denotes the area average over the Atlantic basin between 0° and 60°N . This re-scaling preserves the ratio of variances between grid points and requires estimating only one scaling constant per model; SSS is similarly centred and re-scaled for each model.

3 Methodology

Among the many statistical models that could be used to reconstruct the AMOC, linear models often provide good approximations and useful benchmarks for non-linear models. Accordingly, we consider reconstructions based on linear models of the form

$$\mathbf{y} = \mathbf{X}\boldsymbol{\beta} + \boldsymbol{\epsilon}, \tag{1}$$

where \mathbf{y} is an N -dimensional vector of the AMOC time series, \mathbf{X} is an $N \times P$ matrix of predictors, $\boldsymbol{\beta}$ is a P -dimensional vector of regression coefficients, and $\boldsymbol{\epsilon}$ is a random term representing error in the reconstruction. The predictors \mathbf{X} are time series of either the SST field, SSS field, or both.

To mitigate overfitting, we regularize the regression problem by imposing constraints on the regression coefficients. Specifically, we employ a class of regularizations in which the norm of the regression coefficients is constrained




Fig. 1 Mean square error (red dotted curve) and associated standard errors (error bars) of the reconstructed AMOC index from LASSO, as estimated from the cross-validation procedure described in the text (the precise models and predictors are not important). The λ giving the minimum MSE is indicated by the left vertical dotted line. The λ selected by the “one-standard-error rule” is indicated by the right vertical dotted line. The numbers at the top indicate the number of predictors with non-zero coefficients in the LASSO model.

to be less than a fixed value c . That is, we solve

$$\min_{\boldsymbol{\beta}} \{ \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|^2 \} \quad \text{subject to} \quad R(\boldsymbol{\beta}) \leq c, \tag{2}$$

where c is a tunable parameter and R is defined below. We call this class of methods *linear* machine learning techniques. We consider two choices for the regularization function $R(\boldsymbol{\beta})$, namely the L_1 norm

$$R(\boldsymbol{\beta}) = \sum_p |\beta_p|, \tag{3}$$

which leads to the Least Absolute Shrinkage and Selection Operator (LASSO), and the L_2 norm

$$R(\boldsymbol{\beta}) = \sum_p \beta_p^2, \tag{4}$$

which leads to ridge regression (Hastie et al., 2009). The above optimization problem is equivalent to

$$\min_{\boldsymbol{\beta}} \{ \|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|^2 + \lambda R(\boldsymbol{\beta}) \}, \tag{5}$$

where λ is a regularization parameter that is related to c in Eq. (2). The term $\|\mathbf{y} - \mathbf{X}\boldsymbol{\beta}\|^2$ is a goodness-of-fit measure that decreases as the fit improves, while the term $R(\boldsymbol{\beta})$ is a penalty function that grows with the norm of the regression coefficients. The regularization parameter λ controls the overall magnitude of the penalty function. The above optimization problems are solved using the `cv.glmnet` package in R (Friedman et al., 2010).

Let M be the total number of models under consideration ($M = 10$ for SST-only predictors, or $M = 8$ for SST and

TABLE 2. Configurations of regularized regression discussed in this paper.

Config.	Regression	Predictor	Lagged Predictors	Number of Predictors
1	Ridge	SST	Yes	1036
2	Ridge	SST	No	518
3	Ridge	SSS	Yes	1036
4	Ridge	SST & SSS	Yes	2072
5	LASSO	SST	Yes	1036
6	LASSO	SST	No	518
7	LASSO	SSS	Yes	1036
8	LASSO	SST & SSS	Yes	2072

SSS predictors). One model is identified as the test model and the remaining models are used for training. We call this approach Leave-One-Model-Out (LOMO). After picking a test model, $M - 1$ models remain for training. Our criterion for selecting λ is as follows. For a given value of λ , one model from the $M - 1$ models is withheld, the remaining $M - 2$ models are used to train the regression model, and then the regression model is applied to the withheld model for validation. This procedure is repeated, using each model from the $M - 1$ models in turn for validation to generate a regression prediction for each time step for each of the $M - 1$ models. There are a total of $N = 4500$ time steps for SST, or a total of $N = 3500$ time steps for SST and SSS predictors. Letting ϵ_n denote the (cross-validated) error at the n th time step, the mean square error (MSE) and its standard error are

$$\text{MSE} = \frac{1}{N} \sum_{n=1}^N \epsilon_n^2, \tag{6}$$

$$\text{Standard Error} = \sqrt{\left(\frac{1}{N-1} \sum_{n=1}^N (\epsilon_n^2 - \text{MSE})^2\right) / N}. \tag{7}$$

An example from the above procedure is shown in Fig. 1. The minimum MSE is indicated by the left vertical dotted line. Sometimes, the minimum is shallow and its location depends sensitively on the sample. We use a more stable selection rule known as the one-standard-error rule, in which the “most parsimonious” model (i.e., the model with larger λ) whose MSE is within a standard error of the minimum MSE is selected (Hastie et al., 2009). This λ is indicated by the right vertical dotted line.

After selecting λ in the above manner, the regression model is retrained using all $M - 1$ training models, and the corresponding β is used to reconstruct the AMOC in the test model. Note that no information from the test model is used to inform the regression model. The error of the reconstruction is measured by

$$\hat{\epsilon} = \mathbf{y} - \mathbf{X}\beta, \tag{8}$$

where β comes from $M - 1$ training climate models, and

(\mathbf{y}, \mathbf{X}) comes from the test climate model. When $\text{MSE} > 1$, it means that the MSE is greater than that based on the climatological mean, in which case the reconstruction is said to have no skill. The mean and standard deviation of the squared elements of $\hat{\epsilon}$ are computed as in Eqs (6) and (7), but only over the $N = 500$ time steps of the test model.

4 Results: gridded data

We have examined regularized regressions for a large number of configurations. For this paper, we consider the configurations indicated in Table 2. For a single variable at a single lag, the number of predictors is just the number of grid points in the North Atlantic, namely 518. If the instantaneous values of both SST and SSS are the predictors, then the number of predictors is twice the number of grid points (e.g., 518 points for SST and SSS). Similarly, if a single variable and its time-lagged values are included as predictors, then the number of predictors is also twice the number of grid points. If the predictors include both SST and SSS, each with simultaneous and lagged values, then the number of predictors is four times the number of grid points. The number of grid points in each case is listed in Table 2.

We found that reconstructions were better if SST and SSS predictors lagged AMOC (i.e., the “predictor” value is in the future relative to the AMOC). This makes sense physically: an acceleration of the AMOC would advect warmer and more saline waters from the tropics toward the poles, leading to a positive anomaly in SST and SSS following the acceleration of the AMOC. In the following, we present results only for combinations of simultaneous or lagged predictors. This means that these AMOC reconstructions are not *predictions* but *reconstructions* based on current and subsequent data.

