A Bayesian Approach to Regional Decadal Predictability: Sparse Parameter Estimation in High-Dimensional Linear Inverse Models of High-Latitude Sea Surface Temperature Variability

DALLAS FOSTER

Oregon State University, Corvallis, Oregon, and Computational Physics and Methods (CCS-2), Los Alamos National Laboratory, Los Alamos, New Mexico

DARIN COMEAU AND NATHAN M. URBAN

Computational Physics and Methods (CCS-2), Los Alamos National Laboratory, Los Alamos, New Mexico

(Manuscript received 14 October 2019, in final form 3 March 2020)

ABSTRACT

Stochastic reduced models are an important tool in climate systems whose many spatial and temporal scales cannot be fully discretized or underlying physics may not be fully accounted for. One form of reduced model, the linear inverse model (LIM), has been widely used for regional climate predictability studies—typically focusing more on tropical or midlatitude studies. However, most LIM fitting techniques rely on point estimation techniques deriving from fluctuation–dissipation theory. In this methodological study we explore the use of Bayesian inference techniques for LIM parameter estimation of sea surface temperature (SST), to quantify the skillful decadal predictability of Bayesian LIM models at high latitudes. We show that Bayesian methods, when compared to traditional point estimation methods for LIM-type models, provide better calibrated probabilistic skill, while simultaneously providing better point estimates due to the regularization effect of the prior distribution in high-dimensional problems. We compare the effect of several priors, as well as maximum likelihood estimates, on 1) estimating parameter values on a perfect model experiment and 2) producing calibrated 1-yr SST anomaly forecast distributions using a preindustrial control run of the Community Earth System Model (CESM). Finally, we employ a host of probabilistic skill metrics to determine the extent to which an LIM can forecast SST anomalies at high latitudes. We find that the choice of prior distribution has an appreciable impact on estimation outcomes, and priors that emphasize physically relevant properties enhance the model's ability to capture variability of SST anomalies.

1. Introduction

Interannual variability in sea surface temperatures (SSTs) is an important indicator of the variability of the global climate system and has long been studied as relating to other atmospheric variables (Davis 1976; McKinnon and Deser 2018; Révelard et al. 2018). Forecasting variations in SSTs themselves has been an area of open research: there is interest not only in the time scale of predictability (Davis 1976; Branstator and Teng 2010; Guemas et al. 2012; Stock et al. 2015), but also in regional phenomena such as El Niño–Southern Oscillation (ENSO) and the Pacific decadal oscillation (PDO) (Thomas et al. 2018; DiNezio et al. 2017; Penland 1996; Penland et al. 1993; Hawkins and Sutton 2009;

Wittenberg et al. 2014), which can have global impacts through atmospheric teleconnections. Stochastic reduced models, like the linear inverse model (LIM), have been widely used for regional climate predictability studies. An LIM is a linear stochastic differential equation model for the evolution of a climate field, such as SSTs. Often, as is done in this paper, the empirical orthogonal functions (EOFs) of said field are modeled [see Penland (1989) for a discussion of the use of EOFs in LIMs]. The most prominent application of LIMs are to subannual prediction of SSTs for ENSO forecasting (Penland et al. 1993; Penland and Sardeshmukh 1995), but there are also applications to decadal time scales and midlatitude regions (Alexander et al. 2008; Newman 2007, 2013; Dias et al. 2019; DelSole et al. 2013) and extensions to systems with nonlinear and memory effects (Kravtsov et al. 2009; Kondrashov et al. 2015; Chen et al. 2016), as

Corresponding author: Dallas Foster, fostdall@oregonstate.edu

DOI: 10.1175/JCLI-D-19-0769.1

© 2020 American Meteorological Society. For information regarding reuse of this content and general copyright information, consult the [AMS Copyright Policy](https://www.ametsoc.org/PUBSReuseLicenses) (www.ametsoc.org/PUBSReuseLicenses).

well as a variety of other contexts (Hawkins and Sutton 2009; Penland and Matrosova 2001; Wu et al. 2018; Martinez-Villalobos et al. 2018; Alexander et al. 2008; Barnston et al. 1999).

The high latitudes have long been the subject of numerous climate predictability studies, with Arctic sea ice in particular garnering considerable attention. While practical needs drive predictions of Arctic sea ice on subseasonal to seasonal time scales, many studies have also pushed the forecast horizons to interannual to decadal time scales (Guemas et al. 2016, and references therein). In general, interest in decadal climate predictions has been increasing over the last several years (Meehl et al. 2014), including focus on SSTs (Hawkins et al. 2011) and high latitudes, particularly in the Southern Ocean (Latif et al. 2017; Zhang et al. 2017), where interannual- to decadal-scale variability plays a strong role in oceanic forcing on the Antarctic ice sheet, particularly to the vulnerable West Antarctic Ice Sheet (Jenkins et al. 2016).

The main focus in this paper is the implementation, calibration, and evaluation of the LIM framework as applied to the forecasting of high-latitude SST anomalies using modern Bayesian statistical strategies and probabilistic scoring. We make use of the 1800-yr fully coupled preindustrial control run of the Community Earth System Model version 1 (CESM1) Large Ensemble project (Kay et al. 2015), chosen to ensure that there were enough data to perform validation of the methodology and to avoid complicating factors of non-stationarity, data sparsity, or bias. We adopt a Bayesian perspective because of the regularization effect that the use of prior distribution have in high-dimensional estimation problems (Hastie et al. 2017; Vogel 2002) and the ability to produce better probabilistic calibration when compared to point estimation (Samaniego 2010).

The parameters of an LIM (i.e., the drift and diffusion matrices) must be inferred from data. The inference of the drift term is traditionally (e.g., in Penland 1989) achieved by using fluctuation–dissipation theory (FDT) relating the matrices and the data covariances. Estimation of the drift and incremental noise covariance matrices in this way is also equivalent to the maximum likelihood (ML) method as also shown in Penland (1989) and presented in section 3. To simplify the estimation problem, the dynamics of the system are assumed linear and Markovian so that the eigenfunctions—or so-called principal oscillation patterns (POPs)—of the propagating matrices can be analyzed for spatial patterns (von Storch and Bruns 1988; Hasselmann 1988). This technique produces point estimates for the parameters which allows for scientists to infer the physical characteristics of the EOFs from the system parameters (i.e.,

the spectrum of the drift matrix) and to forecast climate statistics using a reduced model. ML and frequentist approaches can be used to obtain sample statistics like confidence intervals, and skewness and kurtosis estimates, as in Martinez-Villalobos et al. (2018), which provide insight into parameter uncertainty, but we argue that the Bayesian perspective provides a complete framework to specify prior information and propagate parameter uncertainty through the forecasting model.

To facilitate the main focus of this paper, the secondary objective is to consider the LIM problem from a Bayesian perspective and explore implementation strategies. Bayesian parameter estimation diverts from prevailing frequentist statistical estimation by defining the parameters as random variables with probability distributions and the subsequent role of diverse prior probabilities given to the estimands. For a brief review of Bayesian techniques in parameter estimation see Gelman et al. (2013). Notably, filtering approaches, as, for example, in Hansen and Penland (2007) and Zhao et al. (2019), offer convenient approaches to state and parameter estimation at the cost of the assumption of Gaussian or particle filter distributions. The filtering approach can be beneficial for time-dependent parameters, but may also have difficulty with filter divergence or with exploring the entire posterior probability space (Särkkä 2010). Because the LIM parameters are not time dependent, we have chosen in this paper a Markov chain Monte Carlo variant approach to estimate the posterior distribution of the parameters.

In our application there are two primary motivations for a Bayesian approach to parameter estimation. First, we argue that the likelihood distribution induced by an LIM is relatively weak, so that the posterior distribution is highly sensitive to a prior distribution and this sensitivity leads to improvements in parameter recovery and forecast skill. In other words, the problem is high-dimensional and the relative amount of information per parameter is relatively low. The introduction of prior densities in the Bayesian framework is a way to add additional information to, or regularize, the parameter distributions. It is known that, in parameter estimation, the Bayesian estimates will outperform the frequentist estimate as long as the Bayesian prior does not have *both* large error in the mean and a relatively large weight (Samaniego 2010). Second, we argue that propagating parameter uncertainties through the model can produce better calibrated forecast distributions. In the traditional ML approach, only the point estimates are typically used to produce forecast distributions, although cross-validation methods provide another way to regularize parameter estimates and produce an ensemble of forecasts. The Bayesian approach provides a

framework and means to characterize the distribution of parameter uncertainties and their effect on the model and forecast distribution. By fully characterizing this uncertainty, we believe that the Bayesian methods detailed in this paper are better suited to reproduce observation distributions when compared to a traditional point-estimate approach.

The third thrust of this paper is a focus on the performance of the Bayesian perspective (in relation to ML/FDT) as a function of the specific prior implemented. Using nontrivial priors in time series analysis has a long history in economics, where Bayesian methods are used to estimate coefficients (often) of vector auto-regressive (VAR) models (Litterman 1986; Giannone et al. 2015; Partridge and Rickman 1998; Doan et al. 1983). There have been various attempts to introduce Bayesian learning of parameters in stochastic differential equations, but usually require some assumptions or approximations of the underlying likelihood distribution (Eraker 2001; Singer 2004; Beskos et al. 2006; Karimi and McAuley 2016; Batz et al. 2018; Tian et al. 2014; Møller et al. 2011). Introduction of nontrivial prior beliefs about the parameters can complicate the process of describing the necessary parameters since the densities resulting from stochastic differential equations do not generally have closed form expressions. Consideration of complex prior distributions in parametric models is underdeveloped, but the effect of prior densities in nonparametric modeling has been studied in van Zanten (2013).

Statistical estimation of stochastic models has a long history starting with ML and least squares estimates as a simplification; see Le Breton (1977) and DelSole and Yang (2010). In application to climate statistics, reviews can be found in DelSole and Hou (1999) and DelSole and Chang (2003) and comparisons in Mason and Mimmack (2002) and Nummelin et al. (2018). From Penland (1989), the use of FDT approximations for the LIM to derive closed-form parameter estimation formulas makes this approach the method of choice in geoscience applications (Penland 1989; Penland and Sardeshmukh 1995; Nummelin et al. 2018; Barnston et al. 1999; Colman and Davey 2003), particularly over various autoregressive models (DelSole and Chang 2003; Penland et al. 1993). Barnston et al. (2012) compares LIMs and other physically and statistically based models in the International Research Institute for Climate and Society (IRI) ENSO prediction plume. In this paper, we use a statistical framework to show that the same estimates can be achieved from maximizing a natural parameter likelihood function and that the introduction of prior distributions in a Bayesian sense is a natural extension.

