Archetypal Analysis: Mining Weather and Climate Extremes[✉]

A. HANNACHI

Department of Meteorology, Stockholm University, Stockholm, Sweden

N. TRENDAFILOV

Department of Mathematics and Statistics, The Open University, Milton Keynes, United Kingdom

(Manuscript received 9 November 2016, in final form 15 May 2017)

ABSTRACT

Conventional analysis methods in weather and climate science (e.g., EOF analysis) exhibit a number of drawbacks including scaling and mixing. These methods focus mostly on the bulk of the probability distribution of the system in state space and overlook its tail. This paper explores a different method, the archetypal analysis (AA), which focuses precisely on the extremes. AA seeks to approximate the convex hull of the data in state space by finding “corners” that represent “pure” types or archetypes through computing mixture weight matrices. The method is quite new in climate science, although it has been around for about two decades in pattern recognition. It encompasses, in particular, the virtues of EOFs and clustering. The method is presented along with a new manifold-based optimization algorithm that optimizes for the weights simultaneously, unlike the conventional multistep algorithm based on the alternating constrained least squares. The paper discusses the numerical solution and then applies it to the monthly sea surface temperature (SST) from HadISST and to the Asian summer monsoon (ASM) using sea level pressure (SLP) from ERA-40 over the Asian monsoon region. The application to SST reveals, in particular, three archetypes, namely, El Niño, La Niña, and a third pattern representing the western boundary currents. The latter archetype shows a particular trend in the last few decades. The application to the ASM SLP anomalies yields archetypes that are consistent with the ASM regimes found in the literature. Merits and weaknesses of the method along with possible future development are also discussed.

1. Introduction

Weather and climate data are witnessing an explosion in size from both observations and climate models, and as such various decomposition approaches are required to explore and analyze these large-scale data. One of the methods most used in weather and climate is empirical orthogonal function (EOF) analysis (Obukov 1947; Lorenz 1956), also known as principal component (PC) analysis. The EOF method seeks to decompose a space–time dataset into orthogonal EOF patterns and associated uncorrelated time series or PCs by maximizing the explained variance (Jolliffe 2002; Hannachi et al. 2007; Monahan et al. 2009). Other closely related methods

have also been used in atmospheric science (see, e.g., Jolliffe 2002; Hannachi et al. 2007; Wilks 2006). As will be discussed later these methods do not analyze extremes in any particular way. In this paper we present a relatively new method in climate research, the archetypal analysis (AA), whose most important feature is dealing precisely with extremes. AA is a pattern recognition method that finds points on the envelope of the data by minimizing a specific residual function. To better present AA we discuss it below within the context of the conventional methods.

The EOF method, like other techniques, has advantages and drawbacks, which we discuss briefly for convenience and, as a baseline background, to help understand the proposed technique in what follows. For example, the geometrical constraints, such as mixing and scaling (see, e.g., Hannachi et al. 2006; Hannachi 2007), impose limitations to what EOFs can achieve. By construction, EOFs are directions in state space and have no particular units and cannot therefore be compared directly to the actual

[✉] Supplemental information related to this paper is available at the Journals Online website: <http://dx.doi.org/10.1175/JCLI-D-16-0798.s1>.

Corresponding author: A. Hannachi, a.hannachi@misu.su.se

observations (the scaling problem). Only in a few cases can the EOFs be interpreted easily (in a probabilistic sense), and that is when the data are “well behaved” (e.g., normally distributed; see, e.g., Jolliffe 2002). The mixing problem is related to the fact that EOFs tend to mix physical patterns in order to maximize variance (e.g., Dommenges and Latif 2003; Mestas-Núñez 2000). A common method to overcome this is to use rotated EOFs (e.g., Wilks 2006; Hannachi et al. 2007), where a set of EOFs are rotated by maximizing a certain “simplicity” criterion to make the rotated EOFs more local. Another method was presented recently, namely, regularized EOFs (Hannachi 2016). The method overcomes the drawbacks of spatiotemporal orthogonality by solving a generalized eigenvalue problem and helps, in particular, overcome the mixing problem related to the leading mode of variability of sea level pressure anomalies (i.e., the North Atlantic Oscillation vs the Arctic Oscillation; see, e.g., Ambaum et al. 2001).

On the other hand EOFs are quite flexible. This flexibility comes with a price, namely the interpretation (Morup and Hansen 2012), owing to, for example, the existence of efficient algorithms such as the singular value decomposition (SVD) and their nested nature. The latter refers to the fact that EOFs can be ordered naturally, where the set of the leading N EOFs are a subset of the set of the leading M EOFs for $M > N$.

One way to get feature patterns that are similar to the measured data, and therefore lend themselves to an easy interpretation, is through clustering such as the k -means method, which attempts to identify clusters or equivalently the most representative, or typical prototype, entity. This is perhaps one of the most important features of clustering methods, which comes with a price, namely inflexibility related to the binary assignment of the data objects. There are other methods to identify clusters or regimes, notably those that are based on identifying regions of high density in state space via finding peaks in the system probability density function (pdf; e.g., Robertson and Ghil 1999; Hannachi and Turner 2013, hereafter HT13; Christiansen 2003) and mixture models (e.g., Hannachi 2010; Christiansen 2007). For more references and further details on this topic the reader is referred to the recent review of Hannachi et al. (2017). Similar to what we mentioned above, clustering is not the main focus here but is used to help understand the AA method (Morup and Hansen 2012). In clustering, for example, the original observations cannot be obtained as a direct combination of the cluster centroids. Note that this is also the case for EOF analysis, where it is not required that observations be approximated as a mixture (i.e., convex combination; with positive weights adding up to one) of the mode patterns nor that the EOF patterns resemble

the data (Cutler and Breiman 1994; Chan et al. 2003). Note also that using simple linear combinations (i.e., with positive and negative weights) can lead to states outside the data domain and is only useful for identifying directions not locations in state space. The advantage of positive weights (i.e., convex combination) is that they provide “physical” states within the data domain. It is perhaps important to recall here, and as outlined above, that EOFs and clustering do not treat extreme data in any special way. The study of extremes is indeed quite important in weather/climate research because of their impact on the environment, society, and infrastructure (see, e.g., Sura and Hannachi 2015; Sura 2011).

Therefore, expressing the decomposition modes (or basis vectors) directly in terms of the original variables and dealing particularly with extremes in the data are two desirable features that the climate scientist would like to have in any decomposition or analysis approach. AA, suggested by Cutler and Breiman (1994), attempts precisely to achieve this. In their original paper Cutler and Breiman (1994) introduced AA as an intermediate approach between, while combining the virtues of, clustering and PC analysis. As its name suggests, AA seeks to approximate the data in terms of “pure types,” or archetypes, which are themselves required to be a mixture of the observations.

The archetypes are obtained by estimating the convex hull or envelope of the data and are therefore characterized by features favoring representative corners of the data. AA has been applied mostly in pattern recognition, benchmarking and market research, physics (astronomy spectra), computer vision and neuroimaging, and biology. AA, however, is not well known in atmospheric research. In fact, the first application of AA in climate research was only introduced very recently by Steinschneider and Lall (2015), who applied it to daily precipitation in the United States.

This paper is an attempt to present AA to the climate community so that it may be explored alongside other conventional approaches to get the most out of large-scale weather and climate data. Our overall intention here is to present AA as a new tool that can help us learn a different facet of climate data. We perceive AA as a data compression tool suitable for pattern recognition, which may also shed some light on clustering of climate data. The manuscript is organized as follows. Section 2 provides a short review of AA. Section 3 describes briefly the algorithm used in the literature and summarizes the algorithm used in this paper to approximate the archetypes (section 3a) and its application to generated data (section 3b). Most technical details are provided in the supplemental material. Readers that are not familiar with technical background can skip section 3a. An application to sea surface

temperature (SST) and Asian summer monsoon is then presented in [section 4](#). The robustness and general issues of sensitivity of the method in both the generated samples and climate data are discussed in [section 5](#). A summary and conclusions are provided in the last section.

