

Nonlinear Laplacian spectral analysis for time series: Capturing intermittency and low-frequency variability

Dimitrios Giannakis*, and Andrew J. Majda*

*Center for Atmosphere Ocean Science, Courant Institute of Mathematical Sciences, New York University, New York, NY 10012

Submitted to Proceedings of the National Academy of Sciences of the United States of America

Many processes in science and engineering develop multiscale temporal and spatial patterns, with complex underlying dynamics and time-dependent external forcings. Because of the importance in understanding and predicting these phenomena, extracting the salient modes of variability empirically from incomplete observations is a problem of wide contemporary interest. Here, we present a technique for analyzing high-dimensional, complex time series that exploits the geometrical relationships between the observed data points to recover features characteristic of strongly nonlinear dynamics (such as intermittency and rare events), which are not accessible to classical singular spectrum analysis (SSA). The method employs Laplacian eigenmaps, evaluated after suitable time-lagged embedding, to produce a reduced representation of the observed samples, where standard tools of matrix algebra can be used to perform truncated singular value decomposition despite the nonlinear geometrical structure of the data set. We illustrate the utility of the technique in capturing intermittent modes associated with the Kuroshio current in the North Pacific sector of a general circulation model, and dimensional reduction of a low-order atmospheric model featuring chaotic intermittent regime transitions, where classical SSA is already known to fail dramatically.

Laplacian eigenmaps | singular spectrum analysis | decadal variability | dimensional reduction | regime behavior | manifold embedding

The analysis of complex spatio-temporal signals is a central problem in science and engineering, arising in a diverse range of disciplines, including experimental and theoretical fluid dynamics [1, 2], climate science [3, 4, 5], molecular dynamics [6, 7], and astrophysics [8, 9]. In these and other applications, there is a strong interest in extracting physically-meaningful information about the spatial and temporal variability of data from models, experiments, or observations, with the goal of enhancing the understanding of the underlying dynamics, and improving predictive capabilities. In many cases, observations of the system under study are incomplete; i.e., only a subset of the full phase space is accessible.

A classical way of attacking this problem is through singular spectrum analysis (SSA), or one of its variants [3, 10, 11, 12]. Here, a low-rank approximation of a dynamic process is constructed by first embedding a time series of a scalar or multivariate observable in a high-dimensional vector space H (called embedding space) using the method of delays [13, 14, 15], and then performing a truncated singular-value decomposition (SVD) of the matrix X containing the embedded data [16]. In this manner, information about the dynamical process is extracted from the spatial and temporal singular vectors of X with the K largest singular values [12]. The temporal singular vectors form a set of empirical orthogonal functions (EOFs) in H , which, at each instance of time, are weighted by the corresponding principal components (PCs) determined from the temporal singular vectors to yield a rank- K reconstruction of X .

A potential drawback of this approach is that it is based on minimizing an operator norm which may be unsuitable for signals generated by nonlinear dynamical systems. Specifi-

cally, the PCs are computed by projecting onto the principal axes of the k -dimensional ellipsoid that best fits the covariance of the data in H in the least-squares sense. This construction is optimal when the underlying dynamics is linear, but nonlinear processes in general will produce data lying on a curved manifold M in embedding space, with non-Gaussian distributions departing significantly from the ellipsoid defined by the covariance operator in H . Physically, a prominent manifestation of this phenomenon is failure to capture via SSA the intermittent patterns arising in turbulent dynamical systems; i.e., temporal processes that carry low variance, but play an important dynamical role [17, 18].

Despite their inherently nonlinear character, such data sets possess a natural linear structure, namely the Hilbert space $L^2(M, \mu)$ of square-integrable functions on M with inner product inherited from the volume element μ of M (the Riemannian measure). This space may be thought of as the collection of all possible weights that can be assigned to the data samples when making a reconstruction, i.e., it is analogous to the temporal space in SSA. Therefore, it is reasonable to develop algorithms that seek to approximate suitably defined maps from $L^2(M, \mu)$ to the space of spatial patterns H . Such maps, denoted here by A , have the advantage of being simultaneously linear and compatible with the nonlinear manifold comprised by the data.

In this paper, we advocate that this approach, implemented via algorithms developed in machine learning, can reveal important aspects of complex, high-dimensional signals, which are not accessible to classical SSA. In this framework, which we refer to as nonlinear Laplacian spectral analysis (NLSA), an orthonormal basis for $L^2(M, \mu)$ is constructed through eigenfunctions of the Laplace-Beltrami operator on M , computed efficiently via sparse graph-theoretic algorithms [19, 20]. Projecting the data from embedding space H onto these eigenfunctions then gives a matrix representation of A , such that the optimal rank- K reconstruction with respect to the Frobenius norm of maps from $L^2(M, \mu)$ to H is given by standard truncated SVD.