We first discuss a reconstruction based on one of the configurations in Table 2. This will illustrate the degree to which machine learning methods can reconstruct the AMOC. One of the best reconstructions is obtained using both SST and SSS predictors, each with simultaneous and lagged values (configuration #4 in Table 2). Time series of AMOC reconstructions in individual climate models are shown in Fig. 2. The match between actual and reconstructed AMOCs is remarkable given that no data from a climate model are used as training data to reconstruct the AMOC in that climate model. For readers who want a more quantitative argument, the precise p -value is difficult to quantify because of serial correlation, but a crude approach is to estimate the degrees of freedom using the “effective number of independent samples” (Thiébaux & Zwiers, 1984),

$$N_{\text{effective}} = N \left(1 + 2 \sum_{\tau=1}^{50} \hat{\rho}_{\tau} \right)^{-1},$$

where $\hat{\rho}_{\tau}$ is the sample autocorrelation at τ years, $N = 500$,

Reconstructions of AMOC

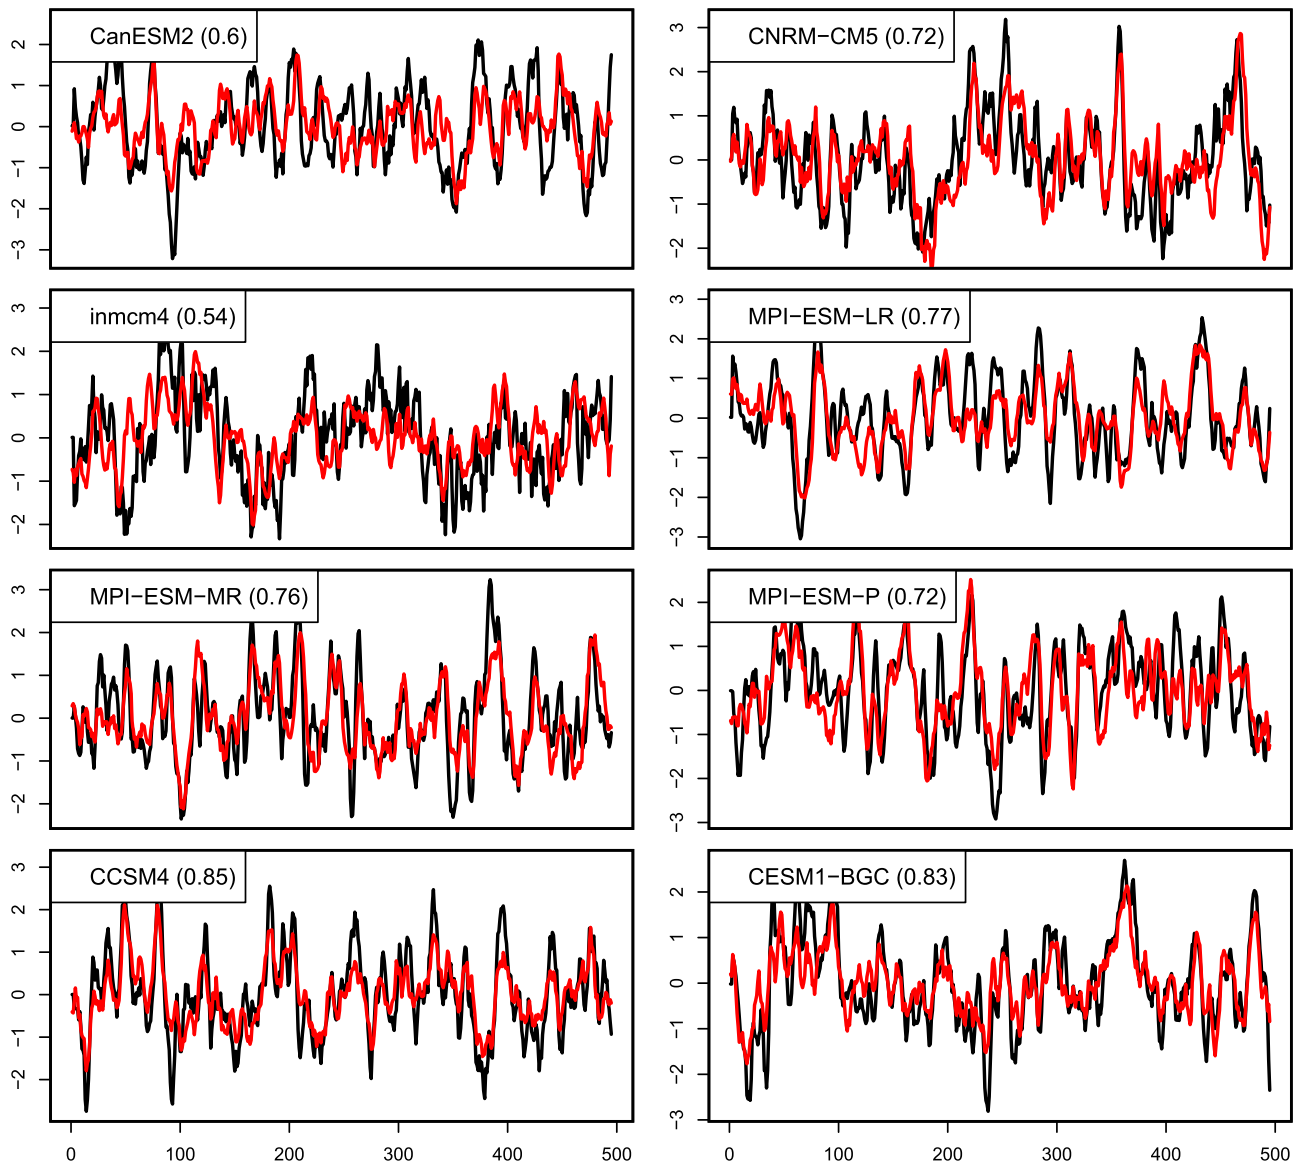


Fig. 2 The AMOC index in each model (black) and the reconstructed AMOC index (red) based on SST and SSS, each with simultaneous and five-year lagged values (corresponding to configuration #4 in Table 2). For each model, the reconstruction algorithm is trained on all other models. The correlation between the two time series is indicated in parentheses next to the model name.

and the upper limit of the sum is cut off at 50 years; $N_{\text{effective}}$ depends on the model, but the smallest value is found to be 32, which corresponds to a correlation 5% significance threshold of 0.35. The correlations shown in Fig. 2 each exceed this value, indicating that the probability of exceeding this value by random chance for all eight models is very small.

The regression coefficients for configuration #4 are shown in Fig. 3 for the case in which MPI-ESM-LR is left out of training. Admittedly, the structure of the coefficients is difficult to interpret, which is not an uncommon situation in machine learning. In general, interpreting regression coefficients is not straightforward because the value of one

coefficient depends on all the others. Furthermore, the spatial structure of the coefficients is sensitive to the choice of model that is left out. In the following, we attempt to understand how the reconstruction works by changing the configuration and examining the resulting change in skill.

The error statistics for different configurations and test models are shown in Fig. 4. As a reminder, the regression model is trained on all but one model, then the resulting regression coefficients are used to predict the AMOC in the withheld model, and the error statistics shown in Fig. 4 are computed only over the withheld model. Each algorithm is based on ridge regression using gridded data of either SST,