2. Archetypal analysis

a. The AA problem

To introduce the AA problem, it is convenient and helpful to briefly recall the conventional analysis tool, namely EOF analysis. Given an $n \times m$ (anomaly) data matrix $\mathbf{X} = (x_{ij})$ containing n observations in an m -dimensional space $\mathbf{x}_k = (x_{k1}, \dots, x_{km})^T$, $k = 1, \dots, n$ (i.e., $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_n]^T$), conventional EOF analysis seeks linear combinations of the variables [i.e., $\mathbf{X}\mathbf{u}$, with $\mathbf{u} = (u_1, \dots, u_m)^T$] that maximize variance {i.e., $\max[(\mathbf{X}\mathbf{u})^T(\mathbf{X}\mathbf{u})]$ }, leading to an eigenvalue problem:

$$\mathbf{X}^T \mathbf{X} \mathbf{u} = \lambda \mathbf{u}. \quad (1)$$

Equation (1) shows clearly that there is no requirement for the eigenvector to directly relate or resemble the original observations or any combination of them.

Unlike EOF analysis AA finds “typical” patterns $\mathbf{z}_1, \dots, \mathbf{z}_p$ (locations in state space) “resembling” the original observations, used as basis vectors to approximate, through a mixture or convex combination, the data matrix by minimizing the residuals Res:

$$\text{Res} = \sum_{t=1}^n \left\| \mathbf{x}_t - \sum_{k=1}^p a_{kt} \mathbf{z}_k \right\|^2. \quad (2)$$

The mixture coefficients in Eq. (2) are required to satisfy

$$\begin{aligned} a_{jt} &\geq 0, \quad t = 1, \dots, n \quad \text{and} \quad j = 1, \dots, p \\ \sum_{k=1}^p a_{kt} &= 1, \quad t = 1, \dots, n. \end{aligned} \quad (3)$$

The patterns $\mathbf{z}_1, \dots, \mathbf{z}_p$ are the archetypes. They themselves are required to be also convex combination of the original observations; that is,

$$\begin{aligned} \mathbf{z}_k &= \sum_{t=1}^n b_{tk} \mathbf{x}_t, \\ b_{tk} &\geq 0 \quad \text{and} \quad \sum_{t=1}^n b_{tk} = 1. \end{aligned} \quad (4)$$

The archetypes and the mixture weights are normally expressed, respectively, in terms of the $m \times p$, $p \times n$, and $n \times p$ matrices as follows:

$$\mathbf{Z} = [\mathbf{z}_1, \dots, \mathbf{z}_p], \quad \mathbf{A} = (a_{ij}), \quad \text{and} \quad \mathbf{B} = (b_{ik}). \quad (5)$$

The AA problem is then cast in terms of the following optimization problem:

$$\{\mathbf{A}, \mathbf{B}\} = \underset{\mathbf{A}, \mathbf{B}}{\operatorname{argmin}} \|\mathbf{X} - \mathbf{A}^T \mathbf{B}^T \mathbf{X}\|_F^2, \quad (6)$$

subject to the convex (or stochasticity) constraints, given by

$$\mathbf{A}, \mathbf{B} \geq 0, \quad \mathbf{A}^T \mathbf{1}_p = \mathbf{1}_n, \quad \text{and} \quad \mathbf{B}^T \mathbf{1}_n = \mathbf{1}_p, \quad (7)$$

where $\mathbf{1}_r = (1, \dots, 1)^T$ is the vector of ones in r dimensions. In Eq. (6) the notation $\|\cdot\|_F$ stands for the Frobenius matrix norm given by

$$\|\mathbf{Y}\|_F^2 = \operatorname{tr}(\mathbf{Y}\mathbf{Y}^T), \quad (8)$$

where $\operatorname{tr}(\cdot)$ is the trace operator. For a given p , Eq. (6) yields the stochastic matrices $\mathbf{A}$ and $\mathbf{B}$. The archetypes $\mathbf{Z} = [\mathbf{z}_1, \dots, \mathbf{z}_p]$ are then given by

$$\mathbf{Z} = \mathbf{X}^T \mathbf{B}. \quad (9)$$

Equation (4) along with the optimization problem [Eq. (6)] reveal that the archetypes are directly related to the observations and can therefore provide an easier interpretation compared to, for example, EOFs and closely related methods. Precisely, and unlike EOFs, Eq. (4) says that the archetypes are convex combinations of the observations. In addition, the observations themselves are also approximately a convex combination of the archetypes. The latter is somehow akin to EOFs except that the weights are not convex.

b. Geometrical interpretation of the archetypes

AA minimizes the residual sum of squares (RSS) $\mathcal{R}$:

$$\mathcal{R} = \|\mathbf{X} - \mathbf{A}^T \mathbf{B}^T \mathbf{X}\|_F^2, \quad (10)$$

with the archetypes being a convex combination of the observations. For a given p [Cutler and Breiman \(1994\)](#) show that the (exact or global) minimizers of RSS $\mathcal{R}$ [Eq. (10)] provide archetypes $\mathbf{Z} = [\mathbf{z}_1, \dots, \mathbf{z}_p]$ that are located on the boundary of the convex hull (or envelope) of the data. The convex hull of a given data point is the smallest convex set containing the data point. Archetypes therefore provide typical representations of the “corners” or extremes of the observations.

[Figure 1](#) shows a simple illustration (with no computation involved) of a two-dimensional example of a set of points with its convex hull and its approximation using five archetypes. The figure shows a set of two dimensional points, and the convex hull is also shown by dashed lines. An illustration of an approximate convex hull with five archetypes (vertices) is also shown. The dots inside this approximate hull, colored with red, do not contribute to the residual sum of squares but only

FIG. 1. Illustration showing a dataset of 25 points with their convex hull (dashed) and a hull approximation (continuous) with five archetypes (square). The set of points contributing to the residual sum of squares $\mathcal{R}$ are shown in purple with their projection on the hull approximation.

the dots outside (colored with blue). Note also that the approximate archetypes are not necessarily located on the edge of the convex hull as it is not known a priori and has to be approximated.¹ The sample mean $\bar{\mathbf{x}} = (1/n)\sum \mathbf{x}_i$ provides the unique archetype for $p = 1$, and that for $p = 2$ the pattern $\mathbf{z}_2 - \mathbf{z}_1$ coincides with the leading EOF of the data. Unlike EOFs archetypes are not required to be nested (Cutler and Breiman 1994; Bauckhage and Thureau 2009). However, like k -means clustering (and unlike EOFs) AA is invariant to translation and scaling and to rotational ambiguity (Morup and Hansen 2012).

A particularly nice, and unique, feature of AA is that it allows visualizing high-dimensional data on a two-dimensional plot using simplex visualization (Seth and Eugster 2016). The $p \times n$ weight matrix $\mathbf{A} = [\mathbf{a}_1, \dots, \mathbf{a}_n]$ is a probability matrix (i.e., $\mathbf{A} \geq 0$ and $\mathbf{A}^T \mathbf{1}_p = \mathbf{1}_n$). Therefore, the points $\mathbf{Z}\mathbf{a}_t$, $t = 1, \dots, n$, of the data matrix $\mathbf{A}^T \mathbf{B}^T \mathbf{X} = \mathbf{A}^T \mathbf{Z}^T$ reside on a $(p - 1)$ simplex² and can be projected onto the two-dimensional plane via a skew orthogonal projection (Coexeter 1973) also known as a Petrie polygon (Fig. 2). In this projection the vertices represent the simplex vertices (i.e., archetypes here). One of the attractive features of this projection operator is that it allows any simplex, with any dimension, to be represented by a (two dimensional) circle containing its vertices and all the vertex pairs are connected by edges.

¹ The complexity of the convex hull computation grows exponentially with the dimension; see, for example, Barber et al. (1996), although Wang et al. (2013) propose a polynomial complexity algorithm.

² A $(p - 1)$ simplex is a $(p - 1)$ -dimensional polytope obtained through convex combination of a given p (affinely) independent points (or vertices) $\mathbf{x}_1, \dots, \mathbf{x}_p$ [in $(p - 1)$ dimensions], that is, combinations of the form $\sum_{k=1}^p \theta_k \mathbf{x}_k$ with $\theta_k \geq 0$, $k = 1, \dots, p$, and $\sum_{k=1}^p \theta_k = 1$. Examples include a line segment, a triangle in three dimensions, and tetrahedron in four dimensions.