We discuss applications of this technique to two problems, each with its own challenges to data analysis. First, we study the variability of the upper-ocean temperature in the North Pacific sector of a comprehensive coupled climate

Reserved for Publication Footnotes

model (NCAR's CCSM3 model) [21, 22, 23]. Here, the dynamics take place in a phase space of very high dimension, is strongly mixing, and exhibits variability over a broad range of time scales, including seasonal, interannual, and decadal time scales. Imposing no a priori assumptions (such as periodicity in the statistics), NLSA recovers three distinct types of temporal processes: (1) periodic processes, including annual and semiannual cycles; (2) decadal-scale variability with spatial patterns resembling the Pacific Decadal Oscillation (PDO) [24]; (3) intermittent processes associated with the Kuroshio current and variations in the strength of the subtropical and subpolar gyres. The latter carry little variance (and are therefore not captured by SSA), yet their dynamical role is expected to be significant. Moreover, we demonstrate the dynamical significance of Galerkin bases constructed through the spatial modes of NLSA in the context of a six-dimensional chaotic low-order model of the atmosphere [25, 26, 18]. This model exhibits strong intermittency and metastability, and its attractor is highly inhomogeneous. Here, we test for the quality of spatial patterns not by their explained variance, but rather by their ability to reproduce chaotic regime behavior in reduced dynamical models. It has been established, for instance, that Galerkin projections onto the spatial modes of SSA or optimal persistence patterns (OPPs) [4] fail to reproduce chaotic regime behavior [18]. In contrast, NLSA leads to models of dimension as low as three featuring chaotic regimes.

Mathematical framework

Setting. We consider that we observe samples of a vector-valued signal x_t sampled uniformly with time step δt . Here, x_t is generated by a dynamical system operating in some phase space, which may only be partially accessible to observations. For instance, in the ocean application ahead, x_t will be a depth-averaged sea temperature field restricted in the North-Pacific sector of CCSM3, i.e., observations of x_t alone are not sufficient to uniquely determine the state of the system in phase space. Our objective is to produce a decomposition of x_t into spatial and temporal modes, taking explicitly into account the fact that the underlying trajectory of the dynamical system lies on a nonlinear manifold M in phase space. A further challenge, which will be highlighted in our atmospheric application, is that sampling of M may be highly inhomogeneous. In this case, regions of M associated with metastable regimes are sampled densely, but the dynamically important transitions between the regimes are observed rarely.

The methodology developed here employs three basic steps to address these issues: (1) Time-lagged embedding of the input data to a high-dimensional Hilbert space H , to ensure that geometrical neighborhoods in H correspond to neighborhoods in phase space; (2) Construction of a regularized linear map A taking a Hilbert space of scalar functions on M representing temporal patterns to the spatial patterns in H , in a manner compatible with the nonlinear geometry of M ; (3) Singular value decomposition of A in a basis of Laplacian eigenfunctions, evaluated via graph-theoretic algorithms with modified Gaussian kernels to deal with sampling inhomogeneity. Below we describe each step of the method in some detail. Pseudocode summarizing the procedure is listed in Algorithm S1.

Time-lagged embedding. A consequence of incomplete observations is that distances between points in data space may be small, even if these points lie far apart in phase space. This is likely to adversely affect the dynamical significance of the modes identified by any algorithm (including SSA) that

analyses the data based on their geometrical relationships. Time-lagged embedding is a familiar technique of alleviating this issue developed in the qualitative theory of dynamical systems [27, 13, 14, 15]. Under generic conditions, the image of x_t in embedding space H under the delayed-coordinate mapping, $x_t \mapsto X_t = (x_t, x_{t-\delta t}, \dots, x_{t-(q-1)\delta t})$ lies in a manifold which is in one-to-one correspondence with M , provided that the dimension of H is sufficiently large. Thus, given a sufficiently-long embedding window $\Delta t = (q-1)\delta t$, we obtain a representation of the data set such that geometrical neighborhoods in embedding space encode similar dynamical features. Augmenting the dimensionality of ambient space is also useful in applications involving non-autonomous dynamical systems, including systems with time-periodic statistics [28]. After data analysis has been completed, a collection of elements $\tilde{X}_t \in H$ may be projected to physical space by writing $\tilde{X}_t = (\hat{x}_{t,0}, \hat{x}_{t,\delta t}, \dots, \hat{x}_{t,(q-1)\delta t})$, and taking the average [29],

$$\tilde{x}_t = \sum \hat{x}_{t',\tau}/q, \quad \text{s.t. } t' - \tau = t. \quad [1]$$

Riemannian formulation. Recall that SSA is essentially an SVD decomposition,

$$X_t = \sum_k u_k \sigma_k v_k(t), \quad u_k \in H, \quad \sigma_k > 0, \quad v_k \in L^2(T), \quad [2]$$

of the signal into bi-orthonormal “topos” (spatial) and “chronos” (temporal) modes, u_k and v_k [10]. In the above, T denotes the set of times at which observations are made, and $L^2(T)$ is the Hilbert space of square-integrable scalar functions on that set. That is, the signal is interpreted as a linear map $X : L^2(T) \mapsto H$ from the topos to the chronos spaces, given by

$$y = X(f) = \int dt X_t f(t), \quad \text{with } f \in L^2(T). \quad [3]$$

The u_k and v_k are eigenfunctions of the spatial and temporal covariance operators,

$$L(y)(\vec{r}) = X X^*(y)(\vec{r}) = \int d\vec{r}' C(\vec{r}, \vec{r}') y(\vec{r}'), \quad [4a]$$

$$R(f)(t) = X^* X(f)(t) = \int dt' \rho(X_t, X_{t'}) f(t'), \quad [4b]$$

respectively, where \vec{r} denotes position in the spatial domain (taking embedding into account), X^* is the adjoint of X , and C and ρ are the spatial and temporal two-point functions in embedding space, $\rho(X_t, X_{t'}) = \int d\vec{r} X_t(\vec{r}) X_{t'}(\vec{r})$ and $C(\vec{r}, \vec{r}') = \int dt X_t(\vec{r}) X_t(\vec{r}')$, respectively. In particular, the temporal patterns $v_k(t)$ correspond to time-dependent projection coefficients along the directions of highest signal variance in H . It is well-known, however, that this type of projection is likely to yield poor representation of the data if M is not flat (linear) [19, 30, 31]. NLSA addresses this shortcoming through a generalization of [3] to explicitly account for the nonlinear geometrical structure of M .