Regression Coefficients for Reconstructing AMOC

RIDGE; grid; MPI-ESM-LR

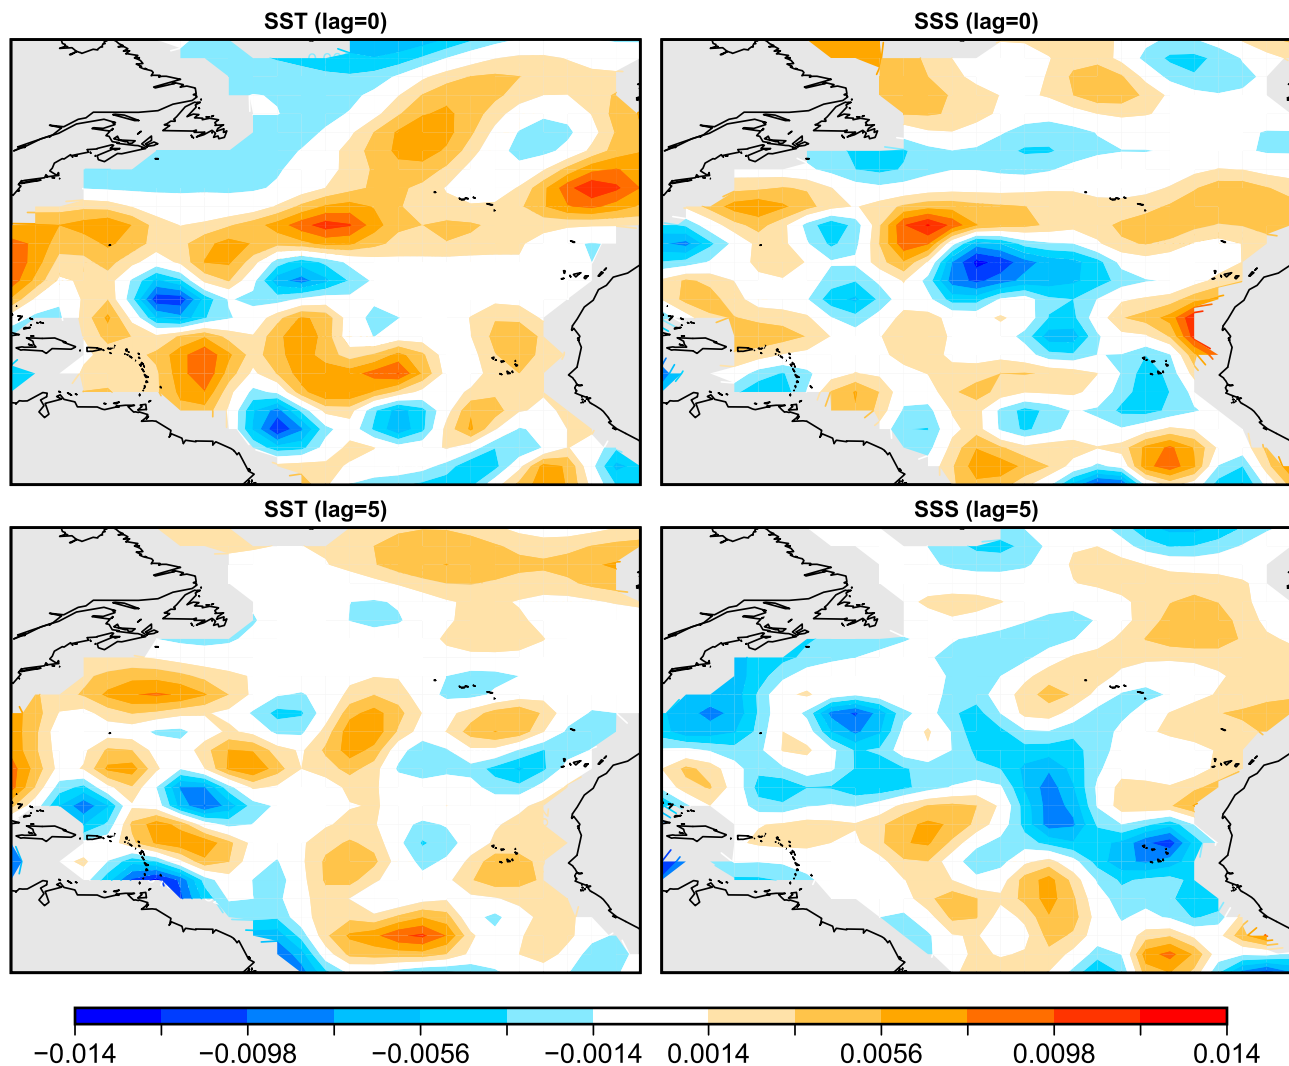


Fig. 3 Regression coefficients in the North Atlantic (0° - 60° N) for configuration #4 in Table 2, trained on all models except MPI-ESM-LR. The coefficients are non-dimensional because of variable normalization.

SSS, or both. The results for LASSO are nearly the same and therefore not shown. It is immediately apparent that the amount of variance explained by the reconstruction varies considerably with validation model. Adding salinity predictors to temperature does not always improve the reconstruction in a cross-model-validation sense (i.e., the blue error bars are not always below the red error bars). On the other hand, reconstructions based on SST only tend to be better than those based on SSS only (i.e., red error bars tend to lie below the gold error bars), and adding time-lagged predictors tends to improve the reconstruction (i.e., red error bars tend to lie below the green error bars).

The best reconstructions (in terms of correlation skill or mean square error) were obtained for the two NCAR models, namely CCSM4 and CESM1-BGC. One might wonder if this is because the two models are very similar. To check this, we

recomputed the above results while leaving out CESM1-BGC from the entire analysis. In this case, the skill of the reconstruction for CCSM4 dropped; the correlation dropped from 0.85 to 0.76, while the mean square error increased from 0.30 to 0.45. The reconstruction skill for the remaining models was nearly the same. This result is consistent with other studies that have shown that models from the same institution tend to behave similarly and therefore do not “diversify” training data as much as adding models from different institutions (Knutti et al., 2013).

5 Interpreting the model

In our application, the regression model (Eq. (1)) has thousands of predictors and is, therefore, difficult to interpret. We seek a comparable reconstruction model that is easier to

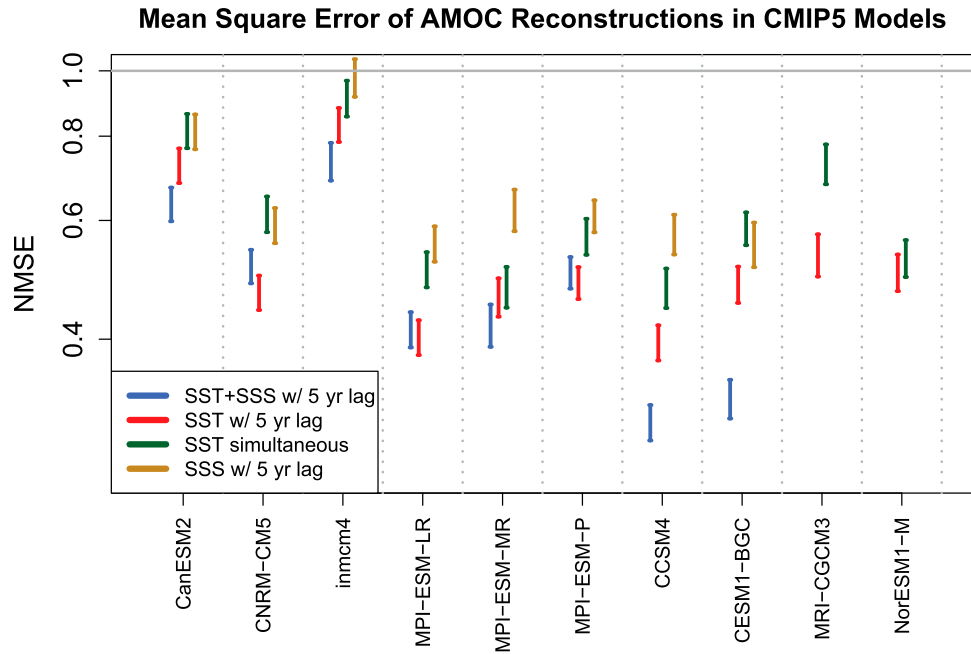


Fig. 4 Mean square error of reconstructing the AMOC index in the CMIP5 model indicated on the x -axis. Training data are based on all models except the model on the x -axis whose AMOC is being reconstructed (LOMO). Error bars show $MSE \pm$ one standard error. The reconstructions are derived from ridge regression and the gridpoint predictors are SST and SSS, each with simultaneous and five-year lagged values (blue); simultaneous and five-year lagged SST (red); simultaneous-only SST (green); simultaneous and five-year lagged SSS (gold).

interpret. That is, we seek a comparable model with just a few predictors. The LASSO method is useful for this purpose because it sets some regression coefficients exactly to zero, which unambiguously identifies the corresponding predictors as unimportant. In our application, this means that LASSO assigns non-zero regression coefficients only to isolated grid points, yielding a spotty regression map (not shown). Unfortunately, the specific grid points chosen by LASSO differ for different test models. This is not surprising because SST and SSS are spatially correlated fields, so very similar predictions can be obtained from predictors that are displaced by only a few degrees of latitude or longitude relative to each other.