This projection can be useful particularly in the context of clustering.

Another nice feature of AA, owing to the convex combination, is the probabilistic interpretation of the weights $\mathbf{A} = (a_{ij})$. Because $\mathbf{A}$ is a probability matrix ($a_{ij} \geq 0$ and $\sum_i a_{ij} = 1$), it permits a (soft) classification of the observation $\mathbf{x}_i$ to one of the archetypes based on the membership probability $a_{ij} = \Pr(\mathbf{z}_i | \mathbf{x}_j)$ of $\mathbf{x}_j$ to archetype $\mathbf{z}_i$. So AA may be used in a number of cases as a clustering tool. But one should keep in mind that archetypes represent corners rather than centroids found in conventional clustering methods.

To recap, AA attempts to identify prototypes of the data extremes residing on the convex hull. They are a convex combination of the data, and the data themselves are approximated by a convex combination of the archetypes. For clustered data the archetypes may be associated with the clusters, particularly if the latter are located near the data border but are distinct from the cluster centroids.

3. Numerical solution of archetypes

a. Manifold-based algorithm

Most methods used to solve the AA problem (e.g., Seth and Eugster 2016; Morup and Hansen 2012; Porzio et al. 2008; Bauckhage and Thureau 2009; Steinschneider and Lall 2015; Chan et al. 2003) are essentially based on the original alternating constrained least squares algorithm of Cutler and Breiman (1994). The algorithm alternates between finding the optimal matrix $\mathbf{A} = (a_{ij})$ for fixed archetypes $\mathbf{Z}$ and finding the optimal archetypes $\mathbf{Z}$ for fixed $\mathbf{A}$. The algorithm has essentially four steps:

- 1) Determine $\mathbf{A}$, for fixed $\mathbf{Z}$, by solving a constrained least squares problem
- 2) Use the obtained $\mathbf{A}$, from step 1, to solve for the archetypes $\mathbf{Z}\mathbf{A} = \mathbf{X}^T$ [i.e., $\mathbf{Z} = \mathbf{X}^T \mathbf{A}^T (\mathbf{A}\mathbf{A}^T)^{-1}$]
- 3) Use the obtained archetypes, from step 2, to estimate the matrix $\mathbf{B}$ again by solving a constrained least squares problem
- 4) Obtain an update of $\mathbf{Z}$ through $\mathbf{Z} = \mathbf{X}^T \mathbf{B}$, then go to step 1 unless the residual sum of squares is smaller than a prescribed threshold

In this paper we propose a (natural) nonalternating algorithm that solves for $\mathbf{A}$ and $\mathbf{B}$ simultaneously. The algorithm is based on differential geometry (i.e., differential calculus on differentiable manifolds also known as Riemannian manifolds).

The topic of optimization on manifolds, also referred to as Riemannian optimization, is gaining ground because many nonlinear problems can be cast in terms of manifolds in addition to the elegance of the theory

FIG. 2. Examples of skew projection of simplexes for a (a) 2 simplex, (b) 4 simplex, and (c) 10 simplex.

behind it (see, e.g., [Smith 1994](#); [Absil et al. 2010](#)). The objective is to seek a solution to the minimization problem

$$\min_{x \in \mathcal{M}} f(x), \quad (11)$$

where $\mathcal{M}$ is a differentiable manifold. Examples of such manifolds include the $(n-1)$ -dimensional sphere S^{n-1} , a submanifold of $\mathbb{R}^n$. Of particular interest here is the set of matrices with unit-vector rows (or columns):

$$\text{Ob}(n, m) = \{\mathbf{X} \in \mathbb{R}^{n \times m}, \text{ddiag}(\mathbf{X}^T \mathbf{X}) = \mathbf{I}_m\},$$

where $\text{ddiag}(\mathbf{Y})$ is the double diagonal operator, which transforms a square matrix $\mathbf{Y}$ into a diagonal matrix with the same diagonal elements as $\mathbf{Y}$. This manifold is known as the oblique manifold and is topologically equivalent to the Cartesian product of spheres (see the online supplemental material), with a natural inner product given by

$$\langle \mathbf{X}, \mathbf{Y} \rangle = \text{tr}(\mathbf{X}\mathbf{Y}^T). \quad (12)$$

Let us denote by $\mathbf{A} \circ \mathbf{C}$ the Hadamard matrix product between two $m \times n$ matrices $\mathbf{A} = (a_{ij})$ and $\mathbf{C} = (c_{ij})$ (i.e., the elementwise product), so $\mathbf{A} \circ \mathbf{C} = (a_{ij}c_{ij})$. The conventional matrix product is denoted in the usual manner (i.e., $\mathbf{AC}$).

We are now in a position to cast our problem using the oblique manifolds [i.e., $\text{Ob}(p, n)$ and $\text{Ob}(n, p)$ as explained in the supplemental material]. Because of the convexity of the weights our problem can be transformed, after replacing $\mathbf{A}$ and $\mathbf{B}$ by $\mathbf{A} \circ \mathbf{A}$ and $\mathbf{B} \circ \mathbf{B}$, respectively (see the supplemental material), to yield

$$\begin{aligned} \mathcal{R} &= \|\mathbf{X} - (\mathbf{A} \circ \mathbf{A})^T (\mathbf{B} \circ \mathbf{B})^T \mathbf{X}\|_F^2 \\ &= \text{tr}(\mathcal{Z}) - 2\text{tr}(\mathcal{Z}\mathbf{W}) + \text{tr}(\mathcal{Z}\mathbf{W}^T\mathbf{W}), \end{aligned} \quad (13)$$

where $\mathcal{Z} = \mathbf{X}\mathbf{X}^T$ and $\mathbf{W} = (\mathbf{A} \circ \mathbf{A})^T (\mathbf{B} \circ \mathbf{B})^T$.

In Eq. (13) the matrices $\mathbf{A}$ and $\mathbf{B}$ are now in $\text{Ob}(p, n)$ and $\text{Ob}(n, p)$, respectively. So instead of solving Eq. (6) we now solve Eq. (13), where $\mathbf{A}$ and $\mathbf{B}$ in Eq. (6)

[satisfying Eqs. (3) and (4)] are now replaced, respectively, by $\mathbf{A} \circ \mathbf{A}$ and $\mathbf{B} \circ \mathbf{B}$ with $\mathbf{A}$ now in $\text{Ob}(p, n)$ and $\mathbf{B}$ in $\text{Ob}(n, p)$. The next step is to compute the gradient of the cost function [Eq. (13)]. The procedure of computing gradients on manifolds consists first in obtaining the gradient on the Euclidean space [e.g., $\mathbb{R}^{n \times m}$, which contains $\text{Ob}(n, m)$] followed by a projection onto the tangent space of the differential manifold, such as $T_{\mathbf{A}}\text{Ob}(p, n)$ of $\text{Ob}(p, n)$ at $\mathbf{A}$. The nice feature here is that all our derivatives can be calculated in matrix form, which greatly simplifies the programming side and helps toward efficient computation.