The regression interpretation of LIMs gives us reasoning to believe why regularization is necessary. The Gauss–Markov theorem ensures that the maximum likelihood estimates for LIMs have the smallest variance of all unbiased estimates; it is known in general that these estimates can be improved by various methods of regularization like ridge regression, LASSO (least absolute shrinkage and selection operator), and cross-validation that reduce estimator variance (Hastie et al. 2017). The need for regularization may be the result of the presence of confounding variables or model error and the lack of data relative to the number of parameters. In fact, the use of a Bayesian prior can be seen as a type of regularization that aims to address these limitations.

In what follows we briefly discuss the theoretical underpinnings and results of LIMs, ML estimation for these models, and Bayesian methods for parameter estimation. The aim of this paper is then to provide insight into the use of a Bayesian framework with nontrivial prior distributions and how they can be used to generate well-calibrated SST anomaly forecasts. To this point, in the results section, we give two experiments to compare the use of ML and Bayesian estimation methods. First, in a perfect model methodology, we compare convergence of parameter estimations. Second, we apply these methods to an LIM of SST anomalies and analyze the forecast performance of each using a variety of traditional and probabilistic skill metrics. We find that the use of Bayesian methods with informative priors is significantly more skillful in forecasts and can better estimate the underlying probability distribution of the data.

2. Linear inverse modeling

Consider the linear stochastic differential equation

$$d\mathbf{x}_t = \mathbf{B}\mathbf{x}_t dt + d\mathbf{W}_t, \quad (1)$$

where $\mathbf{x} \in \mathbb{R}^m$, $\mathbf{B} \in \mathbb{R}^{m \times m}$, and $\mathbf{W}$ is an m -dimension Wiener process with zero mean and covariance given by $\mathbf{Q} \in \mathbb{R}^{m \times m}$. We denote the covariance of the process $\mathbf{x}$: $\langle \mathbf{x}\mathbf{x}^T \rangle = \mathbf{\Lambda}$. The values of $\mathbf{B}$ or $\mathbf{Q}$ are not known a priori and the fitting of an LIM requires the estimation of these $m^2 + (1/2)m(m+1)$ total parameters, since $\mathbf{Q}$ is symmetric.

The solution to (1) at a time $t + \tau$ is given by

$$\mathbf{x}_{t+\tau} = \mathbf{G}(\tau)\mathbf{x}_t + \boldsymbol{\xi}_t, \quad \boldsymbol{\xi}_t \sim \mathcal{N}(0, \boldsymbol{\Sigma}(\tau)), \quad (2)$$

where

$$\mathbf{G}(\tau) = e^{\mathbf{B}\tau} \quad (3)$$

FIG. 1. A stochastic process with known $\mathbf{B}$ and $\mathbf{Q}$ was generated by discretizing (1) using Euler–Maruyama with time step dt 1000 observations from that series were sampled with $\tau = 100dt$ to form $\mathbf{y}_{1:T}$. (left to right): trace plots of the marginal likelihood $p(\mathbf{y}_{1:T}|\mathbf{Q})$, marginal likelihood $p(\mathbf{y}_{1:T}|\mathbf{Q})$ along with prior normal $p(\mathbf{Q})$ and resulting posterior distribution $p(\mathbf{Q}|\mathbf{y}_{1:T})$, marginal likelihood $p(\mathbf{y}_{1:T}|\mathbf{B})$, and the marginal likelihood $p(\mathbf{y}_{1:T}|\mathbf{B})$ along with prior normal $p(\mathbf{B})$ and resulting posterior distribution $p(\mathbf{B}|\mathbf{y}_{1:T})$ are shown. Maximum likelihood and maximum a posteriori are shown for comparison between frequentist and Bayesian point estimates while standard deviations of the distributions are shown to compare uncertainties in the parameters from the Bayesian perspective.

is referred to as the propagating matrix of lead time τ , but in the literature it is also called the Green's function for this differential equation; the incremental noise process satisfies the relation

$$\Sigma(\tau) = \Lambda - \mathbf{G}(\tau)\Lambda\mathbf{G}(\tau), \quad (4)$$

and the fluctuation–dissipation relation gives

$$\mathbf{B}^T\Lambda + \Lambda^T\mathbf{B} = -\mathbf{Q} \quad (5)$$

[see Penland (1989) for detailed calculations]. Both $\mathbf{Q}$ and Σ must be symmetric positive definite, while $\mathbf{B}$ must have eigenvalues with negative real part in order for the solution (2) to be stable. See Penland and Sardeshmukh (1995) for a lengthier discussion on tests for the applicability of the assumptions of LIMs and resulting approximations.

3. Maximum likelihood estimation

Now consider observations of the random process $\mathbf{y}_n = \mathbf{x}(t_n)$ for $n = 1, \dots, T$, abbreviated as $\mathbf{y}_{1:T}$ when considering a range of observations. We assume that the observation process is exact, but the introduction of observational noise can be accommodated with only minor adjustments to the methods detailed below. The conditional probability density function, typically called the conditional marginal likelihood, for these observations conditioned on matrices $\mathbf{G}$ and Σ defined above is Gaussian and given by

$$p(\mathbf{y}_{1:T}|\mathbf{G}, \Sigma) = \prod_{n=1}^{T-\tau} \mathcal{N}(\mathbf{y}_{n+\tau} - \mathbf{G}(\tau)\mathbf{y}_n, \Sigma(\tau)). \quad (6)$$

Taking the derivative of the marginal likelihood with respect to $\mathbf{G}(\tau)$ and $\Sigma(\tau)$, we can derive formulas for the so-called ML estimates

$$\hat{\mathbf{G}}(\tau) = \left(\sum_{n=1}^{T-\tau} \mathbf{y}_{n+\tau} \mathbf{y}_n^T \right) \left(\sum_{n=1}^{T-\tau} \mathbf{y}_n \mathbf{y}_n^T \right)^{-1}, \quad (7)$$

$$\begin{aligned} \hat{\Sigma}(\tau) &= \frac{1}{T-\tau-1} \sum_{n=1}^{T-\tau} (\mathbf{y}_{n+\tau} - e^{\mathbf{B}\tau} \mathbf{y}_n)(\mathbf{y}_{n+\tau} - e^{\mathbf{B}\tau} \mathbf{y}_n)^T \\ &= \langle \mathbf{y}_{\tau:T} \mathbf{y}_{\tau:T}^T \rangle - \hat{\mathbf{G}}(\tau) \langle \mathbf{y}_{1:T-\tau} \mathbf{y}_{1:T-\tau}^T \rangle \hat{\mathbf{G}}(\tau)^T \end{aligned} \quad (8)$$

From these matrices, one calculates the matrices $\hat{\mathbf{B}}$ using (3) and $\hat{\mathbf{Q}}$ using (5) by approximating Λ with ensemble averages. While $\hat{\mathbf{Q}}$ is formally obtained from the fluctuation–dissipation relation, we will also refer to this as the ML estimate of $\mathbf{Q}$ since it is obtained from using the ML estimate of $\mathbf{B}$. Note that (8) is approximately equal to (4) under the assumption of stationarity.

As a numerical example, consider the one-dimensional case. In Fig. 1 a stochastic process $\mathbf{x}_t$ was generated using chosen values $\mathbf{B} = -3$ and $\mathbf{Q} = 0.8$ and sampled to generate 1000 observations $\mathbf{y}_k$. We then compute $p(\mathbf{y}_{1:T}|\mathbf{G}, \Sigma)$, where we simply use $\tau = 1$. For the scalar case, trace plots of $p(\mathbf{y}_{1:T}|\mathbf{B}, \mathbf{Q})$ can be obtained from $p(\mathbf{y}_{1:T}|\mathbf{G}, \Sigma)$ since Σ and $\mathbf{Q}$ have the simple relationship $\Sigma = (e^{2\mathbf{B}\tau} - 1)\mathbf{Q}/(2\mathbf{B})$ from solving (4) and (5). The traceplots for $\mathbf{Q}$ and $\mathbf{B}$ are displayed in the first and third image respectively, adjoined by their Bayesian estimates as described further in section 4a. The location of the maximum of these trace plots give maximum likelihood estimates (MLEs) $\hat{\mathbf{B}}$, and $\hat{\mathbf{Q}}$, whose experimental relative error was found to be approximately 146%, and 6%, respectively, while the relative error for

$\hat{\mathbf{G}}$ was approximately 0.5%. Since the transformation $\mathbf{B} = \log(\mathbf{G})/\tau$ increases the variance of the marginal likelihood and amplifies the relative error, it is conceivable from this experiment alone that reasonable forecasts results can be achieved with $\hat{\mathbf{G}}$ while the numerical value of $\hat{\mathbf{B}}$ differs significantly from the “true” value $\mathbf{B}$. Additionally, observe that the marginal likelihood places significant probability on values of $\mathbf{B}$ that are positive, something we know leads to unstable dynamics and is physically impossible. Regularization of the marginal likelihood can help constrain estimates from undesirable regions; the Bayesian approach can be seen as a particular justification of this regularization.

4. Bayesian methods for parameter estimation

Determining the most likely values and distributions for the LIM parameters is inherently a high-dimensional problem with hundreds, if not thousands, of parameters needing estimation. Because the ratio of the number of parameters to data points may be close to or larger than unity, regularization is needed to constrain possible parameter values and prevent overfitting. In the statistical sense, regularization provides a way to balance bias and variance in the estimates. A Bayesian probabilistic framework provides a mechanism to formally regularize the maximum likelihood distribution by incorporating prior beliefs about the uncertainty of the parameters that is equivalent to Tikhonov regularization (Vogel 2002), also known as ridge regression.