An alternative approach is to use large-scale patterns as predictors. With this aim, we used a novel approach based on the eigenvectors of the Laplace operator. The eigenvectors of the Laplace operator are a complete set of orthogonal vectors that can be ordered by a measure of length scale. Familiar examples of Laplacian eigenvectors include the spherical harmonics and the sines and cosines of a Fourier series. In this paper, we use eigenvectors of the Laplace operator over the Atlantic Ocean basin between 0° and 60°N , computed using the method of DelSole and Tippett (2015). The resulting leading Laplacian eigenvectors are shown in Fig. 5. The first eigenvector is a uniform pattern, and projecting data onto the first Laplacian eigenvector is equivalent to taking the area-weighted average in the basin. In the case of SST, the time series for the first Laplacian eigenvector is merely an index of Atlantic Multidecadal Variability (AMV). The second and third eigenvectors are dipoles that measure the

large-scale gradient across the basin. Subsequent eigenvectors capture smaller scale patterns. Following Laprise (1992), we define the characteristic length scale of the m th Laplacian eigenvector as $\theta_m = \pi a/k_m$, where a is the radius of the Earth, and k_m is the associated total wavenumber. A total of $M = 100$ Laplacian eigenvectors are chosen, which corresponds to resolving length scales greater than about 500 km. The time series for an eigenvector is obtained by projecting model output onto the eigenvector.

To impose smooth structure, the coefficients of small-scale eigenvectors should be penalized more than coefficients for large-scale eigenvectors. Such scale-selective penalization can be incorporated through generalized norms. For instance, a generalized L_1 norm could be

$$R(\boldsymbol{\beta}) = \sum_m w_m |\beta_m|, \quad (9)$$

where w_1, \dots, w_m are suitable positive weights. Penalizing smaller-scale structures more than larger-scale structures corresponds to choosing weights w_m that increase with decreasing θ_m . To explore specific forms for this relation, we first examine the variance spectrum of the Laplacian eigenvectors. Figure 6 shows the percentage of variance of SST and SSS explained by Laplacian eigenfunctions in each model as a function of length scale. Just as is found in the atmosphere (Nastrom & Gage, 1985), the variance of temperature tends to be concentrated in large-scale structures and decreases with length scale according to a -3 power law (this appears to be the first

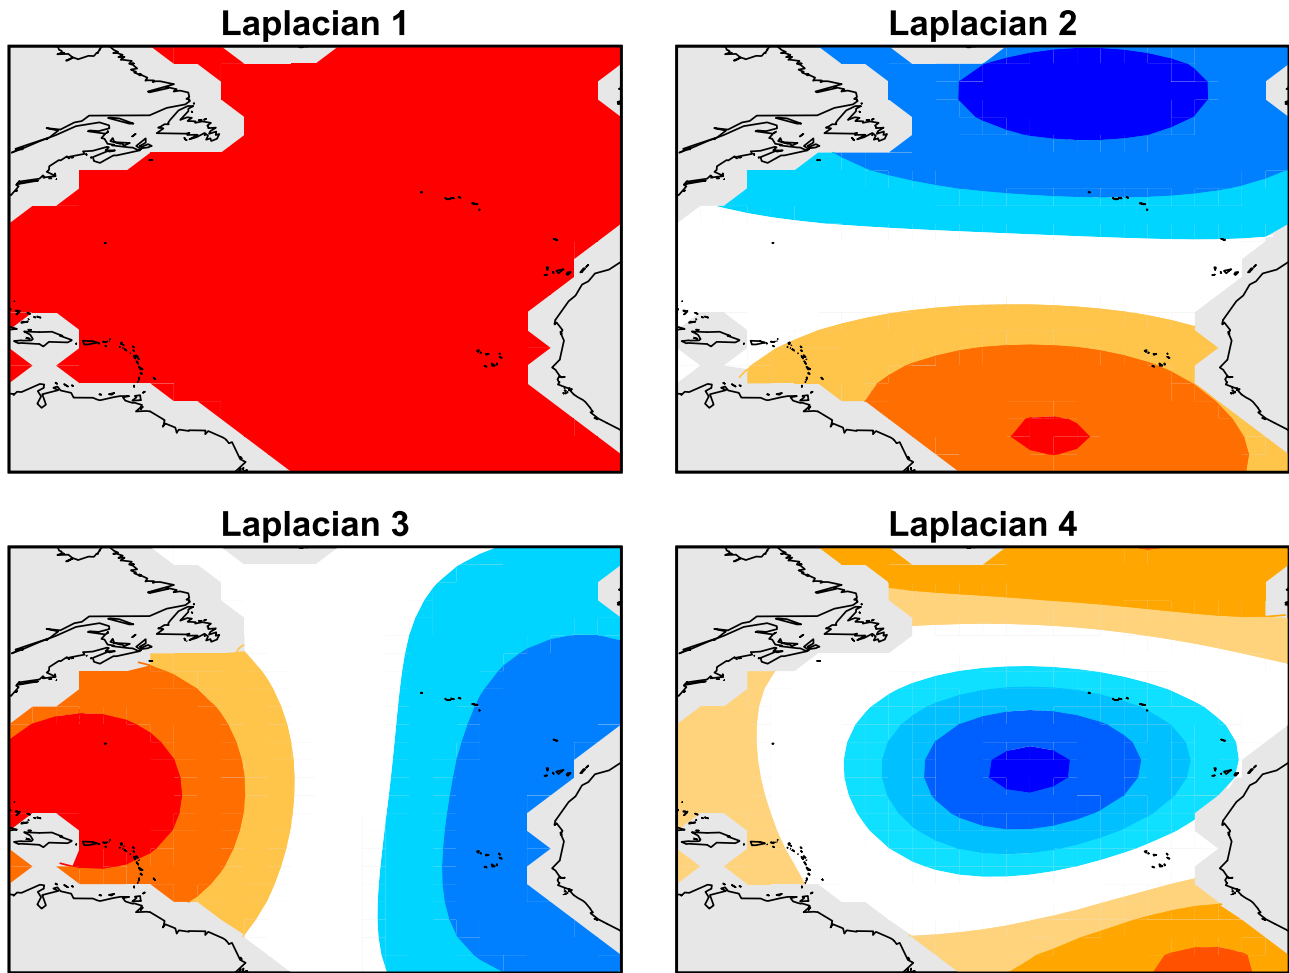


Fig. 5 Laplacian eigenvectors 1, 2, 3, and 4 over the North Atlantic between the equator and 60°N, where dark red and dark blue indicate extreme positive and negative values, respectively.

confirmation of this power law in realistic ocean models). Thus, the decrease in variance with decreasing length scale provides a natural penalty function. However, in most LASSO and ridge regression packages (e.g., *cv.glmnet*), each predictor is standardized by its standard deviation prior to solving the minimization problem (Eq. (5)). As a result of this standardization, the coefficients for the first and the hundredth are penalized equally, which undercuts our goal to penalize small scales more than large scales. To penalize small-scale structures more than large-scale structures, the weights w_m could be set equal to the inverse standard deviation, which would depend

on the model. To avoid model-dependent weights, we use the empirical function

$$f(\theta) = \frac{8}{(4000/\theta)^3 + 7}, \tag{10}$$

which is shown as the blue curve in Fig. 6 (after normalizing by the total sum). As can be seen, the empirical function gives a reasonable fit to the model-generated variance spectra. We have also explored powers of $f(\theta)$, including the square root, corresponding to standard deviation, but these did not give substantially different results. Accordingly, we show results only for $w_m = 1/f(\theta_m)$. With this choice, $w_{100}/w_1 \approx 100$ (i.e., the coefficient for the 100th eigenfunction is penalized about 100 times more strongly than that of the first). We call this approach a scale-selective regularization.