Let us denote $\mathbf{A}^2 = \mathbf{A} \circ \mathbf{A}$ and similarly for $\mathbf{B}$. Then we have the following expression of the gradient of the cost function $\mathcal{R}$ (see supplemental material):

$$\begin{aligned} \nabla_{\mathbf{A}} \mathcal{R} &= 4 \left[(\mathbf{B}^2)^T \mathcal{Z} (-\mathbf{I}_n + \mathbf{W}^T) \right] \circ \mathbf{A} \\ \nabla_{\mathbf{B}} \mathcal{R} &= 4 \left[\mathcal{Z} (-\mathbf{I}_n + \mathbf{W}^T) (\mathbf{A}^2)^T \right] \circ \mathbf{B} \end{aligned} \quad (14)$$

Finally, the projection of the gradient of $\mathcal{R}$ ($\nabla_{\mathbf{A}, \mathbf{B}} \mathcal{R}$) onto the tangent space of the oblique manifolds yields the final gradient $\text{grad}_{\mathbf{A}, \mathbf{B}} \mathcal{R}$:

$$\begin{aligned} \text{grad}_{\mathbf{A}} \mathcal{R} &= \nabla_{\mathbf{A}} \mathcal{R} - \text{Addiag}(\mathbf{A}^T \nabla_{\mathbf{A}} \mathcal{R}) \\ \text{grad}_{\mathbf{B}} \mathcal{R} &= \nabla_{\mathbf{B}} \mathcal{R} - \text{Bddiag}(\mathbf{B}^T \nabla_{\mathbf{B}} \mathcal{R}). \end{aligned} \quad (15)$$

After minimizing the cost function $\mathcal{R}$ using the expression of the gradient in Eq. (15), the archetypes are then given by

$$\mathbf{Z} = \mathbf{X}^T (\mathbf{B} \circ \mathbf{B}) = \mathbf{X}^T \mathbf{B}^2. \quad (16)$$

b. Numerical solution and illustration

The numerical solution of the AA problem [i.e., the minimization of $\mathcal{R}$ in Eq. (13)] is obtained by using either the conjugate gradient approach or the steepest descent method. The algorithm is run in Matlab, and we have used a quite useful toolbox, Manopt ([Boumal et al. 2014](#)). The toolbox requires the expression of the cost function and its gradient in addition to specifying the nonlinear manifold. The method has been illustrated

with simplified examples as shown next. The example consists of (standard) three-dimensional Gaussian clusters centered at the vertices of a three-dimensional polytope located, respectively, at $(0, 3, 0)$, $(-2.1, -2.2, 0)$, $(2.1, -2.1, 0)$, $(0, 0, 3)$, and $(0, 0, -3)$. Basically, the polytope is a reflected triangular pyramid with basis located on an equilateral triangle in the x - y plane. The total sample size used is 4000.

As described in section 2, AA is not nested, and therefore various values of p , the number of archetypes, are tested. For a given value of p the archetypes are obtained and the residual sum of squares are computed. Figure 3a shows a kind of scree plot of the (relative) RSS versus p . As Fig. 3a shows, the most probable number of archetypes can be suggested from the scree plot. Figure 3b shows the cost function $\mathcal{R}$ along with the gradient norm. The cost function reaches its floor during the first few hundred iterations, although the gradient norm continues to decrease with increasing iteration. Figure 3c shows the skew projection of the elements of the probability matrix $\mathbf{A}^2 = \mathbf{A} \circ \mathbf{A}$, namely the two-dimensional simplex projection where the archetypes hold the simplex vertices. The clusters are associated to some extent with the archetypes, although the centroids are, as expected, distinct from the archetypes.

4. Application to SST and Asian monsoon

a. SST archetypes

This section applies the AA method to SST to learn the typical patterns that represent the system. The data come from the Hadley Centre Sea Ice and Sea Surface Temperature³ dataset (HadISST). The data are a combination of globally complete monthly fields of SST and sea ice on a $1^\circ \times 1^\circ$ latitude–longitude grid from 1870 to date (Rayner et al. 2003). We restrict our analysis to the period January 1870–December 2014. The SST anomalies are computed by removing the monthly seasonal cycle climatology. The domain considered here is limited meridionally to the region 45.5°S – 45.5°N .

The conjugate gradient method is applied to compute the archetypes, but the steepest descent method also provides similar results. The algorithm is run with 100 random initial conditions, and the solution with the smallest $\mathcal{R}$ is chosen. The method is first applied to the nondetrended anomalies. The scree plot (Fig. 4) shows a breaklike (or elbowlike) feature at $p = 3$ suggesting three archetypes, which are discussed next. We note at the outset that, once the number p of archetypes is fixed,

FIG. 3. (a) Scree plot of the five Gaussian clusters, (b) cost function (thick) and gradient norm (thin) for five archetypes, and (c) the skew projection of the elements of the probability matrix obtained using five archetypes. Different colors are for different clusters.

the order or ranking of archetypes is arbitrary (i.e., any of the obtained patterns could be labeled archetype 1 for example). One can choose, however, different ways to label them such as the variance of the associated time

³ See www.metoffice.gov.uk/hadobs/hadisst/.

FIG. 4. Scree plot of the SST anomalies showing the relative RSS vs the number of archetypes.

series. Here the archetypes are simply labeled by referring to familiar patterns known in the literature.

Figure 5 shows the spatial patterns of these archetypes. The first two clearly represent El Niño (Fig. 5a) and La Niña (Fig. 5b) phases. Note that the anomalies in Fig. 5 are multiplied by 10. Note also that in EOF analysis the sign of the pattern is irrelevant. Unlike what is expected from EOFs, Figs. 5a and 5b are not reverse of each other, reflecting the asymmetry, or skewness, of El Niño–Southern Oscillation (ENSO) as a whole. In particular, the El Niño signal at the equator and the associated pattern in the central North and South Pacific are more than 1.5 times stronger than their La Niña counterparts.

The third archetype of SST anomalies (Fig. 5c) is quite interesting. It is dominated by the western boundary currents, namely the Kuroshio, Gulf Stream, and Agulhas Current. The Brazil Current and the East Australian Current are also clearly dominant. There is also another small spot of positive SST anomaly in the Gulf of Angola. This anomaly is most probably due to an extension of the Guinea Current where surface water from this current accumulates in the Gulf of Angola. It is also interesting to see signatures of positive SST anomalies near the California coast and Channel Islands, which are most probably due to the Southern California Countercurrent. It is quite interesting to see how such minute details are captured here by AA. It is known that western boundary currents are driven by major gyres, which transport warm tropical waters poleward along narrow, and sometimes deep, currents. These currents are normally fast and are referred to as the western intensification (e.g., Stommel 1948; Munk 1950). This indicates that these water boundary currents represent extreme events and are located on the outer boundary (or corners) in the system state space, and therefore they can be captured by AA as this latter is a method for “mining” the extremes.

The archetypes $\mathbf{Z} = \mathbf{X}^T \mathbf{B}^2$ represent the spatial patterns, somehow “equivalent” to EOF patterns. The matrices $\mathbf{A}^2$ and $\mathbf{B}^2$ are also equally important in providing more information on the data. For instance the matrix $\mathbf{B}^2$ provides the convex weights used to construct the archetypes $\mathbf{Z}$. Since we normally deal with not-too-small sample sizes these weights are expected to be dominated by a much smaller number compared to the sample size of observations.

Figure 6 shows the mixture weights for these archetypes. For the El Niño archetype the contribution comes from various observations scattered over the observation period and most notably from the first half of the record. Those events correspond to prototype El Niños, with largest weights taking place at the end of the nineteenth and early twentieth centuries. The contribution from El Niño events of 1982 and 1997 are also clear.

For the La Niña archetype there is a decreasing contribution with time, with most weights located in the first half of the record, with particularly high contribution from the event of 1916–17. One can also see contributions from La Niña events of 1955 and 1975. It is interesting to note that these contributing weights are clustered (in time) as reflected in Fig. 6. Unlike the previous two archetypes, the third, western current, archetype is dominated by the last quarter of the observational period starting around the late 1970s.

The time series of the archetypes, i.e., the columns of the stochastic matrix $\mathbf{A}^2$ show the “amplitudes” of the archetypes, somehow similar to the PCs, and are shown in Fig. 7. The time series of El Niño (Fig. 7a) shows slight weakening of the archetypes, although the events of early 1980s and late 1990s are clearly showing up. There is a decrease from the 1990s to the end of the record. Prior to about 1945 the signal seemed quite stationary in terms of strength and frequency. The time series of the La Niña archetype (Fig. 7b) shows a general decrease in the last 50 or so years. The signal was somehow “stationary” (with no clear trend) before about 1920. These time series can be compared with, for example, the Southern Oscillation index⁴ (SOI). The difference of these two time series agrees well with the SOI (Fig. S2) with a correlation coefficient of 0.6. There are of course some differences regarding the way the SOI is defined in addition to the fact that the archetypal time series are weights (between 0 and 1).