For simplicity in notation and exposition, we first consider the case where M is a differentiable, compact manifold with no boundary, equipped (through its embedding in H) with a smooth metric tensor g , Riemannian density (volume) $\mu = (\det g)^{1/2}$, and Laplace-Beltrami operator $\Delta f = -\text{div grad } f$ [32]. These conditions guarantee the well-posedness of the eigenvalue problem

$$\Delta \phi_k = \lambda_k \phi_k, \quad \int d\mu(X_t) \phi_i(X_t) \phi_j(X_t) = \delta_{ij}, \quad [5]$$

with $0 = \lambda_0 < \lambda_1 \leq \lambda_2 \leq \dots$, giving rise to an orthonormal basis $\{\phi_0, \phi_1, \dots\}$ for the Hilbert space $L^2(M, \mu)$ of square-integrable functions on M with measure μ . Even though M

may actually be neither differentiable, nor compact, in practice we will always be dealing with graph-theoretic, coarse-grained approximations of M constructed through finite observation samples, where the discrete analog of [5] holds [33], irrespective of the properties of M in the limit of infinite samples. Rigorous theory providing links between the discrete and continuous cases has been developed using heat kernels [34, 20, 35, 36].

With this machinery in place, the key modification of NLSA relative to SSA is to replace the chronos space $L^2(T)$ with the l -dimensional subspace of $L^2(M, \mu)$ spanned by the leading l Laplace-Beltrami eigenfunctions in [5], viz., $V_l = \text{span}\{\phi_0, \phi_1, \dots, \phi_l\}$. Thus, the temporal modes in NLSA are generated by a natural set of basis functions on the nonlinear data set, with length scale (“resolution”) on M controlled by the parameter l . As we discuss in more detail below, l plays a regularization role, to prevent overfitting when sampling of M is not dense. Next, by direct analogy with [3], we introduce a family of linear maps $A^l : V_l \mapsto H$, providing the link between the chronos and topos spaces of NLSA:

$$y = A^l(f) = \int d\mu(X_t) f(X_t) X_t \quad [6]$$

The strategy put forward in this paper is to recover the salient spatio-temporal modes associated with nonlinear signals through spectral decompositions of A , as follows.

Let $\{\phi_0, \dots, \phi_l\}$ be the orthonormal basis of V_l from [5], and $\{e_1, e_2, \dots\}$ an orthonormal basis of H . The matrix elements of A with this choice of bases are

$$A^l_{ij} = \langle A(\phi_j), e_i \rangle = \int d\mu(X_t) \phi_j(X_t) X_t^i, \quad [7]$$

where $\langle \cdot, \cdot \rangle$ is the inner product of H , and $X_t^i = \langle X_t, e_i \rangle$. This leads to a spectral decomposition of A^l ,

$$A^l_{ij} = \sum_{k=1}^r u_{ik} \sigma_k v_{jk}, \quad r = \text{rank } A^l \leq \min\{l, \dim H\}, \quad [8]$$

where u_{ik} and v_{jk} are matrix elements of unitary operators on H and V_l , respectively, and $\sigma_k > 0$ are singular values. The decomposition of the signal X_t in terms of the chronos and topos modes of A^l , analogous to [2], is then $X_t^i = \sum_{k=1}^r \tilde{X}_t^k$, with $\tilde{X}_t^k = u_k \sigma_k v_k(t)$ and

$$u_k = \sum_{i=1}^{\dim H} e_i u_{ik}, \quad v_k(t) = \sum_{j=1}^l v_{jk} \phi_j(X_t). \quad [9]$$

It is a standard result that $\tilde{A}^l_{ij} = \sum_{k=1}^K u_{ik} \sigma_k v_{jk}$ with $K < r$ is the rank- K linear map from V_l to H minimizing the Frobenius norm of the residual, $\|A^l - \tilde{A}^l\|_F^2 = \sum_{i,j} (A^l_{ij} - \tilde{A}^l_{ij})^2$. Moreover, by completeness, X_t^l converges to the input signal X_t as $l \rightarrow \infty$ (or $l = s$ for finite number of samples s).