Under the one-standard error rule described in Section 3, LASSO typically selects 30 or more Laplacian eigenvectors (not shown), which is still too many to interpret. To rank the predictors in order of importance, we count the number of times each predictor is selected for λ values greater than or

TABLE 3. Most important predictors in LASSO using Laplacian eigenvectors under LOMO and scale-selective regularization. The predictors are ranked by the number of times each one is selected for λ values greater than or equal to the value selected by cross-validation.

Rank	Variable	Lag (yrs)	Laplacian
1	SST	0	1
2	SST	5	1
3	SSS	5	3

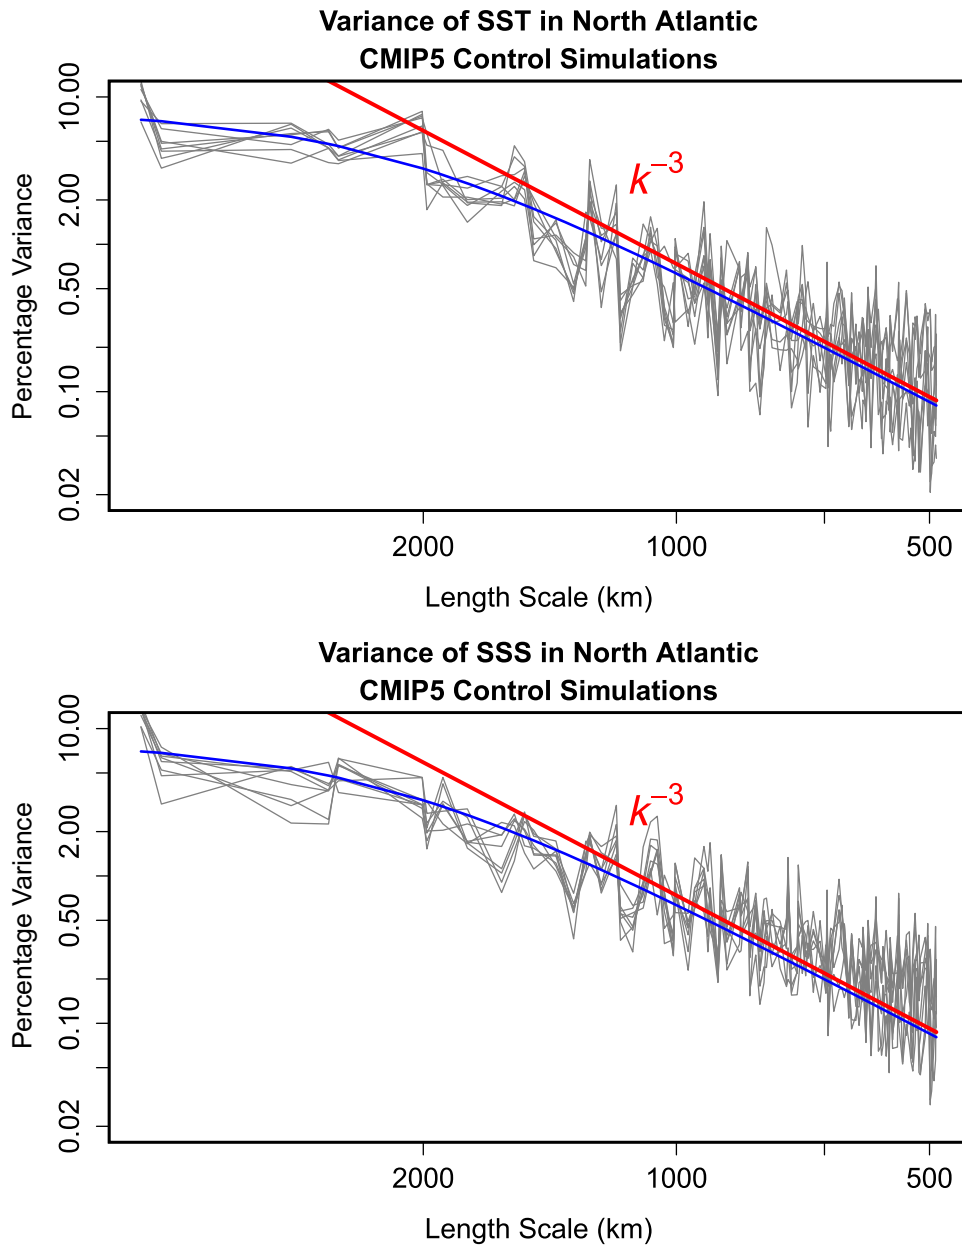


Fig. 6 Temporal variance of Laplacian eigenvectors for SST (top) and SSS (bottom) over the North Atlantic from CMIP5 pre-industrial control runs (grey curves). For reference, a line corresponding to the “ k^{-3} power law” is shown in red (k is the total wavenumber), and the blue curve shows the empirical function Eq. (10).

equal to the selected value. For instance, in the case of Fig. 1, we count the number of times LASSO selects a given predictor for each of the regression models to the right of the selected λ . The top three predictors are shown in Table 3 and correspond to the first Laplacian eigenvector for SST at zero and a five-year lag, and the third Laplacian eigenvector for SSS at a five-year lag. The first two are essentially AMV indices. The third is an index of the east–west gradient of salinity. To gain insight into why this latter pattern was chosen, we show in Fig. 7 the regression pattern between AMOC and the five-year lagged SSS. Each pattern exhibits an east–west gradient, with some models preferring a positive eastern signal and others (e.g.,

NCAR models) preferring a negative western signal. Such dipole structures in the AMOC–SSS relation are commonly seen in similar analyses (Cheng et al., 2013, Fig. 6) (Zhang et al., 2019, their Fig. 9).

Without scale-selective regularization, the top ten predictors include some high-order Laplacian eigenvectors (e.g., the 95th eigenvector; not shown), but the skill is nearly the same as with scale-selective regularization.