Bayesian philosophy assumes that the parameters $\theta \in \mathbb{R}^d$ are random variables with joint prior probability distribution $p(\theta)$. This prior distribution represents any a priori knowledge we have about the nature and structure of the parameters. A state space model with model states $\mathbf{x}(t_k)$, observations $\mathbf{y}_k$, and unknown parameters may be written as

$$\begin{aligned} \theta &\sim p(\theta), \\ \mathbf{x}(t_0) &\sim p[\mathbf{x}(t_0), \theta], \\ \mathbf{x}(t_n) &\sim p[\mathbf{x}(t_n) | \mathbf{x}(t_{n-1}), \theta], \\ \mathbf{y}_n &\sim p[\mathbf{y}_n | \mathbf{x}(t_n), \theta], \end{aligned} \quad (9)$$

where we label the initial uncertainty $p[\mathbf{x}(t_0), \theta]$, the transition probability density $p[\mathbf{x}(t_n) | \mathbf{x}(t_{n-1}), \theta]$, and observation uncertainty $p[\mathbf{y}_k | \mathbf{x}(t_k), \theta]$. To simplify the above relations, it is assumed that the initial condition and observation uncertainty are independent of the parameters θ . The full posterior distribution, with these assumptions, can be written using Bayes’ rule:

$$\begin{aligned} p[\mathbf{x}(t_{0:T}), \theta | \mathbf{y}_{1:T}] \\ = \frac{p[\mathbf{x}(t_0)]p(\theta)}{p(\mathbf{y}_{1:T})} \prod_{n=1}^T p[\mathbf{y}_n | \mathbf{x}(t_n)] p[\mathbf{x}(t_n) | \mathbf{x}(t_{n-1}), \theta]. \end{aligned} \quad (10)$$

Note that (10) gives the joint posterior distribution for both the state $\mathbf{x}_{0:T}$ and the parameters θ and estimating it solves not only the parameter estimation problem but also the filtering/smoothing problem. The distribution of only the parameters requires marginalizing out the state variables:

$$p(\theta | \mathbf{y}_{1:T}) = \int p[\mathbf{x}(t_{0:T}), \theta | \mathbf{y}_{1:T}] d\mathbf{x}(t_{0:T}). \quad (11)$$

Computation of this integral can be difficult, but possible using Monte Carlo methods; some possibilities are presented in the [appendix](#). We make the standard approximation that the marginal distribution is given by

$$p(\theta | \mathbf{y}_{1:T}) \propto p(\mathbf{y}_{1:T} | \theta) p(\theta). \quad (12)$$

Taking the maximum of (12) yields the maximum a posteriori (MAP) point estimate $\hat{\theta}^{\text{MAP}}$:

$$\hat{\theta}^{\text{MAP}} = \text{argmax}[p(\theta | \mathbf{y}_{1:T})]. \quad (13)$$

Point estimates, like the MAP, are only simple statistics of the underlying posterior distribution, and fail to signal the overall variability of likely parameters, but it will be useful to use $\hat{\theta}^{\text{MAP}}$ when comparing results of estimation from ML and Bayesian techniques.

From (12), it will be sufficient to consider only the prior distribution of parameters and the conditional marginal likelihood distribution (6) to form an approximate marginal posterior distribution. Having already defined the marginal likelihood, the remaining factor in estimating the posterior distribution is the choice of prior $p(\theta)$. We discuss some common choices and their properties.

Prior distributions

The prior distribution incorporates our beliefs about the nature of the parameters and their structure. These beliefs are not perfect and may be subject to their own uncertainty. The role of a prior distribution is to weakly prefer certain parameter values when no information is present. Upon inspection, the posterior distribution (10) can be seen as a competition between the likelihood and the prior, where the information from the prior becomes diluted given data. One can also view a prior as a penalty term whose role is to weakly enforce constraints onto the parameters, especially when considering a Gaussian

prior and likelihood where the prior can be viewed as a type of Tikhonov regularization.

There are various labels that have been applied in order to classify priors (see Gelman et al. 2013), but we will not rely on these terms. We distinguish and delineate different prior densities based on whether they inform scale or sign information and whether the constraint informs certain structures. We will also have cause to highlight the number of so-called hyper-parameters (additional parameters needed in order to fully specify the prior distributions). The presence of a large number of hyper-parameters can increase the flexibility of a prior distribution while also increasing the storage and computational costs of producing random samples from the posterior distribution.

We discuss several options of priors for both $\mathbf{B}$ and $\mathbf{Q}$ of varying degree type that we will perform experiments with. A simple prior would be to assign a normal distribution uniformly to all parameters:

$$\mathbf{B}_{ij} \sim \mathcal{N}(\mu, 1) \quad \mathbf{Q}_{ij} \sim \mathcal{N}(0, 1),$$

where $\mathcal{N}(\mu, 1)$ is the normal distribution with mean μ and variance 1. Choosing the normal distribution as a prior does not convey sign information and so it does not enforce $\mathbf{Q}$ to be positive definite, nor does it enforce the eigenvalues of $\mathbf{B}$ to be negative. Instead, this choice of prior is a simple and computationally efficient way to weakly enforce the center of mass of the parameters via location and scale information (i.e., it forces the values of $\mathbf{B}$ to be somewhat close to μ and values of $\mathbf{Q}$ to be small). Variations to inform sign could be the half-normal $\mathcal{N}^+(\mu, \sigma)$, half-Cauchy, $\text{Cauchy}^+(0, \gamma)$, or $\beta(\alpha, \beta)$.

Consider again the one-dimensional example problem of estimating $\mathbf{B}$ and $\mathbf{Q}$ from 1000 sampled data $\mathbf{y}$, but now from a Bayesian approach. In Fig. 1 the priors $\mathbf{Q} \sim \mathcal{N}^+(0, 1)$ and $\mathbf{B} \sim \mathcal{N}(-1, 1)$ are used in conjunction with the marginal likelihood to produce the Bayesian posterior in the second and fourth image respectively. In this example, the truth and maximum a posteriori values are calculated. For $\mathbf{B}$, the error in the MAP estimate of roughly 60% is an improvement over the MLE, whereas for $\mathbf{Q}$, the MAP estimate is marginally worse than the MLE.

Using more sophisticated combinations of distributions, we can leverage knowledge about the structures of the matrix parameters in the LIM. Decompose $\mathbf{Q}$ as

$$\mathbf{Q} = \boldsymbol{\sigma} \mathbf{L} \mathbf{L}^T \boldsymbol{\sigma}, \quad (14)$$

where $\boldsymbol{\sigma}$ is a diagonal matrix with $\sigma_{ii} = \sqrt{\mathbf{Q}_{ii}}$ and $\mathbf{L}$, a lower triangular matrix, is the Cholesky decomposition of the remaining correlation matrix. We impose simple

half-normal or half-Cauchy prior distributions on the diagonal elements of $\boldsymbol{\sigma}$. For the $m(m+1)/2$ elements of $\mathbf{L}$, a popular choice in the literature is to place an LKJ prior (Lewandowski et al. 2009; named after the authors, Lewandowski, Kurowicka, and Joe) on $\mathbf{L} \mathbf{L}^T$, whose distribution is uniform over the space of correlation matrices.

While mathematical constraints on the spectrum of $\mathbf{B}$ pose a possible source of prior information, placing priors directly on an eigenvalue decomposition of $\mathbf{B}$ is difficult because it would require specification of probabilities in the complex plane and increases the number of parameters needed to be estimated. Furthermore, the eigenvalue decomposition is sensitive to random perturbations in the input matrix to the order of the condition number of the matrix times the perturbation (Stoer and Bulirsch 2002)—enough to possibly misidentify elements of the spectrum of $\mathbf{B}$. Instead, we use distributions that indirectly produce these qualities by inducing sparsity (only few nonzero elements) and a (nearly) diagonal matrix. By inducing sparsity, we are hypothesizing that EOFs are weakly coupled and this coupling decays with the explained variance of the EOF. We focus on two types of distributions that have been used to induce this kind of sparsity to varying degrees, the Minnesota (Litterman) and (Finnish) horseshoe priors (see appendix A). The extent to which sparsity is induced is data-driven and determined by hyper-parameter values, meaning that we do not impose a strict prior belief of sparse interaction between EOFs.

The Minnesota-Litterman prior (Litterman 1986; Doan et al. 1983) prior expresses the belief that (2) should be approximately an uncoupled random walk through each dimension. This prior places unit normal priors on the diagonal of $\mathbf{G}(\tau)$ and zero normal priors on the off diagonal components. The variance of these off-diagonal components is then proportional to the variance of the underlying dynamics. This mechanism has the effect of pushing an off-diagonal element to zero if there is not much interaction between the corresponding state variables. There are a total of two hyper-parameters that are used to also control the overall level of sparsity in the matrix, making this a computationally attractive prior.

Similar to the Minnesota-Litterman prior, the horseshoe prior (Carvalho et al. 2009) was developed to allow for greater control of matrix sparsity. With the horseshoe prior, all entries of $\mathbf{B}$ are given a centered normal distribution whose variance is controlled by a global and local sparsity hyper-parameter. The global hyper-parameter sets a certain variance level for all elements of the matrix, while each entry has its own local hyper-parameter. Therefore, using a horseshoe prior for $\mathbf{B}$, say,

requires the learning and sampling of $1 + 2m^2$ total parameters. The horseshoe formulation, while increasing the storage cost and producing a more topologically complicated posterior distribution to sample from, can, in principle, produce a more nuanced sparsity pattern than the Minnesota–Litterman prior. In practice, a regularized, or Finnish horseshoe prior is used.

Once an adequate prior has been chosen, all that remains is to sample the posterior distribution (11) or the approximation (12). There are several methods to achieve this, but in this paper we use a variant of the Markov chain Monte Carlo method (MCMC) (Brooks 2011), Hamiltonian Monte Carlo (HMC) (Neal 2012); see the appendix B for details of the theory and implementation. This method is chosen, in particular, because of the need for an algorithm scalable to large parameter spaces. HMC utilizes gradient information to produce samples from the target distribution quicker than traditional MCMC methods, reducing the storage and computational costs required.