To gain insight on the properties of these modes, as well as their relation to the corresponding modes from SSA, it is instructive to study the action of the spatial and temporal covariance operators, $L = A^l A^{l*}$ and $R = A^{l*} A^l$, on spatial and temporal patterns, $y(\vec{r}) \in H$ and $f(X_t) \in V_l$, respectively. The latter can be expressed as convolution operations, $L(y)(\vec{r}) = \int d\vec{r}' \mathcal{L}(\vec{r}, \vec{r}') y(\vec{r}')$ and $R(f)(X_t) = \int d\mu(\tilde{X}) \mathcal{R}(X_t, \tilde{X}) f(\tilde{X})$, with kernels given by

$$\mathcal{L}(\vec{r}, \vec{r}') = \int d\mu(X_t) d\mu(\tilde{X}) \mathcal{K}_l(X_t, \tilde{X}) X_t(\vec{r}) \tilde{X}(\vec{r}'), \quad [10a]$$

$$\mathcal{R}(X_t, X'_t) = \int d\mu(\tilde{X}) \mathcal{K}_l(X_t, \tilde{X}) \rho(\tilde{X}, X'_t), \quad [10b]$$

and $\mathcal{K}_l(X_t, \tilde{X}) = \sum_{k=1}^l \phi_k(X_t) \phi_k(\tilde{X})$. In particular, take the limit $l \rightarrow \infty$, where $\mathcal{K}_l(X_t, \tilde{X}) = \delta(X_t - \tilde{X})$ becomes a Dirac δ -function by completeness, we have $\mathcal{L}(\vec{r}, \vec{r}') \rightarrow \int d\mu(X_t) X_t(\vec{r}) X_t(\vec{r}')$ and $\mathcal{R}(X_t, X'_t) \rightarrow \rho(X_t, X'_t)$.

Consider now the corresponding kernels in SSA from [4]. If the Riemannian measure μ is uniform, and the system trajectory covers M densely over the observation time interval T (two very strong assumptions), it is possible to replace the integral over time in [4b] with an integral over M . Therefore, in this idealized case, the results from SSA and NLSA coincide. In typical applications, however, neither the Riemannian measure μ is uniform, nor the coverage of M is dense. In NLSA, the latter deficiency is alleviated by truncating V_l to a finite dimensional subspace of $L^2(M, \mu)$. That is, given a finite number of observations, the results of NLSA up to some upper value of l should be close to the corresponding results obtained in the infinite-sample limit (at the same l), and still reflect the coarse-grained nonlinear geometric properties of M . Moreover, non-uniformity of μ , which for finite data sets depends both on the geometrical properties of the embedding $M \mapsto H$ and the dynamics on M , is handled naturally by graph-theoretic algorithms. This will turn out to play a key role in the atmospheric application discussed below, where sampling of the attractor is highly inhomogeneous.

Evaluating spatio-temporal patterns of finite data sets.

In order to implement the spectral decomposition [8] in practical applications involving finite data sets, we proceed by first using graph-theoretic algorithms to evaluate the Laplacian eigenfunction basis for V_l , then performing SVD of the matrix of operator components from [7]. Given a data set consisting of s samples, $G = \{X_{\delta t}, X_{2\delta t}, \dots, X_{s\delta t}\} \subset M$, graph-theoretic algorithms [19, 20] seek to construct a Markov process on G with transition probability matrix P , such that, for large-enough s and small-enough i , the right eigenvectors $\underline{\phi}_i$ of P approximate the corresponding Laplace-Beltrami eigenfunctions ϕ_i in [5]; i.e., $P \underline{\phi}_i = (1 - \lambda_i) \underline{\phi}_i$, with $\underline{\phi}_i = (\phi_{i1}, \dots, \phi_{is})^T$ and $\phi_{ij} \approx \phi_i(X_{j\delta t})$. These eigenvectors satisfy an orthonormality condition which is the discrete analog to [5], $\sum_{k=1}^s \mu_k \phi_{ik} \phi_{jk} = \delta_{ij}$, with μ_k given by the invariant measure (leading left eigenvector) of P ; namely, $\underline{\mu} = \underline{\mu} P$, where $\underline{\mu} = (\mu_1, \dots, \mu_s)$, $\mu_i > 0$, and $\sum_{\mu=1}^s \mu_i = 1$.

In the present work, we evaluate P using the diffusion map (DM) algorithm of Coifman and Lafon [20], with a simple but important modification in the calculation of the Gaussian weight matrix. Specifically, we assign to each sample $X_{i\delta t}$ a local velocity in embedding space, $\xi_i = \|X_{i\delta t} - X_{(i-1)\delta t}\|$, and evaluate the Gaussian weights $W_{ij} = \exp(-\|X_{i\delta t} - X_{j\delta t}\|^2 / (\epsilon \xi_i \xi_j)^{1/2})$, where $\|\cdot\|$ denotes the norm of H . This approach was motivated by the clustering algorithm developed in [37], with the difference that in the latter paper ϵ_i is evaluated using spatial nearest neighbors, rather than the temporal nearest neighbors employed here. In the standard implementation of DM, ϵ must be sufficiently small in order for the diffusion process represented by P to be sensitive only to local neighborhoods around each data point. Here, the normalization by ξ_i enforces geometrical localization even for $\epsilon = O(1)$. The remaining standard steps needed to compute P given W are displayed in Algorithm S2. The scalability of this class of algorithms to large problem sizes has been widely demonstrated in the machine learning and data mining literature. In particular, as a result of Gaussian decay of the weights, the W matrix used in implementations is made sparse, e.g., by truncating W to the largest b nonzero elements per row with $b/s \ll 1$, significantly reducing the cost of the eigenvalue problem for $\underline{\phi}_i$. The least-favorable scaling involves the pairwise

distance calculation between the data samples in embedding space, which scales like $s^2 \dim(H)$ if done in brute force (which is the method employed here). Despite the quadratic scaling with s , the linear scaling with $\dim(H)$ is of key importance, as it guarantees that NLSA does not suffer from a “curse of dimension” as do neural-network-based methods [38].