Having identified the important predictors, the associated coefficients are computed using ordinary least squares (OLS). The coefficients are computed based on all but one dynamical model and then used to reconstruct

Regression Pattern Between AMOC and SSS (lag=5)

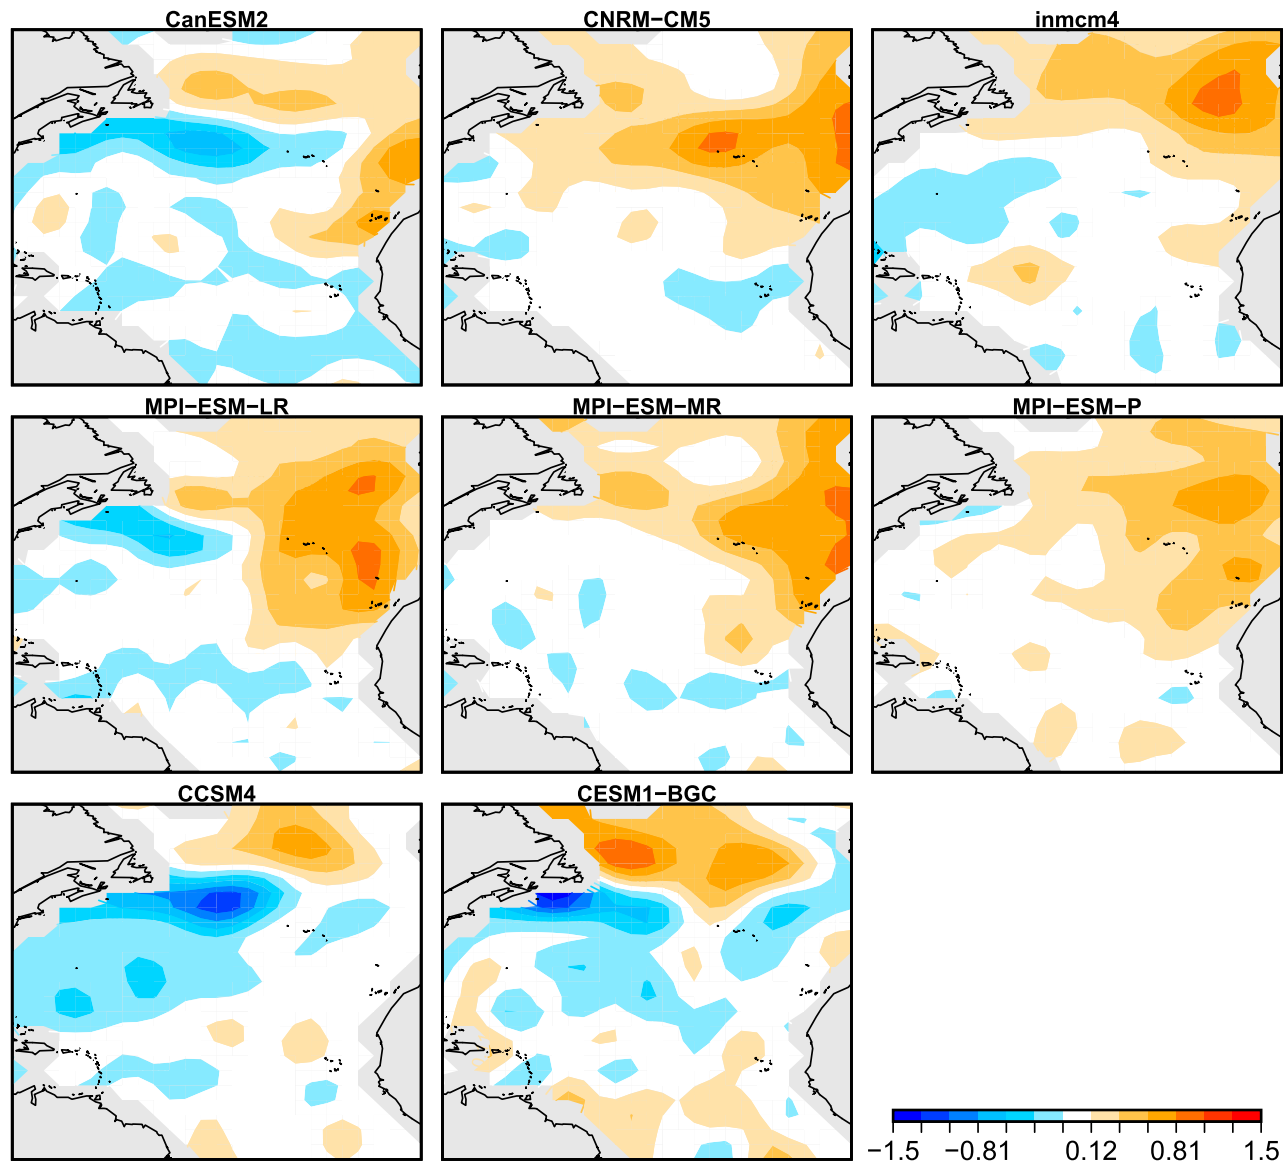


Fig. 7 Regression coefficient between the AMOC index and the local, five-year lagged SSS for each CMIP5 model. The coefficients are dimensionless because of variable normalization, but the contour levels are the same in each panel.

the AMOC in the withheld model. The MSEs of the reconstructions are shown in Fig. 8. Also shown are the reconstructions based on LOMO. For most models, a substantial fraction of the total skill in LOMO can be obtained from two to three predictors. Because these models may be of broad interest, we give the full reconstruction model explicitly:

$$\widehat{\text{AMOC}}(t) = 1.14 \text{AMV}(t) + 0.77 \text{AMV}(t+5) \quad (11)$$

$$\begin{aligned} \widehat{\text{AMOC}}(t) = & 1.02 \text{AMV}(t) + 0.59 \text{AMV}(t+5) \\ & - 0.89 \text{SSS-L3}(t+5), \end{aligned} \quad (12)$$

where $\text{AMV}(t)$ is the re-scaled amplitude of Laplacian 1, which is merely the average Atlantic SST between 0° and 60°N and can be identified as a re-scaled AMV index, and SSS-L3 is the third Laplacian eigenfunction for SSS. The coefficients are the average OLS coefficients obtained from each LOMO training. For reference, Fig. 9 shows time series of the reconstructed AMOC from Eq. (12). We remind the reader that SST and SSS have been rescaled separately for each model so that the sum total variance in the North Atlantic equals one. We further remind the reader that the predictors include lagged values of AMV and SSS, which involve “future” information relative to the AMOC time series. The OLS




Fig. 8 Mean square errors of AMOC reconstructions as a function of validation model (x -axis). For each validation model, the first three dots joined by lines show reconstructions from ordinary least squares by successively including the three predictors in Table 3. The second and third dots correspond to Eqs (11) and (12), except model-dependent coefficients are used in the figure. The square at the end of each curve shows the mean square error of LOMO using LASSO with 100 simultaneous and 100 lagged Laplacian eigenvectors of temperature and salinity.

model performance is close to that of LASSO except for the two NCAR models. In the latter cases, even using the top ten predictors in OLS does not substantially improve the reconstruction.

6 Conclusions

This paper examined machine learning algorithms for reconstructing an AMOC index based on surface data. The algorithms were trained on pre-industrial control simulations from CMIP5 models. As such, the algorithms only account for internal variability. Our goal was to find a reconstruction algorithm that works well in models that were not used for training. Accordingly, we explored several different configurations including LASSO and ridge regression, predictors based on SST and SSS, and simultaneous and lagged predictors. There was a general tendency for out-of-model skill to be best using both simultaneous and lagged values of SST and SSS. Time series from this reconstruction algorithm capture much of the multi-decadal variance of the AMOC although the fraction of variance explained depends on the CMIP5 model used for validation.

The spatial structure of regression coefficients is generally difficult to interpret. To obtain a simpler model, we derived regression models based on a small number of Laplacian eigenvectors. The most important Laplacian eigenvectors were identified by counting how often LASSO selected them over a range of regularization parameters. The first two predictors were simultaneous and lagged values of AMV indices, and the third predictor was the lagged SSS projected onto the third Laplacian, which measures the east–west gradient across the Atlantic basin. These three predictors could recover a substantial fraction of the total skill from regularized regression for most CMIP5 models.