5. Parameter estimation results

We aim to present a test application to demonstrate that the advantages of enhanced estimation and forecast calibration gained from the Bayesian approach is significant. We compare a suite of estimation techniques, including the ML and various MAP estimators corresponding to combinations of the various aforementioned priors. Furthermore, we test the forecast ability using probabilistic metrics to understand the ability of each method to capture the proper forecast distribution.

We consider the task of forecasting yearly SST anomalies. We use SST data from the fully coupled 1800-yr preindustrial control run from the CESM1 Large Ensemble project (Kay et al. 2015), which are averaged to annual values to produce a time series of 1800 data points for each surface ocean location on a roughly 1° global grid. The control run is used here in order to test the impact of data availability on parameter estimation, not necessarily to comment on the associated dynamics from forecasting, and to avoid the effects of uncertainty in atmospheric forcing. This study is intended to showcase the potential benefits from a Bayesian approach by directly comparing it with the classical ML methodology in a fair manner. Further research is necessary to come to a comprehensive conclusion about the benefits of the Bayesian approach on realistic datasets. An empirical orthogonal function (EOF) decomposition is performed to reduce the more than 86 000 spatial points to 10 global EOF spatial patterns that capture approximately 55% of the explained

variance. Let $\{\mathbf{x}_t\}$ represent the time series representing the dynamics of the 10 spatial patterns over the 1800-yr time frame. A roughly 900-yr fraction of this is taken as training data; 500 years are reserved as test data with the remaining portion saved for validation purposes. The matrices $\mathbf{B}$ and $\mathbf{Q}$ require $10 \times 10 = 100$ elements each to be estimated (i.e., the ratio of data points to parameters is 10:1). Increasing the number of EOF spatial patterns reduces this ratio while only obtaining marginal increases in explained variance. We focus our attention on the predictability of the high-latitude regions, and the evidence we display will be concerned with each method's performance in these regions. We believe that using global EOFs may capture possible teleconnections between regions, but recognize that strictly higher-latitude EOFs may capture regional variability using fewer modes.

Here we aim to compare the estimation capabilities of the ML method and the Bayesian approach with a suite of prior distributions. This study performs two experiments: first, a perfect model experiment in order to determine the accuracy of ML and MAP estimates of parameters, and, second, a forecasting experiment in order to test the calibration of the various parameter densities.

a. Perfect model methodology

In this perfect model scenario, we test the ability of each estimation technique to accurately estimate the mean parameter values given data. Using the truncated EOFs, we first estimate MAP values for the parameters $\mathbf{B}$ and $\mathbf{Q}$ using HMC (see the appendix). Using these MAP values, we discretize (1) using the standard Euler-Maruyama numerical method and simulate a path with integration time step $\Delta t = 10^{-3}$ years. To subsample from this path, we take observations at intervals $\tau = k\Delta t = 1$, to construct annual time series of length $N = 50, 100, 1000$. With these time series as data, the mean likelihood or posterior estimates are constructed and compared against the “true” parameter values for $\mathbf{B}$ and $\mathbf{Q}$. The benefit of using data-inspired parameter values, as opposed to specifying values arbitrarily (e.g., diagonal matrices), is to not preimpose prior beliefs on the test data or true parameter values and potentially incur bias in the test results. Because of the randomness in the construction of the sample paths and Markov chain construction, the use of HMC estimated parameters as truth should not introduce bias results toward this method.

For each of the prior distributions that we have discussed, we measure the relative errors (as a percent) of point estimates, $\text{argmax}_p(\theta|\mathbf{y}_{1:T})$, from the “true”

FIG. 2. Relative errors (%) for **B** and **Q** using a variety of prior densities and estimation techniques from the perfect model experiment. Data were generated from an Euler–Maruyama discretization ($\Delta t = 10^{-3}$) of a linear SDE with prescribed drift and diffusion coefficients **B** and **Q**. From this series, N points were uniformly observed for the estimation dataset. Relative errors are taken with respect to the ML or MAP estimates of the parameters. We use the notation PriorB_PriorQ, with ML, N, MINN, LKJ, and HORSE referring to maximum likelihood, normal, Minnesota, LKJ, and horseshoe priors, respectively.

parameter values and results are shown in Fig. 2. There are three main conclusions from this experiment. First, as should be no surprise, the relative errors decrease, although somewhat unevenly, as more data are available. The lack of clear convergence, as a result of the “limited” data, should serve as a precaution to those relying on the central limit theorem for asymptotic guarantees. Second, the Bayesian methodology can produce estimates 4–6 times more accurate than ML alone. The choice of Minnesota prior for **B** and either LKJ or horseshoe for **Q** produced nearly identical results and features slightly less than 100% relative error in the large data scenario. Finally, imposing structure in a concise manner is a critical component when considering a prior distribution. There are obviously some prior combinations that perform just as poorly as MLE, such as the normal priors, but priors with many hyperparameters (like using the horseshoe prior for **B**) can also be underwhelming.

The errors in this experiment may seem unreasonably high to an observer given the relative success of applying LIMs to similar applications. We address this concern in two ways. First, the signal in the yearly averaged CESM control run data is likely much weaker than those used in other papers, such as those from regional or monthly datasets that include atmospheric forcing. The stronger the mean of the signal compared to the variance, the better the parameter estimation results will be. There are also some artificial/ad hoc ways to reduce noise in the data to achieve the same ends, such as smoothing, that we do not perform in this analysis. It should also be noted that while statistical bootstrapping techniques are common ways to improve estimation by inflating the sample size, we did not find discernible advantages of performing such resampling and choose not to report

the details in this paper. Second, as we will see in the next experiment, an accurate ML or MAP estimate is not necessary for producing reasonable forecasts. While this experiment measures the error in the mean of the posterior distribution, it does not comment on the overall calibration of the probability distributions or how their uncertainties are propagated through to state forecasts. Therefore, it is not necessary contradictory in practice for there to be both large relative errors in parameter recovery while forecast errors are relatively low.

b. CESM1 data experiments

When using an LIM, there is traditionally only one source of uncertainty when computing sample paths of the SDE: the noise process with covariance **Q**. In the Bayesian framework there is an additional source of uncertainty: the distribution of the parameters **B** and **Q** themselves. Therefore, we hypothesize that forecasts computed with the MLE alone are too conservative and mis-estimate the variance and resulting distribution of the underlying physical process. The goal of this section is to measure the calibration, that is, how well the method reproduces the true forecast probability distribution, of various sample posterior distributions $p(\mathbf{x}_t | \mathbf{B}, \mathbf{Q}, \mathbf{y}_{t-1})$.

To perform each experiment, we calculate a 100-member ensemble with lead time τ using parameters drawn from their respective posterior distribution (in the ML case this is taken to be only the point estimate). The posterior distributions are estimated from the training set of 900 years. Forecasts are initialized from the test dataset, independent of the training dataset. For this test, the datasets are smoothed with a 3-yr backward-looking moving average. For the Bayesian

TABLE 1. Correlation coefficient between the CESM-LE data and the median path from $p(\mathbf{x}_t|\mathbf{y}_t, \mathbf{B}, \mathbf{Q})$ using a variety of prior densities and estimation techniques. Ensembles of predictions were generated from an Euler–Maruyama discretization ($dt = 10^{-2}$ between data points) of a linear SDE with drift and diffusion coefficients $\mathbf{B}$ and $\mathbf{Q}$ drawn from the associated posterior distribution for each member of the ensemble. From the ensemble, the median path was calculated and compared to the original dataset. ML means that Bayesian estimation was not performed for that variable and only the maximum likelihood value was used.

		Priors for $\mathbf{Q}$			
		ML	N	LKJ	HORSE
Priors for $\mathbf{B}$	ML	0.429	—	—	—
	N	—	0.852	0.852	0.851
	HORSE	—	0.852	0.850	0.851
	MINN	—	0.851	0.852	0.852

estimates, a value for $\mathbf{B}$ and $\mathbf{Q}$ are drawn anew for each ensemble member for each one-step prediction. From this ensemble, we calculate an array of statistics and metrics to judge the induced accuracy and calibration from each choice of prior distribution. Not only will we compute the correlation coefficient, an oft-used statistic used in time series forecasting, but we will also calculate the empirical distribution function for the posterior distribution of sample paths and compare to the distribution of the original data using probabilistic metrics in order to have a holistic view of the ability of each method to capture the underlying uncertainty.

Table 1 displays the correlation coefficient between the median sample path and test data. In practice a correlation coefficient greater than 0.6 represents a heuristic for skillful forecasts (Collins 2002). All fully Bayesian forecasting methods exceed this threshold and hybrid methods with normal and horseshoe priors for $\mathbf{B}$ also produce skillful forecast that rival the fully Bayesian forecasts. Only using a MLE for $\mathbf{B}$ produces poor forecast skill. These two facts together suggest that underspecification of $p(\mathbf{B})$, not $p(\mathbf{Q})$, is one of the major causes of poor forecast performance. The spatial, graphical results for ML and Bayesian estimation using Minnesota and LKJ priors can be seen in Figs. 3 and 4 for 1- and 2-yr lead times, respectively. For forecasts with 1-yr lead time, the ML estimates are generally skillful in the Arctic Ocean with exceptions in the Nares Strait and Beaufort and Kara Seas. In the Southern Ocean, the ML estimates are only skillful off of East Antarctica and in pockets in the Ross and Weddell Seas. The Bayesian forecasts have exceptional skill in both polar regions. For forecasts with 2-yr lead time, the overall skill of either technique is markedly lower. The ML estimates are not skillful in any region. The Bayesian forecasts are considered skillful still in many

areas of both polar regions. In particular we see high correlation skill off of West Antarctica in the Amundsen–Bellingshausen Seas, whose climate is known to be strongly influenced by tropical Pacific SSTs (Steig et al. 2012; Ding et al. 2011; Lachlan-Cope and Connolley 2006). The correlation metric, then, is very sensitive to proper accounting of the uncertainty in $\mathbf{B}$, but this is not a comprehensive view of the calibration of the forecasts.