Unlike the previous two ENSO archetypes the western current archetype time series (Fig. 7c) shows a clear increasing trend starting immediately after an extended

⁴ See www.cgd.ucar.edu/cas/catalog/climind/soi.html.

a) Archetype 1

b) Archetype 2

c) Archetype 3

FIG. 5. The obtained three archetypes of the SST anomalies showing (a) El Niño, (b) La Niña, and (c) the western currents. The color bar interval is $2 \times 10^{-1} \text{ }^{\circ}\text{C}$ from -14 to $24 \times 10^{-1} \text{ }^{\circ}\text{C}$ for (a), from -16 to $8 \times 10^{-1} \text{ }^{\circ}\text{C}$ in (b), and from -4 to $18 \times 10^{-1} \text{ }^{\circ}\text{C}$ in (c).

period of weak activity around 1910. The trend is not gradual, with the existence of a period with moderate activity around the 1960s. The most recent period, from the late 1990s, shows strong activity.

The two-dimensional simplex projection is shown in Fig. 8. The whole sample comprises both the light-gray-shaded and colored points. The archetypes are represented by the vertices of the simplex, and the states with the largest weights (matrix $\mathbf{B}^2$) are color coded by the archetype number. So basically the colors represent the 200 points that are closest to each archetype. Figure 8, with no evidence of clusters, represents an example of archetypal analysis where the archetypes are not necessarily associated with clusters.

The analysis was also applied to the detrended SST anomalies after removing a linear trend. The corresponding scree plot is shown in Fig. 9. There is an elbow at $k = 2$ archetypes. The two archetypes (not shown) represent, respectively, El Niño and La Niña. Note that

in this case the direction linking the two archetypes represents the first EOF. The analysis suggests that El Niño and La Niña represent the primary extremal patterns and that the western boundary currents represent secondary extremal phenomena that are enhanced by the trend. Nevertheless, when the algorithm is applied to the detrended SST anomalies with $p = 3$, the result (not shown) indicates that the three archetypes represent, respectively, El Niño, La Niña, and a pattern containing the main western boundary currents with a cold tongue in the equatorial SST anomalies combined with positive anomalies in the South (around 30°S) and North (around 35°N) Pacific Ocean.

Research pertaining to the western boundary currents was virtually nonexistent and has just started to attract researchers' attention. A very recent paper by Yang et al. (2016) [see also the discussion by Seager and Simpson (2016)] analyzed and discussed the western boundary currents intensification. Yang et al. (2016)

FIG. 6. Mixture weights of the three archetypes of SST anomalies: (a) El Niño, (b) La Niña, and (c) the western currents.

performed a simple linear trend analysis and suggest that the intensification could be due to global warming. As pointed out by [Seager and Simpson \(2016\)](#) more research is needed to investigate and disentangle the reasons behind this intensification. A comprehensive study of the western boundary currents using other variables such as bathymetry (e.g., sea surface heights) and ocean net heat flux is under investigation but goes beyond the scope of the present paper.

b. Monsoon archetypes

We consider here the sea level pressure (SLP) field from the 40-yr European Centre for Medium-Range Weather Forecasts (ECMWF) Re-Analysis (ERA-40) project. The data span the period 1958–2001 over the monsoon region (20°S – 35°N , 50° – 150°E) during the summer monsoon season from June to September (JJAS). Daily SLP anomalies are constructed as in [HT13](#), by subtracting the mean daily seasonal cycle. The data are provided on a Gaussian grid with a total sample size of 5368. [HT13](#) used isometric mapping (Isomap) to study the nonlinear behavior of Asian monsoon and characterize the different monsoon regimes. They identified three regimes, namely, active and break

monsoon phases over the Indian continent and an active monsoon phase over the China Sea. Here we wish to investigate the archetypes of Asian summer monsoon and analyze any potential link between these archetypes and monsoon regimes. We follow [HT13](#) by using the detrended data obtained by removing a linear trend from the daily SLP anomalies.

The same AA procedure is applied as explained above. The scree plot (Fig. S3) shows some indication of a small break at $p = 3$, then a steeper decrease of the relative RSS with a (slight) second elbow at $p = 4$ archetypes with about $\text{RSS} = 62\%$. Note, however, that the elbow here is weaker than that of [Fig. 4](#) or [Fig. 9](#) owing to the noise level in the monsoon daily SLP anomalies compared with the monthly SST anomalies. We explore both ($p = 3$ and $p = 4$) solutions and also discuss the five-archetype solution. The three-archetype solution provides a kind of baseline for comparison with the three monsoon regimes of [HT13](#). [Figure 10](#) shows the three archetypes obtained with the detrended SLP anomalies. The first archetype ([Fig. 10a](#)) shows a high pressure anomaly located mainly over the East China Sea and eastern China. The second archetype ([Fig. 10b](#)) has a low pressure center sitting over most of the Indian

FIG. 7. As in Fig. 6, but for the time series of the archetypes.

Ocean and South Asia with its center of action sitting over India, the Arabian Sea, and the Bay of Bengal. Associated with this low pressure is the small high pressure center over the northwest Pacific. The last archetype (Fig. 10c) shows mainly a low pressure center over most of China, Cambodia, and the western North Pacific (WNP) but not over India and surrounding seas.

To compare these archetypes with the regimes of HT13, Fig. 11, which is similar to their Fig. 11, provides a summary of their regimes. Figure 11a shows the pdf of monsoon SLP anomalies within the two-dimensional reduced space spanned by the leading two Isomap time series x_1 and x_2 . The pdf is obtained using a three-Gaussian mixture model where the three regimes represent the centers of the three Gaussian components of the model. The individual Gaussian components of the mixture model are shown by their covariance ellipse shown by the solid line contour. These regimes represent, respectively, the WNP active phase (Fig. 11b) and the active phase (Fig. 11c) and the break phase (Fig. 11d) of the Asian summer monsoon. These SLP anomaly maps are obtained by averaging the 300 states closest to the individual centers of the Gaussian components of the mixture model. To have a full picture of

the monsoon dynamics, the corresponding averages of the 850-hPa wind anomalies are also computed and are overlaid on the SLP anomaly maps. The anticyclonic wind velocity, or negative relative vorticity (Fig. 11b), is

FIG. 8. Two-dimensional simplex projection (shaded and colored points) of the SST anomalies using three archetypes. The 200 points that are closest to each of the three archetypes are colored, respectively, red, blue, and black, and the remaining points are shown by light gray shading in the background.

FIG. 9. Scree plot of the detrended SST anomalies showing the relative RSS vs the number of archetypes.

clearly associated with the high pressure system, with clear enhancement over the northwest Pacific. The Somali jet is clearly enhanced for the active phase (Fig. 11c) with a particular low pressure system associated with cyclonic wind (positive vorticity) over the Bay of Bengal and most of the Indian continent. The break phase (Fig. 11d), on the other hand, is associated with anticyclonic wind and positive pressure anomaly over the Indian continent and the Arabian Sea. Note that the break phase over the Indian continent is associated with a low pressure system over NWP and the eastern parts of China.

There is a clear similarity between the three archetypes and the monsoon regimes. Archetype 1 represents the WNP active phase whereas archetypes 2 and 3 represent, respectively, the active and break phases of the Asian summer monsoon. Note the difference in amplitude between the archetypes and the regimes, reaching up to 6 times, as the latter represent centroids (Fig. 11a) whereas the former represent extremal states.

The association between the archetypes and the regimes can be analyzed again via the simplex projection. Figure 12 shows the simplex projection of the three-archetype model (gray-shaded dots), and the colors represent the 800 SLP anomaly states that are closest to the regimes of HT13 (Fig. 12). The red, blue, and black colors represent, respectively, the WNP, active, and break phases of the Asian monsoon. Figure 12 bears some resemblance of concept to Fig. 3c. Each cluster can be associated to one archetype but is, of course, distinct from it. The clusters are on the faces as well as inside the simplex, with a clear dominance of the active phase (Fig. 11a) of the Asian summer monsoon (blue color).