Ultimately, our goal is to reconstruct the AMOC based on observational data. Without any strong reason to believe one dynamical model is better than another, we desire a reconstruction algorithm that works well across all the dynamical models. We found that LOMO yielded a reconstruction that was always skilful in the withheld model. From the simplified models, Eqs. (11) and (12), one can anticipate that a reconstruction derived from observational data would be a smoothed version of the AMV index. However, the observed AMV index also contains a forced component that should be removed before reconstructing the AMOC based on our model. Unfortunately, the precise forced component of the AMOC is a matter of debate and will be considered in future work.

To our knowledge, past reconstructions have not utilized lagged predictors and have not combined temperature and salinity data. Our results suggest that including such information can improve reconstructions of the AMOC.

Acknowledgements

We thank Martha Buckley for comments on an earlier draft that led to a more thorough and clearer presentation. We acknowledge the World Climate Research Programmes Working Group on Coupled Modelling, which is responsible for CMIP, and we thank the climate modelling groups (listed in Table 1 of this paper) for producing and making available their model output.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Reconstructions of AMOC

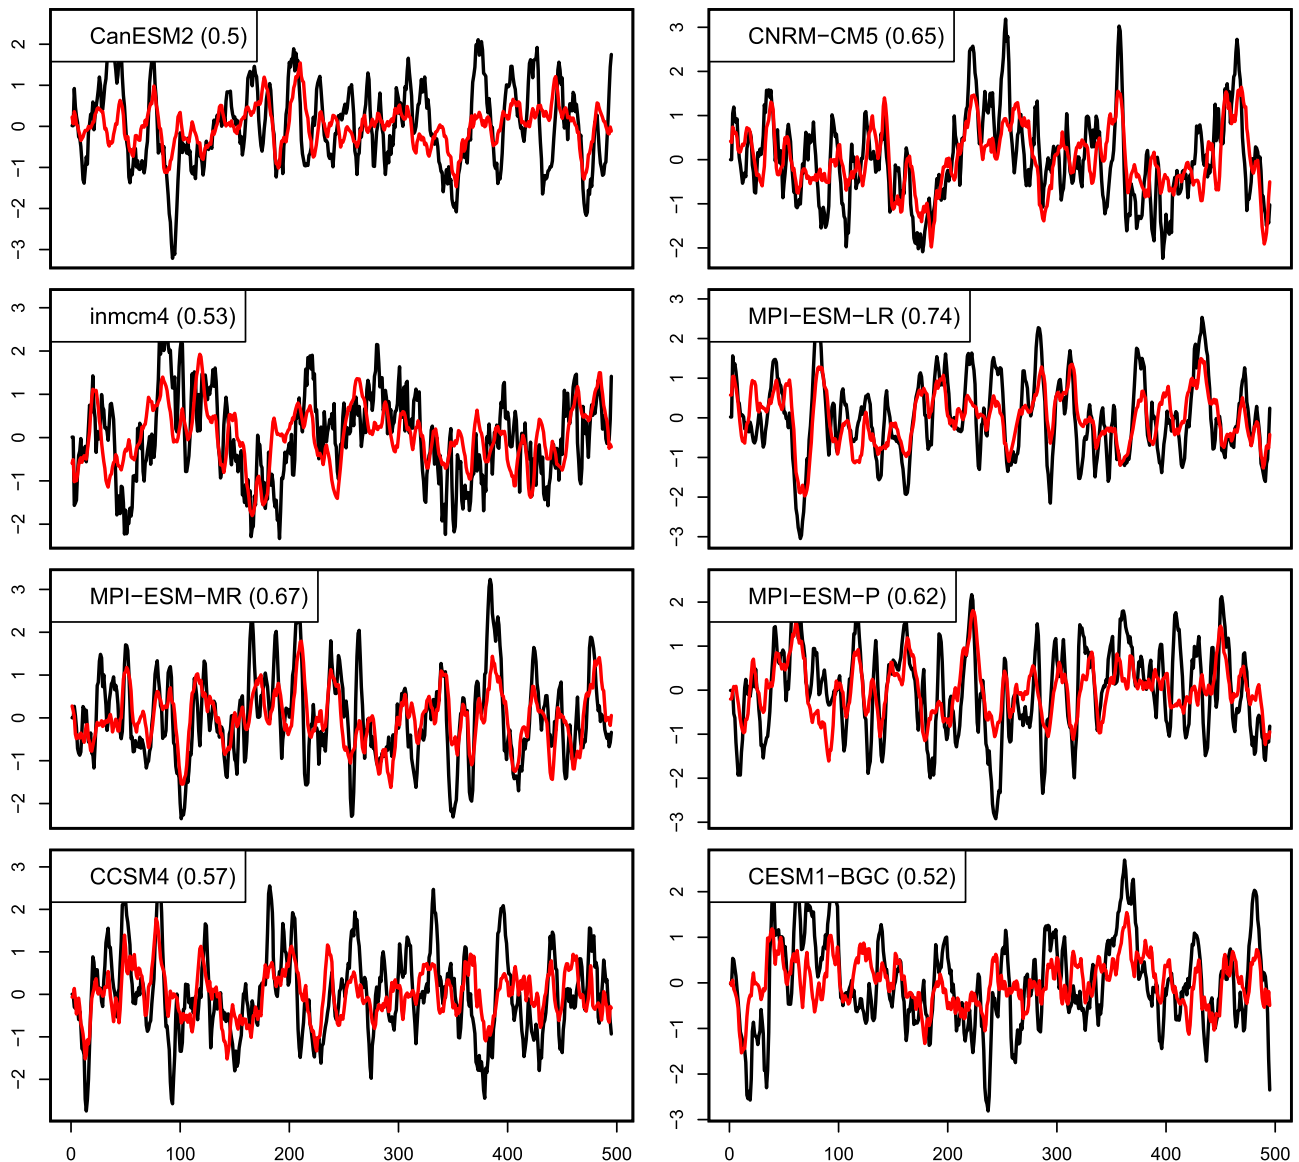


Fig. 9 As in Fig. 2, but for AMOC reconstructions based on the three-predictor model (Eq. (12)).

Funding

This research was supported by the National Oceanic and Atmospheric Administration, Climate Program Office [grant number NA16OAR4310175]. The views expressed herein are those of the authors and do not necessarily reflect the views of the funding agencies.

ORCID

Timothy DelSole  <http://orcid.org/0000-0003-2041-3024>

Douglas Nedza  <http://orcid.org/0000-0003-1492-6675>

References

- Ayala-Solares, J. R., Wei, H.-L., & Bigg, G. R. (2018). The variability of the Atlantic Meridional Circulation since 1980, as hindcast by a data-driven nonlinear systems model. *Acta Geophysica*, 66(4), 683–695. <https://doi.org/10.1007/s11600-018-0165-7>
- Booth, B. B. B., Dunstone, N. J., Halloran, P. R., Andrews, T., & Bellouin, N. (2012). Aerosols implicated as a prime driver of twentieth-century North Atlantic climate variability. *Nature*, 484(7393), 228–232. <https://doi.org/10.1038/nature10946>