A view of the correlation coefficient of various methods as a function of lead time is given in Fig. 5. Here, we plot both the global mean (dashed) and global maximum (solid) correlation coefficients of the ML forecasts, Bayesian forecasts with Minnesota and LKJ priors, and vector auto-regressive (VAR) forecasts with two and three terms. These VAR models are fit with a optimization-based maximum likelihood approach, and are included as a means of comparison of LIMs against other typical time series forecasting techniques. In general, the ML LIM forecasts are much less skillful than either the Bayesian LIM or VAR models, with the latter models having roughly identical skill. In particular, the low global correlation coefficients and sharp decline of skill at small lead times from the ML forecasts are extreme in comparison to the other methods. Notice that while the ML forecasts have skillful globally averaged correlation coefficients, the low quality of the average correlation coefficients indicates that there is much greater variance of region dependent skill when compared to the other methods, especially for short lead times. All methods, however, produce forecasts that are not skillful on average with a lead time of greater than 3 years.

To begin to measure the calibration of probabilistic forecasts, we use the ensemble forecast members to construct conditional cumulative distribution functions, from which quantiles can be calculated and compared to the empirical distribution function for the observations:

$$\hat{F}(t) = \frac{1}{T} \sum_{n=1+\tau}^T \mathbf{1}_{(\mathbf{y}_n - \mathbf{y}_{n-\tau}) < t} \quad (15)$$

To be clear, under the assumption of stationarity, this average is performed temporally, but for clarity in presentation we also integrate spatially to consider a one-dimensional distribution function. Figure 6 plots the observed and forecast distribution functions against each other, where the discrepancy between forecasted and observed frequencies is plotted. A perfect method would have 1:1 agreement between the forecasts and the observations and have zero residual forecast probability, as denoted in the plot by the dashed (red) line. The

FIG. 3. Correlation coefficient between mean forecasts with 1-yr lead time and test data; higher correlation coefficients are more predictive. (top) ML estimates were used for the parameter values of $\mathbf{B}$ and $\mathbf{Q}$. (bottom) Minnesota and LKJ priors were used for $p(\mathbf{B})$ and $p(\mathbf{Q})$, respectively.

distributions of the ML forecasts (dashed) are generally below the dotted line, especially for more frequent events, representing that the method underestimates the uncertainty in the data. The forecasts generated with a Minnesota prior for $\mathbf{B}$ and a horseshoe prior for $\mathbf{Q}$ have generally good agreement with the observations, but slightly overestimate the uncertainty in unlikely events. We include vector auto-regressive models for a wider comparison of methods; these VAR models are fitted with an optimization based maximum likelihood model. The VAR models suffer more strongly from over- and underestimation than the Bayesian methods. All of the Bayesian methods outperform the ML techniques in this metric, showing again that regularizing using reasonable prior densities and having access to the entire probability distribution produces more skillful forecasts.

Figure 7 presents the relative L_2 error between the distribution of forecasts at a lead time of 1 year and

of observations (what we call the “calibration error”), as a function of space. Here a lower calibration error represents a smaller percent discrepancy between the forecast probability distribution from the empirical observational distribution. Generally, the ML forecasts are at least 1.5 to 2 times as less calibrated than the Bayesian forecasts.

6. Conclusions

In this paper we advocate for a Bayesian approach to LIMs as a way to 1) address the ill-conditioned nature of high-dimensional parameter estimation by providing a formal way to regularize the parameter distributions and 2) produce well-calibrated predictive uncertainties. This strategy has been tested and compared to traditional LIM modeling at recovering LIM parameter values and at forecasting high-latitude SST anomalies. The evidence presented in this paper suggests that a Bayesian

FIG. 4. Correlation coefficient between mean forecasts with 2-yr lead time and test data; higher correlation coefficients are more predictive. (top) ML estimates were used for the parameter values of $\mathbf{B}$ and $\mathbf{Q}$. (bottom) Minnesota and LKJ priors were used for $p(\mathbf{B})$ and $p(\mathbf{Q})$, respectively.

approach to statistical forecasting may produce significant recovery and calibration benefits, especially in terms of regional predictability, when compared to simple point estimation.

Bayesian approaches to parameter estimation for stochastic differential equations are rare in the literature for a number of factors: there is a lack of closed form solutions to key probability densities, prohibitive computational cost make estimation difficult, and the decent historic performance of various levels of simplifying approximations has not spurred the need to apply the Bayesian approach. The increasing amount of computational resources now commonly available, development of efficient sampling algorithms, and evolving view of statistical estimation have set a fertile field for the adoption of the Bayesian framework. This paper advocates for the intelligent use of a Bayesian framework with nontrivial prior distributions to formulate the LIM problem and provides an accounting of how well these

approaches reconstruct and forecast the appropriate probability distributions.

Maximum likelihood and least squares approaches to parameter estimation in LIMs have asymptotic convergence guarantees, but may not be optimal when the available data are sparse, the underlying signal is not

FIG. 5. Global mean (dashed) and maximum (solid) correlation coefficients of forecasts from using a variety of methods as a function of the lead or lag time τ .

FIG. 6. Comparison of empirical distribution function of observed data against ensembles of forecasts using different prior probabilities as well as ML-estimated VAR(2, 0) and VAR(3, 0) methods. The horizontal axis represents the SST anomalies that occur at an observed frequency based on the data, and the vertical axis represents the difference between the forecasted and observed frequency of those values. The red dashed line represents an ideal 1:1 correspondence between forecast and observation distribution. Positive (negative) residuals represent overestimation (underestimation) of events of that observed frequency.

strong, or there are confounding variables present, or when confronted with model error. As demonstrated by the simple example in Figs. 1 and 2, researchers should be cautious about using such estimators in the LIM framework. We found that, even in a relatively large dataset, relative errors often exceeded 100%. Furthermore, our results also suggest that it would be hazardous to conclude that forecast skill is a reflection of parameter reconstruction (Table 1).

The main hypothesis of this paper is that using ML point estimates can result in overly conservative forecasts, especially when compared to results from a Bayesian framework. Using both traditional and probabilistic forecast metrics, we found that predictions from a Bayesian approach are more accurate than their ML competitors. Comparing the empirical distribution functions from forecasts and the metrics in Table 1, Figs. 6 and 7 affirm that the Bayesian approach better captures the probability distribution of the dynamics and makes clear that the benefits from using a Bayesian framework in forecasting can be significant.

There are several possible avenues to continue and extend the analysis in this paper. Immediately, this analysis should be extended to observational datasets so that comparisons can start to be used with other papers (e.g., Penland and Sardeshmukh 1995). It is still an open question of whether a Bayesian framework can be used to reliably improve the performance of LIMs in forecasting decadal predictability of SSTs (Hawkins and

Sutton 2009; Branstator and Teng 2010; Guemas et al. 2012), and this area can be pursued with either a global or regional scope. Second, further exploration of relevant prior distributions should be pursued with the goal of producing physically meaningful prior densities. In particular, we can explore the use of global climate models in the Bayesian framework. There are several extensions to LIMs where one could feasibly implement a Bayesian approach. First, investigation of LIM frameworks that incorporate state-dependent noise, as in Sardeshmukh et al. (2015) and Martinez-Villalobos et al. (2019), can also be formulated in a Bayesian way given an appropriate approximate likelihood distribution. Finally, one could depart from the parametric form that a linear stochastic differential equation imposes and attempt a Bayesian framework with more complicated, nonparametric physically constrained statistical models such as in Majda and Harlim (2013) and Wikle and Holan (2011) or machine learning techniques such as Bayesian networks in McDermott and Wikle (2019a,b).

Acknowledgments. This research was supported in part by the Regional and Global Model Analysis (RGMA) component of the Earth and Environmental System Modeling (EESM) program of the U.S. Department of Energy's Office of Science, as contribution to the HiLAT-RASM project. DF, NMU, and DC received support for this project from the U.S. Department of Energy, Biological and Environmental Research (BER) division (Regional and Global Modeling and Applications program); NMU received additional support under the BER Early Career Research program.

APPENDIX A

Prior Matrix Distributions

a. LKJ

The so-called LKJ prior (Lewandowski et al. 2009; named after the authors, Lewandowski, Kurowicka, and Joe) is the result of work that extended methods designed to sample random correlation matrices in a computationally efficient manner. In implementation, there is a single hyper-parameter to the distribution η and the distribution can be defined as

$$\text{LKJ}(\Sigma|\eta) \propto \det(\Sigma)^{\eta-1}. \quad (\text{A1})$$

From the definition, if $\eta = 1$, the LKJ distribution is uniform over all correlation matrices. For $\eta > 1$, the distribution places mass closer to 1, resulting in samples close to the unit matrix. In other words, η controls the

FIG. 7. Calibration errors that represent the relative difference between forecast cumulative distribution and the empirical observation distribution. Smaller errors represents a better calibrated forecast. (top) ML estimates were used for parameter values. (bottom) Minnesota and LKJ priors were used for $p(\mathbf{B})$ and $p(\mathbf{Q})$, respectively.

expected amount of correlation, with $\eta > 1$ favoring less correlation, while $\eta < 1$ expecting more correlation.

b. Minnesota–Litterman

The Minnesota–Litterman prior (Litterman 1986; Doan et al. 1983) has two hyper-parameters λ and θ that allow for controlling the degree of sparsity in the matrix and differentiating between diagonal and off-diagonal terms. The prior distributions are given by

$$\begin{aligned} \mathbf{G}_{ii} &\sim \mathcal{N}(1, \lambda), \\ \mathbf{G}_{ij} &\sim \mathcal{N}\left(0, \lambda \theta \frac{\Sigma_{ii}}{\Sigma_{jj}}\right) \quad i \neq j, \\ \lambda &\sim \text{Cauchy}^+(0, 1) \\ \theta &\sim \mathcal{U}(0, 1), \end{aligned} \quad (\text{A2})$$

where $\mathcal{U}(a, b)$ is the uniform distribution on $[a, b]$.