The archetypes have contribution from a number of observations as is reflected by the weights from the $\mathbf{B}^2$ matrix (Fig. 13). The WNP active phase archetype (Fig. 13a) has a main contribution from the years 1979,

FIG. 10. Solutions of the three-archetype problem applied to the detrended SLP anomalies over the Asian summer monsoon region. Contour interval is 1 hPa.

1983, 1967, and 1997. The active phase archetype (Fig. 13b) has the largest contributions from 1958, 1988, and other observations from the early 1960s and 1990s. The break phase archetype (Fig. 13c) has contributions mainly from the 1960s, early 1970s, late 1980s, and 1990s. It is interesting to note that the contributing periods for one archetype do not overlap with the contributing periods of another.

The (amplitude) time series of the archetypes are shown in Fig. 14. By inspecting Figs. 13 and 14 the contributing observations (Fig. 13) can be identified also

FIG. 11. (a) Three-component Gaussian mixture model of the leading two Isomap time series of the monsoon SLP anomalies showing the three centers A, B, and C and the SLP anomalies obtained by averaging (or compositing over) the closest 300 states to these centers and showing, respectively, (b) the WNP active phase and (c) the active phase and (d) the break phase regimes of the Asian monsoon. The centers A, B, and C in (a) correspond to the centers of the individual Gaussians (solid) of the three-component mixture model. Overlaid on the SLP anomaly maps are the associated 850-hPa wind anomaly averages (or composites). Contour interval is 0.5 hPa, and the maximum wind arrow represents 2.5 m s^{-1} . Adapted from HT13.

in Fig. 14 through their relatively high amplitudes. It is particularly interesting to see that the WNP active phase archetype (Fig. 14a) has a pronounced low-frequency signature (on the order of 4–5 years) compared to the active (Fig. 14b) or break (Fig. 14c) monsoon phases, associated possibly with a low-frequency signal originating from the Pacific Ocean and generated by ENSO. The break monsoon phase (Fig. 14c), in particular, does not show any low-frequency signature.

The four- and five-archetype solutions are also investigated for completeness and comparison. Figure 15 shows the four archetypes from the four-archetype model of the Asian summer monsoon. Despite the fact that archetypes are not nested, Fig. 15 shows that the three archetypes identified earlier are approximately reproduced where archetypes 1, 3, and 4 (Fig. 15) are

associated, respectively, with archetypes 2, 1, and 3 (Fig. 10). Note that archetype 2 (Fig. 15b) shows mainly a low pressure system over most of the contiguous South Asian mainland associated most probably with a thermal low due to landmass heating. The five-archetype solution has also been examined (not shown). Three archetypes (among the five) look quite similar to those of the three-archetype solution, suggesting that the three regimes of HT13 are consistent and represent an integral part of the Asian monsoon dynamics.

As for the three-archetype solution Fig. 16 shows the simplex projection of the SLP data, with colors referring to the monsoon regimes, associated with the four- and five-archetype solutions. As for Fig. 12, the colors refer again to the 800 SLP anomaly states closest to the regime centroids of HT13 (Fig. 11), and the remaining

FIG. 12. Simplex projection of the three-archetype solution of the daily detrended sea level pressure anomalies over the Asian monsoon region. The colors refer to the 800 states closest to each of the regime centroids of HT13 (WNP active phase in red, the active phase in blue, and the break phase in black) and the remaining points are shown by light gray shading in the background.

points are shown by light gray shading in the background. It can be seen that for both the four- (Fig. 16a) and five-simplex (Fig. 16b) projections the clusters are still associated with the corresponding archetype. For example, the blue cluster is associated with (or closer to, compared with the others) the top vertex in Fig. 16a and to the top-left vertex in Fig. 16b. Compared to Fig. 12, with the optimum number of archetypes, Fig. 16 has more than three archetypes and the simplexes are different, which affect the projection. In particular, the clusters get somehow stretched. We stress again, however, that the clusters are still reasonably associated with, but separated from, the three archetypes, particularly with four archetypes and to some extent with five archetypes (Fig. 16).

5. Discussion of the robustness of AA

In this section we discuss the sensitivity of the method. Since archetypes seek to approximate the convex hull and are therefore connected directly to extremes it is important to investigate the sensitivity of the method especially to outliers and sample size.

We start with the clustered sample of Fig. 3 discussed in section 3b. To investigate the effect of sample size we consider subsamples of different sizes ranging from $n = 4000$ (original size) to 250 by increments of 250, in addition to two extra subsamples with respective sizes 100 and 50, adding up to 18 subsamples. Figure S4a shows the scree plot of these 18 subsamples. The plot for the whole (original) sample (Fig. 3a) is also shown. One

can notice that the relative RSS at five archetypes for small samples is slightly larger than that of the whole sample. Overall, however, there is a clear consistency between the different samples on five archetypes. These archetypes for the different subsamples are calculated and plotted in Fig. S4b. There is clear consistency between the archetypes of the different samples.

To investigate the effect of outliers the original sample was first centered then sorted based on the norm of the data points. Extreme points, defined as those points above the top 5th percentile, are multiplied by a factor of 3 making a new sample with 5% outliers (Fig. S5). These outliers are made deliberately exaggerated to make the point clear. The scree plot of this new set (with outliers) is also shown in Fig. S4a. The relative RSS is small, but nonvanishing, for four archetypes. However, with five archetypes the relative RSS is very similar to that of the original sample. The archetypes of the new set are shown in Fig. S4b. Although these archetypes are different from those of the original sample (because of the outliers), they have some sort of association with those archetypes (and with the clusters).

A similar procedure is applied to the data used in section 4. The effect of sample size is investigated by using eight subsamples of sizes ranging from 1500 to 100 by steps of 200. The SST data were randomized to avoid any trend effect. Figure S6a shows the relative RSS in the scree plot of these eight subsamples along with the curve of the original SST anomaly data. The curves are very similar and clearly show an obvious break at $k = 3$. The associated archetypes are computed and, to ease visualization, are projected onto the leading three EOFs of the SST anomalies (Fig. S6b). There is clear evidence of consistency between the archetypes of the subsamples and those of the whole sample. To investigate the effect of extremes we construct three new datasets by removing, respectively, the top 1st, 2.5th, and 5th percentiles (defined as in the above example). The associated relative RSS are shown in the scree plot of Fig. S6a. They are hardly distinguishable from the remaining curves and are consistent with the relative RSS of the whole dataset. The corresponding archetypes (projected onto the SST EOFs) are displayed in Fig. S6b. The archetypes are again consistent with those of the original data and the subsamples. A similar procedure is applied to the detrended SST anomalies and also the monsoon data and similar results of consistency are obtained (not shown).

Last, we would like to discuss the AA results when the data have no structure (e.g., multivariate Gaussian). Christiansen (2007), for example, showed that a number

FIG. 13. Mixture weights of the three archetypes of the Asian monsoon (detrended) SLP anomalies: (a) WNP, (b) active, and (c) break monsoon phases.

of clustering methods detect fictitious clusters in such data. We reiterate here that AA is not meant to find clusters but finds an optimum set of archetypes approximating the convex hull. It is only when the data are clustered, with clusters located near the data border, that the archetypes are associated, but not identical, to the cluster centroids. So the application (e.g., to a multinormal sample) provides a priori an appropriate number of archetypes. The application to (standard) multivariate normal samples yields scree plots that depend on the problem dimension m . For small m , less than about $m_0 = 5$, the relative RSS plots show that the optimum number of archetypes is about $2m$ (Fig. S7). For those dimensions the multinormal sample has m principal axes (which also represent EOFs), and the archetypes come in pairs and each pair corresponds to two (extreme) states (along the corresponding axis). For larger than m_0 dimensions the curse of dimensionality affects the results, and the scree plots show gradually smooth relative RSS curves with increasing m with no clear optimum value. Of course, in the latter case one can still choose an appropriate number of archetypes,

based on a given threshold of relative RSS, such as 50% or 70%.