North Pacific data from a coupled climate model

We first study the variability in the North Pacific sector of CCSM3; specifically, variability of the mean upper 300 m sea temperature field in the 700-year equilibrated integration used in [22, 23] in work on the initial-value and forced-response predictability in that model. Here, our objective is to diagnose the prominent modes of variability in a comprehensive climate model. In this analysis, the x_t observable is the mean upper 300 m temperature field sampled every month at $n = 534$ gridpoints (native ocean grid mapped to the model’s T42 atmosphere) in the region 20°N – 65°N and 120°E – 110°W . Below, we work with an embedding window $\Delta t = 2$ years; i.e., the dimension of embedding space is $\dim(H) = n \times 24 = 12,816$. After embedding, $x_t \mapsto X_t$, we are left with $s = 700 \times 12 - 24 = 8376$ samples for data analysis. Here, we present results obtained using the parameter values $l = 23$, $\epsilon = 2$, and $b = 1000$. Figures 1 and 2 display representative singular values and temporal patterns of the A^l operator in [8], respectively, and snapshots of the temperature anomaly fields \tilde{x}_t^k (evaluated by applying [1] to the spatio-temporal patterns \tilde{X}_t^k from [9]) are shown in Figure 3 and Movie S1. We find that the spatio-temporal patterns fall into three distinct families of periodic, low-frequency, and intermittent modes, described below. We tested for robustness of these results in a series of calculations with $l \in [20, 100]$, $\epsilon \in [1, 10]$ and $b \in [100, 2000]$. Note that time-lagged embedding is essential to the separability of the spatio-temporal patterns into qualitatively-distinct processes; their character is mixed if no embedding is performed. Also, Δt is not directly related to the lowest frequencies captured by the temporal patterns $v_k(t)$, but must be long-enough to incorporate information about major sources of non-Markovianity in the data—here the annual cycle of solar forcing. We have tested our results for robustness using several $\Delta t \in [2, 10]$ years.

Periodic modes. Modes in this family come in doubly-degenerate pairs, and have the structure of sinusoidal waves with phase difference $\pi/2$ and frequency equal to integer multiples of 1/year. The leading periodic modes, $\{v_1, v_2\}$ [Figure 2(a)], represent the seasonal cycle in the data. In the spatial domain [Figure 3(a), Movie S1(c)], these modes generate an annual oscillation of the temperature anomaly, whose amplitude is largest ($\sim 1^\circ\text{C}$) in the western part of the basin (130°E – 160°E) and for latitudes in the range 30°N – 45°N . The semiannual modes [Figure 3(c), Movie S1(e)] exhibit significant amplitude in the western part of the domain, but also along the West Coast of North America; the latter is consistent with semiannual variability of the upper ocean associated with the California current [39].

Low-frequency modes. These modes are characterized by high spectral power over interannual to interdecadal timescales. Featuring strongly suppressed power over annual or shorter time scales, they represent the low-frequency variability of the upper ocean, which has been well-studied in the North Pacific sector of CCSM3 [21, 22]. The leading mode in this family $\{v_3\}$; see Figure 2(b)], gives rise to a typical PDO pattern [Figure 3(b), Movie S1(d)], where the most prominent basin-scale structure is a horseshoe-like temperature anomaly pattern developing eastward along the Kuroshio extension, to-

gether with an anomaly of the opposite sign along the west coast of North America. The higher modes in this family gradually develop smaller spatial features and spectral content over shorter time scales than v_3 , but have no spectral peaks on annual or shorter timescales. As illustrated in Movie S1(b), these modes also arise in classical SSA.

Intermittent modes. A major result of this analysis, which highlights the importance of taking explicitly into account the nonlinear geometrical structure of complex spatio-temporal data sets, is the existence of intermittent patterns of variability in the North Pacific sector of CCSM3, which are not accessible through SSA. As illustrated in Figure 2(c), the key feature of modes of this family is temporal intermittency, arising out of oscillations at annual or higher frequency, which are modulated by relatively sharp envelopes with a temporal extent in the 2–10-year regime. Like their periodic counterparts, the intermittent modes form nearly degenerate pairs, and their base frequency of oscillation is an integer multiple of 1/year. The resulting Fourier spectrum is dominated by a peak centered at at the base frequency, exhibiting some skewness towards lower frequencies. In the physical domain, these modes describe processes with relatively fine spatial structure, which are activated during the intermittent bursts, and become quiescent when the amplitude of the envelopes is small. The most physically-recognizable aspect of these processes is enhanced transport along the Kuroshio extension region, shown for the leading-two intermittent modes $\{v_9, v_{10}\}$ in Figure 3(d) and Movie S1(f). This process features sustained eastward propagation of small-scale, $\sim 0.2^\circ\text{C}$ temperature anomalies during the intermittent bursts. The intermittent modes higher in the spectrum also encode rich spatiotemporal patterns, including retrograde (westward) propagating anomalies, and gyre-like patterns resembling the subpolar and subtropical gyres.