This formulation allows for the distinct treatment of diagonal from off-diagonal parameters. The λ acts as a global covariance hyper-parameter, controlling the overall uncertainty in the elements of the matrix. If $\lambda = 0$, then the prior would collapse to a distribution around the identity matrix. For $\lambda \neq 0$, the off-diagonal terms $\lambda \theta \Sigma_{ii}/\Sigma_{jj}$ shrinks λ by a multiple, θ , of the relative size of the respective variances. This combination is a natural choice since it controls the sparsity of the matrix according to the relative uncertainty and size of the dynamics of elements of $\mathbf{x}_t$.

c. (Finnish) horseshoe

Similar to the Minnesota–Litterman prior, the horseshoe prior (Carvalho et al. 2009) was originally derived to allow for specification of sparsity and has hyper-parameters λ and τ . However, $\lambda \in \mathbb{R}^{m \times m}$ has m^2 elements λ_{ij} —one for each entry of $\mathbf{B}$. The distributions are

$$\begin{aligned}\mathbf{B}_{ij} &\sim \mathcal{N}(0, \tau^2 \lambda_{ij}^2), \\ \lambda_{ij} &\sim \text{Cauchy}^+(0, 1),\end{aligned}\quad (\text{A3})$$

where $\tau > 0$ is some chosen constant. The two hyper-parameters τ and λ are interpreted at global and local sparsity controls. The value of τ should represent our overall belief about the size of elements of $\mathbf{B}$. Then, each λ_{ij} is used to tune the size of individual elements $\mathbf{B}_{ij}$. Each λ_{ij} is given a half-Cauchy prior in order to allow for a large spread of values. If a value of τ is not known a priori, then one can use a noninformative prior like a half-Cauchy instead. Of concern with this scheme is the introduction of $m^2 + 1$ hyper-parameters that must be estimated along with the n^2 elements of $\mathbf{B}$. While the large number of hyper-parameters can help control the size of individual elements $\mathbf{B}_{ij}$, it increases the computational complexity of estimation and may require relatively more data to get reliable estimates than simpler priors. Furthermore, it is occasionally possible during estimation for $\tau^2 \lambda^2$ to become exceedingly large. The Finnish horseshoe introduces a transformation, $\bar{\lambda}_{ij}$, of λ_{ij} and additional hyper-parameter c^2 , such that

$$\bar{\lambda}_{ij}^2 = \frac{\lambda_{ij}^2}{1 + \tau^2 \lambda_{ij}^2 / c^2}. \quad (\text{A4})$$

If $\tau^2 \lambda^2 \gg c^2$, then $\bar{\lambda} \approx \lambda$ and the regularized horseshoe behaves similarly to the original horseshoe prior. On the other-hand, if $\tau^2 \lambda^2 \ll c^2$ then $\bar{\lambda} \ll \lambda$ and $\bar{\lambda}$ is a regularization. This transformation penalizes values of λ and τ that are near some predetermined threshold c^2 . Like τ , c can be estimated by considering the understanding of the size of the parameters, or can be given some non-informative prior and estimated along with the other parameters. Unlike the Minnesota prior, this prior can be applied to both $\mathbf{B}$ and $\mathbf{Q}$. Of course, when estimating $\mathbf{Q}$, we follow the transformation (14) and need only estimate L .

APPENDIX B

Bayesian Computational Techniques

To make full use of the Bayesian framework, instead of relying on maximal parameter estimates, we must be able to draw samples of the parameters from their respective probability distributions. If the marginal likelihood function $p(\mathbf{y}_{1:T}|\theta)$ is normal and the prior distribution $p(\theta)$ is chosen conjugate to the likelihood, then the posterior distribution can be calculated analytically. There are numerous books written on this

technique (see [Gelman et al. 2013](#)). In general, however, numerical techniques are needed to sample from an arbitrary posterior distribution. In particular, Markov chain Monte Carlo (MCMC) algorithms are standard practice to generate independent samples from the posterior. We will review MCMC in order to motivate Hamiltonian Monte Carlo (HMC), which is the algorithm we use for the studies in this paper. This [appendix](#) is only meant as a brief description of the types of methods employed in this paper; for more detailed and comprehensive discussions see [Brooks \(2011\)](#).

a. Markov chain Monte Carlo

A Markov chain is a sequence of random variables $\{\epsilon_1, \epsilon_2, \dots\}$ where

$$p(\epsilon_1, \epsilon_2, \dots, \epsilon_N) = p(\epsilon_1) \prod_{k=2}^N p(\epsilon_k | \epsilon_{k-1}), \quad (\text{B1})$$

that is, a member of the sequence, given the previous value, is conditionally independent of all other elements in the sequence. The goal of MCMC is to generate a Markov chain such that elements of the chain come from the posterior distribution. The elements of the Markov chain are then the requested samples. The basic MCMC algorithm consists of specifying an initial distribution $p(\epsilon_1)$, proposing a next element using a transition distribution $p(\epsilon_k | \epsilon_{k-1})$, and a method for accepting or rejecting the proposed value. In the LIM case, the initial distribution could be specified by the ML distribution since one has analytical representations for the mean and standard deviation. For this method to generate a proper Markov chain that has $p(\theta | \mathbf{y}_{1:T})$ as its equilibrium distribution the transition density and acceptance criterion must meet certain theoretical conditions. If these conditions are met, then $p(\epsilon_k)$ is stationary, $p(\epsilon_k) = p(\epsilon_{k'})$ for all k, k' , and ϵ_k is independent of $\epsilon_{k'}$. This stationary distribution is then guaranteed to be equal to $p(\theta | \mathbf{y}_{1:T})$. In practice, ensuring that generated Markov chains are independent and come from the required distribution is nontrivial. We will not discuss these theoretical concerns, but assume such conditions have been met and refer the reader to the literature. To avoid some of these technical difficulties, we rely on using the Hamiltonian Monte Carlo algorithm.

b. Hamiltonian Monte Carlo

Hamiltonian Monte Carlo (HMC) uses the dynamics from an associated Hamiltonian dynamical system to construct a transition density $p(\epsilon_k | \epsilon_{k-1})$ that is theoretically guaranteed to produce the required posterior distribution. These dynamics are defined by the pair of differential equations

$$\begin{aligned}\frac{d\mathbf{p}}{dt} &= -\frac{\partial \mathbf{H}}{\partial \mathbf{q}}, \\ \frac{d\mathbf{q}}{dt} &= \frac{\partial \mathbf{H}}{\partial \mathbf{p}},\end{aligned}\quad (\text{B2})$$

where $(\mathbf{p}, \mathbf{q})$ are coordinates of the state system (representing momentum and position), and $\mathbf{H}$ is the Hamiltonian defined as

$$\mathbf{H}(\mathbf{p}, \mathbf{q}) = \mathbf{K}(\mathbf{p}, \mathbf{q}) + \mathbf{U}(\mathbf{p}, \mathbf{q}), \quad (\text{B3})$$

where $\mathbf{K}$ is called the “kinetic energy” and $\mathbf{U}$ is the “potential energy.” HMC utilizes this framework by introducing a conjugate momentum variable $\rho \sim \mathcal{N}(0, M)$ and seeking to draw from the joint density $p(\rho, \theta) = p(\rho|\theta)p(\theta|\mathbf{y}_{1:T})$. This joint density defines the Hamiltonian:

$$\begin{aligned}\mathbf{H}(\rho, \theta) &= -\log p(\rho, \theta) \\ &= -\log p(\rho|\theta) - \log p(\theta|\mathbf{y}_{1:T}) \\ &= \mathbf{K}(\rho, \theta) + \mathbf{U}(\theta)\end{aligned}\quad (\text{B4})$$

We assume that the momentum is independent of the target density, so $p(\rho|\theta) = p(\rho)$ and $\mathbf{K}(\rho, \theta) = \mathbf{K}(\rho) = 1/2\rho^T M^{-1}\rho$. To generate transitions for the Markov chain we evolve the system (B2) starting from an initial value for θ and a draw ρ from $p(\rho)$. In particular, we solve

$$\begin{aligned}\frac{d\theta}{dt} &= \frac{\partial \mathbf{K}}{\partial \rho}, \\ \frac{d\rho}{dt} &= -\frac{\partial \mathbf{U}}{\partial \theta}.\end{aligned}\quad (\text{B5})$$

Solving (B5) in general requires discretization of ρ and θ into their respective Markov chains and numerical integration. While there are numerous numerical integrators for Hamiltonian systems, the most common in practice is the leapfrog, or Störmer–Verlet integrator with step size Δt :

$$\begin{aligned}\rho_{n+1/2} &= \rho_n - \frac{\Delta t}{2} \frac{\partial \mathbf{U}}{\partial \theta}, \\ \theta_{n+1} &= \theta_n + \Delta t M^{-1} \rho_{n+1/2}, \\ \rho_{n+1} &= \rho_{n+1/2} - \frac{\Delta t}{2} \frac{\partial \mathbf{U}}{\partial \theta}.\end{aligned}\quad (\text{B6})$$

The probability that these values are accepted is given by

$$\min\{1, \exp[\mathbf{H}(\rho_n, \theta_n) - \mathbf{H}(\rho_{n+1}, \theta_{n+1})]\}.$$

If the generated values are not accepted, then $\theta_{n+1} = \theta_n$ and the algorithm continues for a total number of L time steps.

It is proven that this algorithm produces samples that converge to the target distribution given certain conditions. Often this convergence is quicker, in terms of samples needed, than standard MCMC algorithms. As a tradeoff for speed, the HMC requires calculation of derivatives at each time step. Furthermore, the time step Δt may need to be taken excessively small in order to ensure stability of the scheme, meaning that more time steps are required to adequately explore the state space. Modifications to the original HMC algorithm, the no U-turn sampler in particular, help automatically adjust the step size and number of steps. We leave the details of these algorithms to (Neal 2012). For out-of-the-box implementation of this algorithm, see the programming language Stan (Carpenter et al. 2017).