- Buckley, M. W., & Marshall, J. (2016). Observations, inferences, and mechanisms of the Atlantic Meridional Overturning Circulation: A review. *Reviews of Geophysics*, 54(1), 5–63. <https://doi.org/10.1002/2015RG000493>
- Cheng, W., Chiang, J. C. H., & Zhang, D. (2013). Atlantic Meridional Overturning Circulation (AMOC) in CMIP5 models: RCP and historical simulations. *Journal of Climate*, 26(18), 7187–7197. <https://doi.org/10.1175/JCLI-D-12-00496.1>
- DelSole, T., & Tippett, M. K. (2015). Laplacian eigenfunctions for climate analysis. *Journal of Climate*, 28(18), 7420–7436. <https://doi.org/10.1175/JCLI-D-15-0049.1>
- Delworth, T., Manabe, S., & Stouffer, R. J. (1993). Interdecadal variations of the thermohaline circulation in a coupled ocean-atmosphere model. *Journal of Climate*, 6(11), 1993–2011. [https://doi.org/10.1175/1520-0442\(1993\)006<1993:IVOTTC>2.0.CO;2](https://doi.org/10.1175/1520-0442(1993)006<1993:IVOTTC>2.0.CO;2)
- Frajka-Williams, E., Anson, I. J., Baehr, J., Bryden, H. L., Chidichimo, M. P., Cunningham, S. A., Danabasoglu, G., Dong, S., Donohue, K. A., Elipot, S., Heimbach, P., Holliday, N. P., Hummels, R., Jackson, L. C., Karstensen, J., Lankhorst, M., Le Bras, I. A., Lozier, M. S., McDonagh, E. L., ... Wilson, C. (2019). Atlantic Meridional Overturning Circulation: Observed transport and variability. *Frontiers in Marine Science*, 6, 260. <https://doi.org/10.3389/fmars.2019.00260>
- Friedman, J., Hastie, T., & Tibshirani, R. (2010). Regularization paths for generalized linear models via coordinate descent. *Journal of Statistical Software, Articles*, 33(1), 1–22. <https://doi.org/10.18637/jss.v033.i01>
- Hastie, T., Tibshirani, R., & Friedman, J. H. (2009). *Elements of statistical learning* (2nd ed.). Springer.
- Karspeck, A. R., Stammer, D., Köhl, A., Danabasoglu, G., Balmaseda, M., Smith, D. M., Fujii, Y., Zhang, S., Giese, B., Tsujino, H., & Rosati, A. (2017). Comparison of the Atlantic Meridional Overturning Circulation between 1960 and 2007 in six ocean reanalysis products. *Climate Dynamics*, 49(3), 957–982. <https://doi.org/10.1007/s00382-015-2787-7>
- Kirtman, B., Power, S. B., Adedoyin, J. A., Boer, G. J., Bojariu, R., Camilloni, I., Doblas-Reyes, F. J., Fiore, A. M., Kimoto, M., Meehl, G. A., Prather, M., Sarr, A., Schar, C., Sutton, R., van Oldenborgh, G. J., Vecchi, G., & Wang, H. (2013). Near-term climate change: Projections and predictability. In T. Stocker, D. Qin, G.-K. Plattner, M. Tignor, S. Allen, J. Boschung, A. Nauels, Y. Xia, V. Bex, & P. Midgley (Eds.), *Climate Change 2013: The physical science basis. Contribution of Working Group I to the fifth assessment report of the Intergovernmental Panel on Climate Change* (Chapter 11, pp. 953–1028). Cambridge University Press.
- Knutti, R., Masson, D., & Gettelman, A. (2013). Climate model genealogy: Generation CMIP5 and how we got there. *Geophysical Research Letters*, 40(6), 1194–1199. <https://doi.org/10.1002/grl.50256>
- Laprise, R. (1992). The resolution of global spectral models. *Bulletin of the American Meteorological Society*, 73(9), 1453–1455. <https://doi.org/10.1175/1520-0477-73.9.1453>
- Latif, M., Roeckner, E., Botzet, M., Esch, M., Haak, H., Hagemann, S., Jungclaus, J., Legutke, S., Marsland, S., Mikolajewicz, U., & Mitchell, J. (2004). Reconstructing, monitoring, and predicting multi-decadal-scale changes in the North Atlantic thermohaline circulation with sea surface temperature. *Journal of Climate*, 17(7), 1605–1614. [https://doi.org/10.1175/1520-0442\(2004\)017<1605:RMAPMC>2.0.CO;2](https://doi.org/10.1175/1520-0442(2004)017<1605:RMAPMC>2.0.CO;2)
- Mahajan, S., Zhang, R., Delworth, T. L., Zhang, S., Rosati, A. J., & Chang, Y.-S. (2011). Predicting Atlantic Meridional Overturning Circulation (AMOC) variations using subsurface and surface fingerprints. *Deep Sea Research Part II: Topical Studies in Oceanography*, 58(17–18), 1895–1903. <https://doi.org/10.1016/j.dsr2.2010.10.067>
- Nastrom, G. D., & Gage, K. S. (1985). A climatology of atmospheric wavenumber spectra of wind and temperature observed by commercial aircraft. *Journal of the Atmospheric Sciences*, 42(9), 950–960. [https://doi.org/10.1175/1520-0469\(1985\)042<0950:ACOWS>2.0.CO;2](https://doi.org/10.1175/1520-0469(1985)042<0950:ACOWS>2.0.CO;2)
- Rahmstorf, S., Box, J. E., Feulner, G., Mann, M. E., Robinson, A., Rutherford, S., & Schaffernicht, E. J. (2015). Exceptional twentieth-century slowdown in Atlantic Ocean overturning circulation. *Nature Climate Change*, 5(5), 475–480. <https://doi.org/10.1038/nclimate2554>
- Taylor, K. E., Stouffer, R. J., & Meehl, G. A. (2012). An overview of CMIP5 and the experiment design. *Bulletin of the American Meteorological Society*, 93(4), 485–498. <https://doi.org/10.1175/BAMS-D-11-00094.1>
- Thiébaux, H. J., & Zwiers, F. W. (1984). The interpretation and estimation of effective sample size. *Journal of Climate and Applied Meteorology*, 23(5), 800–811. [https://doi.org/10.1175/1520-0450\(1984\)023<0800:TIAEOE>2.0.CO;2](https://doi.org/10.1175/1520-0450(1984)023<0800:TIAEOE>2.0.CO;2)
- Wunsch, C., & Heimbach, P. (2006). Estimated decadal changes in the North Atlantic Meridional Overturning Circulation and heat flux 1993–2004. *Journal of Physical Oceanography*, 36(11), 2012–2024. <https://doi.org/10.1175/JPO2957.1>
- Zhang, L., & Wang, C. (2013). Multidecadal North Atlantic sea surface temperature and Atlantic Meridional Overturning Circulation variability in CMIP5 historical simulations. *Journal of Geophysical Research: Oceans*, 118(10), 5772–5791. <https://doi.org/10.1002/jgrc.20390>
- Zhang, R. (2008). Coherent surface-subsurface fingerprint of the Atlantic Meridional Overturning Circulation. *Geophysical Research Letters*, 35, L20705. <https://doi.org/10.1029/2008GL035463>
- Zhang, R., Delworth, T. L., Sutton, R., Hodson, D. L. R., Dixon, K. W., Held, I. M., Kushnir, Y., Marshall, J., Ming, Y., Msadek, R., Robson, J., Rosati, A. J., Ting, M., & Vecchi, G. A. (2013). Have aerosols caused the observed Atlantic multidecadal variability? *Journal of the Atmospheric Sciences*, 70(4), 1135–1144. <https://doi.org/10.1175/JAS-D-12-0331.1>
- Zhang, R., Sutton, R., Danabasoglu, G., Kwon, Y.-O., Marsh, R., Yeager, S. G., Amrhein, D. E., & Little, C. M. (2019). A review of the role of the Atlantic Meridional Overturning Circulation in Atlantic multidecadal variability and associated climate impacts. *Reviews of Geophysics*, 57(2), 316–375. <https://doi.org/10.1029/2019RG000644>