A chaotic intermittent low-order atmosphere model

The spatial modes u_k from [2] arising in SSA are optimal from the point of view of capturing the maximal covariance in H using the smallest number of modes. However, this does not necessarily imply that the dynamical significance of these modes is correspondingly high. For instance, the SSA-based models of the Kuramoto-Sivashinsky (KS) equation by Aubry et al. [17] capture 99.99995% of the variance using only the leading six modes, but projecting the KS equation onto the subspace spanned by those modes fails to reproduce the right dynamics. Similar issues were encountered by Crommelin and Majda [18] in a study of a six-mode chaotic low-order model for the atmosphere [25, 26]. Those authors explored a suite of reduced-modeling approaches, among which only spatial patterns derived through principal interacting patterns (PIPs) [40, 5] were able to reproduce chaotic regime transitions. PIPs, however, make use of significant information about the dynamical system generating the observed data. Yet, spatial modes derived via purely data-driven methods, including SSA and OPPs, were found in [18] to exhibit weak dynamical significance. In particular, Galerkin-truncated models onto those bases typically decayed to fixed points, or locked on orbits of unrealistically high temporal regularity [18]. Here, we demonstrate that the spatial patterns from [8], evaluated through NLSA using the same input data as SSA or OPPs, are able to reproduce chaotic regime transitions.

Specifically, setting $x_t = (x_1, \dots, x_6) \in \mathbb{R}^6$, we consider a deterministic dynamical system of the form $\dot{x}_t = F(x_t)$, with the governing equations and parameter values given in [18], and reproduced in Table S1 for convenience. Here, the x_i ’s are low-order Fourier expansion coefficients of the streamfunction in the barotropic vorticity equation for the atmosphere,

and non-dimensional time corresponds to one day. Physically, this model describes chaotic transitions between zonal and blocked regimes of the atmospheric streamfunction, caused by large-scale topography (e.g., continental mountain ranges) and barotropic instability [26]. For our purposes, the modes that most-closely represent these transitions are x_1 and x_4 . The latter correspond to purely zonal wind (East-West), with no variation in the meridional direction, forced by a zonal atmospheric jet profile as observed in the real atmosphere; see Figures 4(a) and 5(a) for representative time series.

The goal of dimensional reduction is to produce a set of orthonormal spatial patterns for Galerkin projection, $u_k \in \mathbb{R}^6$ with $k \leq K < 6$, such that the reduced model projected onto these patterns, $\dot{z}_t = U^T F(Uz_t)$ with $z_t \in \mathbb{R}^K$, exhibits chaotic transitions between zonal and blocked states. Note that by the orthogonality property, $U^T U = I$, a solution z_t generates a trajectory $x_t = Uz_t$ in the phase space of the full model. Here, we constructed Galerkin bases using NLSA with a training time series x_t of length 10^6 days, sampled daily. Note that because the full phase space of the model is accessible, no embedding was performed in these calculations, i.e., $\dim(H) = 6$. We ran NLSA (Algorithm S1) using $\epsilon = 1$ throughout, and representative values of $b \in [30, 500]$ and $l \in [3, 20]$. In each case, we carried out the decomposition in [8] and formed $6 \times K$ projection matrices $U = [u_{i1}, \dots, u_{iK}]$ for $K = 3, 4$, or 5 . We tested the resulting reduced models over long, 10^6 -day integrations with initial conditions chosen randomly from the training data.

With moderate experimentation in the choice of b and l , it was possible to construct reduced models featuring chaotic regime transitions for all $K \in \{3, 4, 5\}$; examples of these models are illustrated in Figures 4 and 5 and Movie S2. Qualitatively, the fidelity of these models increases with K ; e.g., the $K = 5$ model reproduces excursions towards large (x_1, x_4) values, i.e., away from the regions of high joint probability $p(x_1, x_4)$ [26], though with some bias. Note that the objective here is not to reproduce the trajectory of the full model in a pointwise sense; that would be futile (and of little use) given the system's chaotic dynamics. The fact that models with dimension as low as $K = 3$ feature chaotic transitions between zonal and blocked regimes strongly supports the dynamical significance of NLSA modes.

A further important point concerns the behavior of the discrete Riemannian measure $\bar{\mu}$. As illustrated in Figure 5(c), when plotted as a function of time, the weights μ_i exhibit a bursting behavior, with spikes corresponding to transitions between regimes. Thus, the algorithm assigns high weight to transitory events in the operator components in [6], which compensates for the fact they are observed only rarely. Heuristically, μ_i corresponds to the volume in M occupied by the sample X_t at time $t = i\delta t$, which tends to be large for states visited infrequently by the system. In contrast, the corresponding time integral [4b] in classical SSA, carries uniform temporal weights, and the influence of the short-time tran-

sitions in the resulting spatial patterns is suppressed. The non-uniform Riemannian weighting thus enhances the capability of NLSA of capturing rare events, which frequently turn out to be crucial in reproducing the right dynamics [17, 18].

Concluding remarks

Combining techniques from machine learning and the qualitative theory of dynamical systems, we have presented a method for time series analysis, referred to as NLSA, which explicitly accounts for the nonlinear geometrical structure of data sets generated by complex dynamical systems. Like classical SSA, NLSA utilizes time-lagged embedding and SVD to produce a spatio-temporal decomposition of a signal generated by partial observations of high-dimensional dynamical systems. However, the linear operator used here in the SVD step differs crucially from SSA in that its domain of definition is the Hilbert space of square-integrable functions on the nonlinear manifold M comprised by the data (in a suitable coarse-grained graph representation). These functions, analogous to the chronos modes in SSA [10], are tailored to the geometry of M through its Riemannian measure μ .