REFERENCES

- Alexander, M. A., L. Matrosova, C. Penland, J. D. Scott, and P. Chang, 2008: Forecasting Pacific SSTs: Linear inverse model predictions of the PDO. *J. Climate*, **21**, 385–402, <https://doi.org/10.1175/2007JCLI1849.1>.
- Barnston, A. G., Y. He, M. H. Glantz, A. G. Barnston, M. H. Glantz, and Y. He, 1999: Predictive skill of statistical and dynamical climate models in SST forecasts during the 1997–98 El Niño episode and the 1998 La Niña onset. *Bull. Amer. Meteor. Soc.*, **80**, 217–243, [https://doi.org/10.1175/1520-0477\(1999\)080<0217:PSOSAD>2.0.CO;2](https://doi.org/10.1175/1520-0477(1999)080<0217:PSOSAD>2.0.CO;2).
- , M. K. Tippett, M. L. L’Heureux, S. Li, and D. G. Dewitt, 2012: Skill of real-time seasonal ENSO model predictions during 2002–11: Is our capability increasing? *Bull. Amer. Meteor. Soc.*, **93**, 631–651, <https://doi.org/10.1175/BAMS-D-11-00111.1>.
- Batz, P., A. Ruttner, and M. Opper, 2018: Approximate Bayes learning of stochastic differential equations. *Phys. Rev.*, **98**, 022109, <https://doi.org/10.1103/PHYSREVE.98.022109>.
- Beskos, A., O. Papaspiliopoulos, G. O. Roberts, and P. Fearnhead, 2006: Exact and computationally efficient likelihood-based estimation for discretely observed diffusion processes. *J. Roy. Stat. Soc.*, **68**, 333–382, <https://doi.org/10.1111/j.1467-9868.2006.00552.x>.
- Branstator, G., and H. Teng, 2010: Two limits of initial-value decadal predictability in a CGCM. *J. Climate*, **23**, 6292–6311, <https://doi.org/10.1175/2010JCLI3678.1>.
- Brooks, S., 2011: *Handbook of Markov Chain Monte Carlo*. CRC Press/Taylor & Francis, 592 pp.
- Carpenter, B., and Coauthors, 2017: Stan: A probabilistic programming language. *J. Stat. Software*, **76** (1), 1–32, <https://doi.org/10.18637/jss.v076.i01>.
- Carvalho, C. M., N. G. Polson, and J. G. Scott, 2009: Handling sparsity via the horseshoe. *Proc. 12th Int. Conf. on Artificial Intelligence and Statistics*, Clearwater Beach, Florida, AISTATS, 73–80, <http://proceedings.mlr.press/v5/carvalho09a/carvalho09a.pdf>.
- Chen, C., M. A. Cane, N. Henderson, D. E. Lee, D. Chapman, D. Kondrashov, and M. D. Chekroun, 2016: Diversity, non-linearity, seasonality, and memory effect in ENSO simulation and prediction using empirical model reduction. *J. Climate*, **29**, 1809–1830, <https://doi.org/10.1175/JCLI-D-15-0372.1>.
- Collins, M., 2002: Climate predictability on interannual to decadal time scales: The initial value problem. *Climate Dyn.*, **19**, 671–692, <https://doi.org/10.1007/s00382-002-0254-8>.

- Colman, A. W., and M. K. Davey, 2003: Statistical prediction of global sea-surface temperature anomalies. *Int. J. Climatol.*, **23**, 1677–1697, <https://doi.org/10.1002/joc.956>.
- Davis, R. E., 1976: Predictability of sea surface temperature and sea level pressure anomalies over the North Pacific Ocean. *J. Phys. Oceanogr.*, **6**, 249–266, [https://doi.org/10.1175/1520-0485\(1976\)006<0249:POSSTA>2.0.CO;2](https://doi.org/10.1175/1520-0485(1976)006<0249:POSSTA>2.0.CO;2).
- DelSole, T., and A. Y. Hou, 1999: Empirical stochastic models for the dominant climate statistics of a general circulation model. *J. Atmos. Sci.*, **56**, 3436–3456, [https://doi.org/10.1175/1520-0469\(1999\)056<3436:ESMFTD>2.0.CO;2](https://doi.org/10.1175/1520-0469(1999)056<3436:ESMFTD>2.0.CO;2).
- , and P. Chang, 2003: Predictable component analysis, canonical correlation analysis, and autoregressive models. *J. Atmos. Sci.*, **60**, 409–416, [https://doi.org/10.1175/1520-0469\(2003\)060<0409:PCACCA>2.0.CO;2](https://doi.org/10.1175/1520-0469(2003)060<0409:PCACCA>2.0.CO;2).
- , and X. Yang, 2010: State and parameter estimation in stochastic dynamical models. *Physica D*, **239**, 1781–1788, <https://doi.org/10.1016/j.physd.2010.06.001>.
- , L. Jia, and M. K. Tippett, 2013: Decadal prediction of observed and simulated sea surface temperatures. *Geophys. Res. Lett.*, **40**, 2773–2778, <https://doi.org/10.1002/grl.50185>.
- Dias, D. F., A. Subramanian, L. Zanna, and A. J. Miller, 2019: Remote and local influences in forecasting Pacific SST: A linear inverse model and a multimodel ensemble study. *Climate Dyn.*, **52**, 3183–3201, <https://doi.org/10.1007/S00382-018-4323-Z>.
- DiNezio, P. N., C. Deser, Y. Okumura, and A. Karspeck, 2017: Predictability of 2-year La Niña events in a coupled general circulation model. *Climate Dyn.*, **49**, 4237–4261, <https://doi.org/10.1007/s00382-017-3575-3>.
- Ding, Q., E. J. Steig, D. S. Battisti, and M. Küttel, 2011: Winter warming in west Antarctica caused by central tropical Pacific warming. *Nat. Geosci.*, **4**, 398–403, <https://doi.org/10.1038/ngeo1129>.
- Doan, T., R. Litterman, and C. Sims, 1983: Forecasting and conditional projection using realistic prior distributions. NBER working paper series 1202, 100 pp., <https://www.nber.org/papers/w1202.pdf>.
- Eraker, B., 2001: MCMC analysis of diffusion models with application to finance. *J. Bus. Econ. Stat.*, **19**, 177–191, <https://doi.org/10.1198/073500101316970403>.
- Gelman, A., J. B. Carlin, H. S. Stern, D. B. Dunson, A. Vehtari, and D. B. Rubin, 2013: *Bayesian Data Analysis*. 3rd ed. Chapman & Hall/CRC Texts in Statistical Science, 646 pp.
- Giannone, D., M. Lenza, and G. E. Primiceri, 2015: Prior selection for vector autoregressions. *Rev. Econ. Stat.*, **97**, 436–451, https://doi.org/10.1162/REST_a_00483.
- Guemas, V., F. J. Doblas-Reyes, F. Lienert, Y. Soufflet, and H. Du, 2012: Identifying the causes of the poor decadal climate prediction skill over the North Pacific. *J. Geophys. Res.*, **117**, D20111, <https://doi.org/10.1029/2012JD018004>.
- , and Coauthors, 2016: A review on Arctic sea-ice predictability and prediction on seasonal to decadal time-scales. *Quart. J. Roy. Meteor. Soc.*, **142**, 546–561, <https://doi.org/10.1002/qj.2401>.
- Hansen, J. A., and C. Penland, 2007: On stochastic parameter estimation using data assimilation. *Physica D*, **230**, 88–98, <https://doi.org/10.1016/j.physd.2006.11.006>.
- Hasselmann, K., 1988: PIPs and POPs: The reduction of complex dynamical systems using principal interaction and oscillation patterns. *J. Geophys. Res.*, **93**, 11015, <https://doi.org/10.1029/JD093iD09p11015>.
- Hastie, T., R. Tibshirani, and J. Friedman, 2017: *The Elements of Statistical Learning Data Mining, Inference, and Prediction*. 2nd ed. Springer-Verlag, 745 pp., <https://doi.org/10.1007/978-0-387-84858-7>.
- Hawkins, E., and R. Sutton, 2009: Decadal predictability of the Atlantic Ocean in a coupled GCM: Forecast skill and optimal perturbations using linear inverse modeling. *J. Climate*, **22**, 3960–3978, <https://doi.org/10.1175/2009JCLI2720.1>.
- , J. Robson, R. Sutton, D. Smith, and N. Keenlyside, 2011: Evaluating the potential for statistical decadal predictions of sea surface temperatures with a perfect model approach. *Climate Dyn.*, **37**, 2495–2509, <https://doi.org/10.1007/s00382-011-1023-3>.
- Jenkins, A., P. Dutrieux, S. Jacobs, E. J. Steig, G. H. Gudmundsson, J. Smith, and K. J. Heywood, 2016: Decadal ocean forcing and Antarctic ice sheet response: Lessons from the Amundsen Sea. *Oceanography*, **29**, 106–117, <https://doi.org/10.5670/oceanog.2016.103>.
- Karimi, H., and K. B. McAuley, 2016: Bayesian estimation in stochastic differential equation models via laplace approximation. *IFAC-PapersOnLine*, **49**, 1109–1114, <https://doi.org/10.1016/j.ifacol.2016.07.351>.
- Kay, J. E., and Coauthors, 2015: The Community Earth System Model (CESM) large ensemble project: A community resource for studying climate change in the presence of internal climate variability. *Bull. Amer. Meteor. Soc.*, **96**, 1333–1349, <https://doi.org/10.1175/BAMS-D-13-00255.1>.
- Kondrashov, D., M. D. Chekroun, and M. Ghil, 2015: Data-driven non-Markovian closure models. *Physica D*, **297**, 33–55, <https://doi.org/10.1016/j.physd.2014.12.005>.
- Kravtsov, S., D. Kondrashov, and M. Ghil, 2009: Empirical model reduction and the modelling hierarchy in climate dynamics and the geosciences. *Stochastic Physics and Climate Modelling*, T. Palmer and P. Williams, Eds., Cambridge University Press, 35–72.
- Lachlan-Cope, T., and W. Connolley, 2006: Teleconnections between the tropical Pacific and the Amundsen-Bellinghousen Sea: Role of the El Niño/Southern Oscillation. *J. Geophys. Res.*, **111**, D23101, <https://doi.org/10.1029/2005JD006386>.
- Latif, M., T. Martin, A. Reintges, and W. Park, 2017: Southern Ocean decadal variability and predictability. *Curr. Climate Change Rep.*, **3**, 163–173, <https://doi.org/10.1007/S40641-017-0068-8>.
- Le Breton, A., 1977: Parameter estimation in a linear stochastic differential equation. *Transactions of the Seventh Prague Conference on Information Theory, Statistical Decision Functions, Random Processes and of the 1974 European Meeting of Statisticians*, J. Kozesnik, Ed., Springer Netherlands, 353–366, https://doi.org/10.1007/978-94-010-9910-3_36.
- Lewandowski, D., D. Kurowicka, and H. Joe, 2009: Generating random correlation matrices based on vines and extended onion method. *J. Multivariate Anal.*, **100**, 1989–2001, <https://doi.org/10.1016/j.jmva.2009.04.008>.
- Litterman, R. B., 1986: Forecasting with Bayesian vector autoregressions—Five years of experience. *J. Bus. Econ. Stat.*, **4**, 25–38, <https://doi.org/10.1080/07350015.1986.10509491>.
- Majda, A. J., and J. Harlim, 2013: Physics constrained nonlinear regression models for time series. *Nonlinearity*, **26**, 201–217, <https://doi.org/10.1088/0951-7715/26/1/201>.
- Martinez-Villalobos, C., D. J. Vimont, C. Penland, M. Newman, and J. D. Neelin, 2018: Calculating state-dependent noise in a linear inverse model framework. *J. Atmos. Sci.*, **75**, 479–496, <https://doi.org/10.1175/JAS-D-17-0235.1>.
- , M. Newman, D. J. Vimont, C. Penland, and J. D. Neelin, 2019: Observed El Niño–La Niña asymmetry in a linear model. *Geophys. Res. Lett.*, **46**, 9909–9919, <https://doi.org/10.1029/2019GL082922>.