6. Summary and conclusions

The increase in the size and quality of climate data calls for more advanced data analysis tools to be explored. This call is motivated by the need to understand not only the bulk of the probability distribution function of the climate system but also its tail or extremes. Conventional methods used in climate science are associated mostly with EOF analysis one way or another. EOFs have a number of drawbacks and do not treat extremes (e.g., with heavy-tailed distribution) in any special way. This paper explores a different analysis tool, the archetypal analysis (AA), to identify “pure” types or archetypes of the data where these archetypes are expressed as a mixture (or convex combination) of the observations, and the observations are an approximate mixture of the archetypes. This, plus the fact that the archetypes constitute a subset of the data corners (or outskirts of the observations in the data sample space),

FIG. 14. As in Fig. 13, but for the time series of the archetypes.

constitute the main features of AA, which combines the virtues of clustering and EOF analysis. AA is relatively “new” in climate science, although it was introduced about two decades ago in pattern recognition (Cutler and Breiman 1994).

AA is essentially based on estimating the convex hull of the data and attempts to find representative corners or extremes of the data. In addition, AA allows a representation of high-dimensional data using simplex visualization and provides transition probabilities between states, which may be useful (e.g., in clustering). Conventionally, the AA problem is solved using the original multistep alternating constrained least squares (Cutler and Breiman 1994; Seth and Eugster 2016). Here we have developed a new approach to solve the AA problem using a Riemannian manifold-based algorithm. The method is based on a simultaneous optimization of both the weight matrices **A** and **B** using the oblique manifold by calculating the gradient of the cost function on this manifold and optimizing via the conjugate gradient or the steepest descent using the Manopt toolbox. The method has been illustrated with a simple example and then applied to real data of SSTs and the Asian summer monsoon.

The method was first applied to the monthly SST anomalies (with respect to the mean annual cycle) from the HadISST between 45.5°S and 45.5°N. Both the nondetrended and the detrended anomalies were used. For the nondetrended anomalies the analysis suggests three archetypes. The first two archetypes are associated, respectively, with El Niño and La Niña states with anomalies ranging from -2° to 2.5°C . These two states are not opposite of each other as would be expected from an EOF analysis, reflecting the skewness of the tropical Pacific SSTs. The third archetype came out as the western boundary currents including the Kuroshio and Gulf Stream but also other features of these currents such as the East Australian Current, Gulf of Angola, and Southern California Countercurrent. The SST anomalies of these currents go up to about 2°C . The contributions to this pattern come essentially from the last two decades. The (amplitude) time series of the western currents archetype shows an increasing trend starting from around the mid-1930s. There is also a decreasing trend of La Niña archetype time series, but the El Niño archetype time series does not show a pronounced trend. When the SST anomalies are detrended the analysis

FIG. 15. Solutions of the four-archetype problem applied to the detrended SLP anomalies over the Asian summer monsoon region. Contour interval is 1 hPa.

suggests that the data can be explained well with two archetypes, El Niño and La Niña.

AA was also applied to the Asian summer monsoon variability. The data consisted of the detrended SLP anomalies from ERA-40 over the Asian monsoon region (20° – 35° N, 50° – 150° E) for the summer season June–September (JJAS) 1958–2001 as in the study of HT13, which is used for comparison. The scree plot of the relative RSS suggests two elbows, one at $p = 3$ and a weaker one at $p = 4$ archetypes. We have investigated the different solutions corresponding to $p = 3, 4$, and 5 for consistency. The three-archetype solution yielded similar patterns to the regime centroids of HT13, namely the WNP active phase and the active and break

monsoon phases. The magnitudes of the archetypes are, however, larger than those of the regimes by up to 6 times. The simplex projection also suggests that the archetypes are somewhat associated with the regime centroids of HT13. Contributions to the archetypes come from various observations across the observed period, particularly the 1980s for the WNP active phase and from the 1960s and 1990s for the break phase. There is some signature of low-frequency variation in the WNP active phase and the Indian active phase (on the order of 4–5 years), due possibly to oceanic influence but with much less signature for the (Indian) break phase. The four- and five-archetype solutions have also been examined. Despite the nonnested nature of AA, the

FIG. 16. As in Fig. 12, but for simplex projection of the (a) four- and (b) five-archetype solution of the daily detrended SLP anomalies over the Asian monsoon region.

patterns provided by the three-archetype model are approximately reproduced in the four- and five-archetype models. This was also confirmed from the simplex projection, suggesting that the Asian monsoon dynamics are consistent with the nonlinear regimes of HT13.

The robustness of the method was also analyzed by investigating the sensitivity to the sample size and outliers using both the generated samples and the real data. The method shows robustness to the sample size and also extremes. The method is slightly sensitive to outliers, but the archetypes of the data with outliers are not dissociated with those of the data without outliers. The investigation of data without any structure (multinormal) shows that for small dimensions the optimum number of archetypes is twice the dimension, but for large dimensions no such optimum is suggested.

AA constitutes a powerful tool that can be used to analyze large-scale climate data. One of the outstanding features of AA that is lacking from other conventional methods, such as EOF analysis, is the direct link to the observations, which facilitates the interpretation. The approximation using the “corners” of the data is unique to AA. This can be used, for example, to study the temporal expansion of the convex hull of the system within its state space resulting, for instance, from global warming. This can be applied, for example, to surface temperature in the polar region in association with the polar amplification.

The algorithm developed here scales with the temporal dimension (or sample size) and the problem dimension, an advantage that can be explored to analyze, for example, simulations from high-resolution climate models. Perhaps one of the challenging issues of AA is the determination of the “right” number of archetypes. The scree plot was shown to be quite useful here but may not be so in other cases, and methods based, for example, on information criteria such as Akaike information criteria (AIC; Akaike 1973, 1974) or Bayesian information procedure (BIC; Akaike 1978; Schwarz 1978) or even a cross-validation approach (e.g., Ramsay and Silverman 2006) could be explored in the future. Another point that was not investigated further is the probabilistic nature of the weight matrices. These could be used to analyze the temporal transition probabilities between archetypes and could be explored, for example, in combination with forecasting models.

Acknowledgments. The authors thank two anonymous reviewers for their constructive comments and Leon Chafik for the discussion we had.

REFERENCES

- Absil, P.-A., R. Mahony, and R. Sepulchre, 2010: Optimization on manifolds: Methods and applications. *Recent Advances in Optimization and Its Applications in Engineering*, M. Diehl, F. Glineur, and E. J. W. Michiels, Eds., Springer, 125–144.
- Akaike, H., 1973: Information theory and an extension of the maximum likelihood principle. *Proc. Second Int. Symp. on Information Theory*, Budapest, Hungary, Akademiai Kiado, 267–281.
- , 1974: A new look at the statistical model identification. *IEEE Trans. Autom. Control*, **19**, 716–723, doi:10.1109/TAC.1974.1100705.
- , 1978: Time series analysis and control through parametric models. *Applied Time Series Analysis*, D. F. Findley, Ed., Academic Press, 1–23.
- Ambaum, M. H. P., B. J. Hoskins, and D. B. Stephenson, 2001: Arctic Oscillation or North Atlantic Oscillation. *J. Climate*, **14**, 3495–3507, doi:10.1175/1520-0442(2001)014<3495:AOONAO>2.0.CO;2.
- Barber, C. B., D. P. Dobkin, and H. Huhdanpaa, 1996: The quickhull algorithm for convex hulls. *ACM Trans. Math. Software*, **22**, 469–483, doi:10.1145/235815.235821.
- Bauckhage, C., and C. Thureau, 2009: Making archetypal analysis practical. *Pattern Recognition*, J. Denzler, G. Notni, and H. Süße, Eds., Lecture Notes in Computer Science, Vol. 5748, Springer, 272–281.
- Boumal, N., B. Mishra, P.-A. Absil, and R. Sepulchre, 2014: Manopt: A toolbox for optimization on manifolds. *J. Mach. Learn. Res.*, **15**, 1455–1459.
- Chan, B. H. P., D. A. Mitchell, and L. E. Cram, 2003: Archetypal analysis of galaxy spectra. *Mon. Not. Roy. Astron. Soc.*, **36**, 338–347, doi:10.1046/j.1365-8711.2003.06099.x.
- Christiansen, B., 2003: Evidence for nonlinear climate change: Two stratospheric regimes and a regime shift. *J. Climate*, **16**, 3681–3690, doi:10.1175/1520-0442(2003)016<3681:EFNCT>2.0.CO;2.
- , 2007: Atmospheric circulation regimes: Can cluster analysis provide the number? *J. Climate*, **20**, 2229–2250, doi:10.1175/JCLI4107.1.
- Coexeter, H. S. M., 1973: *Regular Polytopes*. Dover Publications, 339 pp.
- Cutler, A., and L. Breiman, 1994: Archetypal analysis. *Technometrics*, **36**, 338–347, doi:10.1080/00401706.1994.10485840.
- Dommenget, D., and M. Latif, 2003: Reply. *J. Climate*, **16**, 1094–1097, doi:10.1175/1520-0442(2003)016<1094:R>2.0.CO;2.
- Hannachi, A., 2007: Pattern hunting in climate: A new method for finding trends in gridded climate data. *Int. J. Climatol.*, **27**, 1–15, doi:10.1002/joc.1375.
- , 2010: On the origin of planetary-scale extratropical winter circulation regimes. *J. Atmos. Sci.*, **67**, 1382–1401, doi:10.1175/2009JAS3296.1.
- , 2016: Regularised empirical orthogonal functions. *Tellus*, **68A**, 31723, doi:10.3402/tellusa.v68.31723.
- , and A. G. Turner, 2013: 20th century intraseasonal Asian summer monsoon dynamics viewed from isomap. *Nonlinear Processes Geophys.*, **20**, 725–741, doi:10.5194/npg-20-725-2013.
- , I. T. Jolliffe, D. B. Stephenson, and N. Trenafilov, 2006: In search of simple structures in climate: Simplifying EOFs. *Int. J. Climatol.*, **26**, 7–28, doi:10.1002/joc.1243.
- , —, and —, 2007: Empirical orthogonal functions and related techniques in atmospheric science: A review. *Int. J. Climatol.*, **27**, 1119–1152, doi:10.1002/joc.1499.