Applying this scheme to the upper-ocean temperature in the North Pacific sector of a comprehensive climate model, we found a family of intermittent processes which are not captured by SSA. These processes describe eastward-propagating, small-scale temperature anomalies in the Kuroshio current, as well as retrograde-propagating structures at high latitudes and in the subtropics. These modes carry little variance of the raw signal, and display burst-like behavior characteristic of strongly nonlinear dynamics [17]. The remaining identified modes include the familiar PDO pattern of decadal variability, as well as annual and semiannual periodic processes.

Besides demonstrating NLSA's skill in capturing qualitatively-distinct spatio-temporal patterns, we have examined its utility in building high-quality bases for model reduction through Galerkin projection. As a particularly challenging example, involving strong intermittency and inhomogeneous sampling of the attractor, we considered a six-mode low-order model for the atmosphere featuring chaotic transitions between zonal and blocked regimes [26]. While other data-driven approaches fail to reproduce chaotic regime behavior [18], Galerkin modes based on NLSA led to reduced models of dimension as low as 3 featuring the salient aspects of regime behavior of the full model. More generally, NLSA is well-poised to capture short-time transitions and rare events via non-uniform weighting of the data through the Riemannian volume. This feature is encouraging for potential future research on NLSA-based empirical models for prediction in large-scale models, such as coupled climate models.

ACKNOWLEDGMENTS. This work was supported by NSF grant DMS-0456713, ONR DRI grants N00014-10-1-0554 and N0014-11-1-0306. We thank G. Branstator and H. Teng for providing access to the CCSM3 data set used in this analysis.

- Holmes, P., Lumley, J. L., & Berkooz, G. (1996) *Turbulence, Coherent Structures, Dynamical Systems and Symmetry*. (Cambridge University Press, Cambridge), p. 440.
- Sreenivasan, K. R. & Antonia, R. A. (1997) The phenomenology of small-scale turbulence. *Annu. Rev. Fluid Mech.* 29, 435–472.
- Vautard, R. & Ghil, M. (1989) Singular spectrum analysis in nonlinear dynamics, with applications to paleoclimatic time series. *Phys. D* 35, 395–424.
- DeSole, T. (2001) Optimally persistent patterns in time-varying fields. *J. Atmos. Sci.* 58, 1341–1356.
- Kwasniok, F. (2004) Empirical low-order models of barotropic flow. *J. Atmos. Sci.* 61, 235–245.
- Deuffhard, P., Dellnitz, M., Junge, O., & Schütte, C. (1999) Computation of essential molecular dynamics by subdivision techniques I: basic concept. *Lect. Notes Comp. Sci. Eng.* 4, 98.
- Nadler, B., Lafon, S., Coifman, R. R., & Kevrikedes, I. (2006) Diffusion maps, spectral clustering, and reaction coordinates of dynamical systems. *Appl. Comput. Harmon. Anal.* 21, 113–127.
- Brunt, C. M. & Heyer, M. H. (2002) Interstellar turbulence. I. Retrieval of velocity field statistics. *Astrophys. J.* 566, 276–288.
- Cacadaid, A. C., Lawrence, J. K., & Ruzmaikin, A. (2008) Principal components and independent component analysis of solar and space data. *Solar Phys.* 248, 247–261.
- Aubry, N., Guyonnet, R., & Lima, R. (1991) Spatiotemporal analysis of complex signals: Theory and applications. *J. Stat. Phys.* 64, 683–739.
- Golyandina, N., Nekrutkin, V., & Zhigljavsky, A. (2001) *Analysis of Time Series Structure: SSA and Related Techniques*. (CRC Press, Boca Raton), p. 320.
- Ghil, M. et al. (2002) Advanced spectral methods for climatic time series. *Rev. Geophys.* 40.