- Mason, S. J., and G. M. Mimmack, 2002: Comparison of some statistical methods of probabilistic forecasting of ENSO. *J. Climate*, **15**, 8–29, [https://doi.org/10.1175/1520-0442\(2002\)015<0008:COSSMO>2.0.CO;2](https://doi.org/10.1175/1520-0442(2002)015<0008:COSSMO>2.0.CO;2).
- McDermott, P. L., and C. K. Wikle, 2019a: Bayesian recurrent neural network models for forecasting and quantifying uncertainty in spatial-temporal data. *Entropy*, **21**, 184, <https://doi.org/10.3390/e21020184>.
- , and —, 2019b: Deep echo state networks with uncertainty quantification for spatio-temporal forecasting. *Environmetrics*, **30**, e2553, <https://doi.org/10.1002/env.2553>.
- McKinnon, K. A., and C. Deser, 2018: Internal variability and regional climate trends in an observational large ensemble. *J. Climate*, **31**, 6783–6802, <https://doi.org/10.1175/JCLI-D-17-0901.1>.
- Meehl, G. A., and Coauthors, 2014: Decadal climate prediction an update from the trenches. *Bull. Amer. Meteor. Soc.*, **95**, 243–267, <https://doi.org/10.1175/BAMS-D-12-00241.1>.
- Møller, J. K., H. Madsen, and J. Carstensen, 2011: Parameter estimation in a simple stochastic differential equation for phytoplankton modelling. *Ecol. Modell.*, **222**, 1793–1799, <https://doi.org/10.1016/j.ecolmodel.2011.03.025>.
- Neal, R. M., 2012: MCMC using Hamiltonian dynamics, <https://doi.org/10.1201/B10905-6>.
- Newman, M., 2007: Interannual to decadal predictability of tropical and North Pacific sea surface temperatures. *J. Climate*, **20**, 2333–2356, <https://doi.org/10.1175/JCLI4165.1>.
- , 2013: An empirical benchmark for decadal forecasts of global surface temperature anomalies. *J. Climate*, **26**, 5260–5269, <https://doi.org/10.1175/JCLI-D-12-00590.1>.
- Nummelin, A., S. Jeffress, and T. Haine, 2018: Statistical inversion of surface ocean kinematics from sea surface temperature observations. *J. Atmos. Oceanic Technol.*, **35**, 1913–1933, <https://doi.org/10.1175/JTECH-D-18-0057.1>.
- Partridge, M. D., and D. S. Rickman, 1998: Generalizing the Bayesian vector autoregression approach for regional inter-industry employment forecasting. *J. Bus. Econ. Stat.*, **16**, 62, <https://doi.org/10.1080/07350015.1998.10524735>.
- Penland, C., 1989: Random forcing and forecasting using principal oscillation pattern analysis. *Mon. Wea. Rev.*, **117**, 2165–2185, [https://doi.org/10.1175/1520-0493\(1989\)117<2165:RFAFUP>2.0.CO;2](https://doi.org/10.1175/1520-0493(1989)117<2165:RFAFUP>2.0.CO;2).
- , 1996: A stochastic model of IndoPacific sea surface temperature anomalies. *Physica D*, **98**, 534–558, [https://doi.org/10.1016/0167-2789\(96\)00124-8](https://doi.org/10.1016/0167-2789(96)00124-8).
- , and P. D. Sardeshmukh, 1995: The optimal growth of tropical sea surface temperature anomalies. *J. Climate*, **8**, 1999–2024, [https://doi.org/10.1175/1520-0442\(1995\)008<1999:TOGOTS>2.0.CO;2](https://doi.org/10.1175/1520-0442(1995)008<1999:TOGOTS>2.0.CO;2).
- , and L. Matrosova, 2001: Expected and actual errors of linear inverse model forecasts. *Mon. Wea. Rev.*, **129**, 1740–1745, [https://doi.org/10.1175/1520-0493\(2001\)129<1740:EAAEOL>2.0.CO;2](https://doi.org/10.1175/1520-0493(2001)129<1740:EAAEOL>2.0.CO;2).
- , T. Magorian, C. Penland, and T. Magorian, 1993: Prediction of Niño 3 sea surface temperatures using linear inverse modeling. *J. Climate*, **6**, 1067–1076, [https://doi.org/10.1175/1520-0442\(1993\)006<1067:PONSST>2.0.CO;2](https://doi.org/10.1175/1520-0442(1993)006<1067:PONSST>2.0.CO;2).
- Révilard, A., C. Frankignoul, and Y. O. Kwon, 2018: A multivariate estimate of the cold season atmospheric response to North Pacific SST variability. *J. Climate*, **31**, 2771–2796, <https://doi.org/10.1175/JCLI-D-17-0061.1>.
- Samaniago, F. J., 2010: *A Comparison of the Bayesian and Frequentist Approaches to Estimation*. 1st ed. Springer-Verlag, 225 pp., <https://doi.org/10.1007/978-1-4419-5941-6>.
- Sardeshmukh, P. D., G. P. Compo, and C. Penland, 2015: Need for caution in interpreting extreme weather statistics. *J. Climate*, **28**, 9166–9187, <https://doi.org/10.1175/JCLI-D-15-0020.1>.
- Särkkä, S., 2010: *Bayesian Filtering and Smoothing*. Cambridge University Press, 232 pp., <https://doi.org/10.1017/cbo9781139344203>.
- Singer, H., 2004: A survey of estimation methods for stochastic differential equations. *Sixth Int. Conf. on Social Science Methodology*, Amsterdam, Netherlands, RC33, 12 pp.
- Steig, E. J., Q. Ding, D. S. Battisti, and A. Jenkins, 2012: Tropical forcing of circumpolar deep water inflow and outlet glacier thinning in the Amundsen Sea Embayment, West Antarctica. *Ann. Glaciol.*, **53**, 19–28, <https://doi.org/10.3189/2012AoG60A110>.
- Stock, C. A., and Coauthors, 2015: Seasonal sea surface temperature anomaly prediction for coastal ecosystems. *Prog. Oceanogr.*, **137**, 219–236, <https://doi.org/10.1016/j.pocean.2015.06.007>.
- Stoer, J., and R. Bulirsch, 2002: Eigenvalue problems. *Introduction to Numerical Analysis*, Springer, 364–464, https://doi.org/10.1007/978-0-387-21738-3_6.
- Thomas, E. E., D. J. Vimont, M. Newman, C. Penland, and C. Martínez-Villalobos, 2018: The role of stochastic forcing in generating ENSO diversity. *J. Climate*, **31**, 9125–9150, <https://doi.org/10.1175/JCLI-D-17-0582.1>.
- Tian, T., Y. Zhou, Y. Wu, and X. Ge, 2014: Estimation of parameters in mean-reverting stochastic systems. *Math. Probl. Eng.*, **2014**, 317059, <https://doi.org/10.1155/2014/317059>.
- van Zanten, H., 2013: Nonparametric Bayesian methods for one-dimensional diffusion models. *Math. Biosci.*, **243**, 215–222, <https://doi.org/10.1016/j.mbs.2013.03.008>.
- Vogel, C. R., 2002: *Computational Methods for Inverse Problems*. Society for Industrial and Applied Mathematics, 199 pp., <https://doi.org/10.1137/1.9780898717570>.
- von Storch, H., and T. Bruns, 1988: Principal oscillation pattern analysis of the 30- to 60-day oscillation in general circulation model equatorial troposphere. *J. Geophys. Res.*, **93**, 11 022–11 036, <https://doi.org/10.1029/JD093ID09P11022>.
- Wikle, C. K., and S. H. Holan, 2011: Polynomial nonlinear spatio-temporal integro-difference equation models. *J. Time Ser. Anal.*, **32**, 339–350, <https://doi.org/10.1111/j.1467-9892.2011.00729.x>.
- Wittenberg, A. T., A. Rosati, T. L. Delworth, G. A. Vecchi, and F. Zeng, 2014: ENSO modulation: Is it decadally predictable? *J. Climate*, **27**, 2667–2681, <https://doi.org/10.1175/JCLI-D-13-00577.1>.
- Wu, S., M. Notaro, S. Vavrus, E. Mortensen, R. Montgomery, J. de Piérola, and P. Block, 2018: Efficacy of tendency and linear inverse models to predict southern Peru's rainy season precipitation. *Int. J. Climatol.*, **38**, 2590–2604, <https://doi.org/10.1002/joc.5442>.
- Zhang, L., T. L. Delworth, and L. Jia, 2017: Diagnosis of decadal predictability of Southern Ocean sea surface temperature in the GFDL CM2.1 model. *J. Climate*, **30**, 6309–6328, <https://doi.org/10.1175/JCLI-D-16-0537.1>.
- Zhao, Y., Z. Liu, F. Zheng, and Y. Jin, 2019: Parameter optimization for real-world ENSO forecast in an intermediate coupled model. *Mon. Wea. Rev.*, **147**, 1429–1445, <https://doi.org/10.1175/MWR-D-18-0199.1>.