- , D. M. Straus, C. L. E. Franzke, S. Corti, and T. Woollings, 2017: Low-frequency nonlinearity and regime behavior in the Northern Hemisphere extratropical atmosphere. *Rev. Geophys.*, **55**, 199–234, doi:[10.1002/2015RG000509](https://doi.org/10.1002/2015RG000509).
- Jolliffe, I. T., 2002: *Principal Component Analysis*. 2nd ed. Springer, 487 pp.
- Lorenz, E. N., 1956: Empirical orthogonal functions and statistical weather prediction. MIT Dept. of Meteorology Tech. Rep., 49 pp.
- Mestas-Núñez, A. M., 2000: Orthogonality properties of rotated empirical modes. *Int. J. Climatol.*, **20**, 1509–1516, doi:[10.1002/1097-0088\(200010\)20:12<1509::AID-JOC553>3.0.CO;2-Q](https://doi.org/10.1002/1097-0088(200010)20:12<1509::AID-JOC553>3.0.CO;2-Q).
- Monahan, A. H., J. C. Fyfe, M. H. P. Ambaum, D. B. Stephenson, and G. R. North, 2009: Empirical orthogonal functions: The medium is the message. *J. Climate*, **22**, 6501–6514, doi:[10.1175/2009JCLI3062.1](https://doi.org/10.1175/2009JCLI3062.1).
- Morup, M., and L. K. Hansen, 2012: Archetypal analysis for machine learning and data mining. *Neurocomputing*, **80**, 54–63, doi:[10.1016/j.neucom.2011.06.033](https://doi.org/10.1016/j.neucom.2011.06.033).
- Munk, W. H., 1950: On the wind-driven ocean circulation. *J. Meteor.*, **7**, 80–93, doi:[10.1175/1520-0469\(1950\)007<0080:OTWDOC>2.0.CO;2](https://doi.org/10.1175/1520-0469(1950)007<0080:OTWDOC>2.0.CO;2).
- Obukov, A. M., 1947: Statistically homogeneous fields on a sphere. *Usp. Mat. Nauk*, **2**, 196–198.
- Porzio, G. C., G. Ragozini, and D. Vistocco, 2008: On the use of archetypes as benchmarks. *Appl. Stochastic Models Data Anal.*, **24**, 419–437, doi:[10.1002/asmb.727](https://doi.org/10.1002/asmb.727).
- Ramsay, J. O., and B. W. Silverman, 2006: *Functional Data Analysis*. 2nd ed. Springer, 426 pp.
- Rayner, N. A., D. E. Parker, E. B. Horton, C. K. Folland, L. V. Alexander, D. P. Rowell, E. C. Kent, and A. Kaplan, 2003: Global analyses of sea surface temperature, sea ice, and night marine air temperature since the late nineteenth century. *J. Geophys. Res.*, **108**, 4407, doi:[10.1029/2002JD002670](https://doi.org/10.1029/2002JD002670).
- Robertson, A. W., and M. Ghil, 1999: Large-scale weather regimes and local climate over the western United States. *J. Climate*, **12**, 1796–1813, doi:[10.1175/1520-0442\(1999\)012<1796:LSWRAL>2.0.CO;2](https://doi.org/10.1175/1520-0442(1999)012<1796:LSWRAL>2.0.CO;2).
- Schwarz, G., 1978: Estimating the dimension of a model. *Ann. Stat.*, **6**, 461–464, doi:[10.1214/aos/1176344136](https://doi.org/10.1214/aos/1176344136).
- Seager, R., and I. R. Simpson, 2016: Western boundary currents and climate change. *J. Geophys. Res. Oceans*, **121**, 7212–7214, doi:[10.1002/2016JC012156](https://doi.org/10.1002/2016JC012156).
- Seth, S., and M. Eugster, 2016: Probabilistic archetypal analysis. *Mach. Learn.*, **102**, 85–113, doi:[10.1007/s10994-015-5498-8](https://doi.org/10.1007/s10994-015-5498-8).
- Smith, S. T., 1994: Optimization techniques on Riemannian manifolds. *Hamiltonian and Gradient Flows, Algorithm and Control*, A. Bloch, Ed., Field Institute Communications, Vol. 3, American Mathematical Society, 113–136.
- Steinschneider, S., and U. Lall, 2015: Daily precipitation and tropical moisture exports across the eastern United States: An application of archetypal analysis to identify spatiotemporal structure. *J. Climate*, **28**, 8585–8602, doi:[10.1175/JCLI-D-15-0340.1](https://doi.org/10.1175/JCLI-D-15-0340.1).
- Stommel, H., 1948: The westward intensification of wind-driven ocean currents. *Eos, Trans. Amer. Geophys. Union*, **29**, 202–206, doi:[10.1029/TR029i002p00202](https://doi.org/10.1029/TR029i002p00202).
- Sura, P., 2011: A general perspective of extreme events in weather and climate. *Atmos. Res.*, **101**, 1–21, doi:[10.1016/j.atmosres.2011.01.012](https://doi.org/10.1016/j.atmosres.2011.01.012).
- , and A. Hannachi, 2015: Perspectives of non-Gaussianity in atmospheric synoptic and low-frequency variability. *J. Climate*, **28**, 5091–5114, doi:[10.1175/JCLI-D-14-00572.1](https://doi.org/10.1175/JCLI-D-14-00572.1).
- Wang, D., H. Qiao, B. Zhang, and M. Wang, 2013: Online support vector machine based on convex hull vertices selection. *IEEE Trans. Neural Networks Learn. Syst.*, **24**, 593–609, doi:[10.1109/TNNLS.2013.2238556](https://doi.org/10.1109/TNNLS.2013.2238556).
- Wilks, D., 2006: *Statistical Methods in the Atmospheric Sciences*. Academic Press, 676 pp.
- Yang, H., G. Lohmann, W. Wei, M. Dima, M. Ionita, and J. Liu, 2016: Intensification and poleward shift of subtropical western boundary currents in a warming climate. *J. Geophys. Res. Oceans*, **121**, 4928–4945, doi:[10.1002/2015JC011513](https://doi.org/10.1002/2015JC011513).