13. Takens, F. (1981) in *Dynamical Systems and Turbulence*, Warwick 1980, Lecture Notes in Mathematics. (Springer, Berlin) Vol. 898, pp. 366–381.
14. Sauer, T, Yorke, J. A. & Casdagli, M. (1991) *Embedology*. *J. Stat. Phys.* 65, 579–616.
15. Deyle, E. R & Sugihara, G. (2011) Generalized theorems for nonlinear state space reconstruction. *PLoS ONE* 6, e18295.
16. Broomhead, D. S & King, G. P. (1986) Extracting qualitative dynamics from experimental data. *Phys. D* 20, 217–236.
17. Aubry, N, Lian, W.-Y, & Titi, E. S. (1993) Preserving symmetries in the proper orthogonal decomposition. *SIAM J. Sci. Comput.* 14, 483–505.
18. Crommelin, D. T & Majda, A. J. (2004) Strategies for model reduction: Comparing different optimal bases. *J. Atmos. Sci.* 61, 2206.
19. Belkin, M & Niyogi, P. (2003) Laplacian eigenmaps for dimensionality reduction and data representation. *Neural Comput.* 15, 1373–1396.
20. Coifman, R. R & Lafon, S. (2006) Diffusion maps. *Appl. Comput. Harmon. Anal.* 21, 5–30.
21. Alexander, M et al. (2006) Extratropical atmosphere–ocean variability in CCSM3. *J. Climate* 19, 2496–2525.
22. Teng, H & Branstator, G. (2010) Initial-value predictability of prominent modes of North Pacific subsurface temperature in a CGCM. *Climate Dyn.* 36, 1813–1834.
23. Branstator, G & Teng, H. (2010) Two limits of initial-value decadal predictability in a CGCM. *J. Climate* 23, 6292–6311.
24. Mantua, N. J & Hare, S. R. (2002) The pacific decadal oscillation. *J. Oceanogr.* 58, 35–44.
25. De Swart, H. E. (1989) Analysis of a six-component atmospheric spectral model: Chaos, predictability and vacillation. *Phys. D* 36, 222–234.
26. Crommelin, D. T, Opsteegh, J. D, & Vehrlost, F. (2004) A mechanism for atmospheric regime behavior. *J. Atmos. Sci.* 61, 1406.
27. Packard, N. H et al. (1980) Geometry from a time series. *Phys. Rev. Lett.* 45, 712–716.
28. Majda, A. J & Wang, X. (2010) Linear response theory for statistical ensembles in complex systems with time-periodic forcing. *Comm. Math. Sci.* 8, 145–172.
29. Horenko, I. (2008) On simultaneous data-based dimension reduction and hidden phase identification. *J. Atmos. Sci.* 65, 1941–1954.
30. Ham, D, Lee, D. D, Mika, S, & Schölkopf, B. (2003) A kernel view of the dimensionality reduction of manifolds, (Max Planck Institute for Biological Cybernetics), Technical Report TR-110.
31. Coifman, R. R et al. (2005) Geometric diffusions as a tool for harmonic analysis and structure definition on data. *Proc. Natl. Acad. Sci.* 102, 7426–7431.
32. Bérard, P. H. (1989) *Spectral Geometry: Direct and Inverse Problems*, Lecture Notes in Mathematics. (Springer-Verlag, Berlin) Vol. 1207, p. 272.
33. Chung, F. R. K. (1997) *Spectral Graph Theory*, CBMS Regional Conference Series in Mathematics. (American Mathematical Society, Providence) Vol. 97, p. 207.
34. Belkin, M & Niyogi, P. (2003) Using manifold structure for partial labelled classification. *Adv. NIPS* 15.
35. Belkin, M & Niyogi, P. (2008) Towards a theoretical foundation for Laplacian-based manifold methods. *J. Comput. Syst. Sci.* 74, 1289–1308.
36. Jones, P. W, Maggioni, M, & Schul, R. (2008) Manifold parametrizations by eigenfunctions of the laplacian and heat kernels. *Proc. Natl. Acad. Sci.* 105, 1803.
37. Zelnik-Manor, L & Perona, P. (2004) Self-tuning spectral clustering. Vol. 17, p. 1601.
38. Hsieh, W. W. (2007) Nonlinear principle component analysis of noisy data. *Neural Networks* 20, 434–443.
39. Mendelsohn, R, Schwing, F. B, & Bograd, S. J. (2004) Nonstationary seasonality of upper ocean temperature in the California Current. *J. Geophys. Res.* 109.
40. Kwasniok, F. (1997) Optimal Galerkin approximations of partial differential equations using principal interaction patterns. *Phys. Rev. E* 55, 5365–5375.




Fig. 1. Singular values σ_i (normalized so that $\sigma_1 = 1$) evaluated through NLSA and SSA for the CCSM3 ocean data, indicating the periodic, low-frequency, and intermittent modes.




Fig. 2. Temporal patterns $v_k(t)$ from [9], corresponding to the singular values in Figure 1, plotted in the temporal and frequency domains. (a) Annual periodic mode, v_1 ; (b) Leading low-frequency (PDO) mode, v_3 ; (c) Leading intermittent (Kuroshio) mode, v_9 .

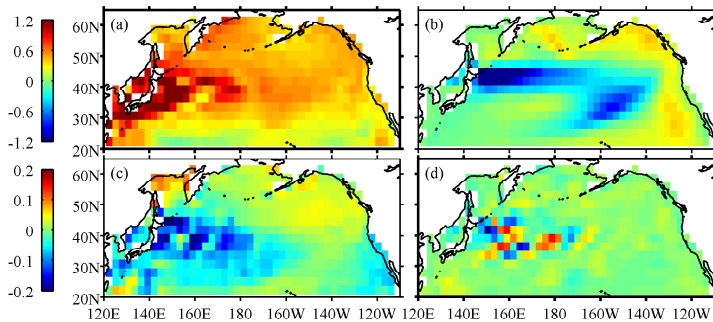


Fig. 3. Spatio-temporal patterns $\sum_k \tilde{x}_t^k$ of the upper 300 m temperature anomaly field (annual mean subtracted at each gridpoint, color-coded in $^{\circ}\text{C}$) evaluated using [9] and [1] for August of year 39 of Figure 2, using (a) the annual modes, $k \in \{1, 2\}$; (b) the first low-frequency mode, $k = 3$, describing the PDO; (c) the semiannual modes, $k \in \{6, 7\}$; (d) the leading two-fold degenerate set of intermittent modes, $k \in \{9, 10\}$, describing variability of the Kuroshio current. The dynamic evolution is much more revealing; see Movie S1.




Fig. 4. 1000-day trajectories for the zonal modes $\{x_1, x_4\}$ in (a) the full model (Table S1), showing approximate locations of the zonal and blocked states; (b–d) reduced models constructed by projection onto the NLSA modes u_k from [9]. See also Movie S2.




Fig. 5. 2000-day trajectories for mode x_1 in (a) the $K = 4$ reduced model; (b) the full model. (c) the Riemannian measure μ of the reduced model.