Published in final edited form as:

Biometrics. 2008 September ; 64(3): 931–939. doi:10.1111/j.1541-0420.2007.00959.x.

Order-Preserving Dimension Reduction Procedure for the Dominance of Two Mean Curves with Application to Tidal Volume Curves

Sang Han Lee^{1,*}, Johan Lim^{2,**}, Marina Vannucci^{3,***}, Eva Petkova^{4,****}, Maurice Preter^{5,*****}, and Donald F. Klein^{5,*****}

¹Nathan S. Kline Institute for Psychiatric Research, 140 Old Orangeburg Road, Orangeburg, New York 10962, U.S.A

²Department of Applied Statistics, Yonsei University, Seoul 120-749, Korea

³Department of Statistics, Rice University, Houston, Texas 77251-1892, U.S.A

⁴Child Study Center, School of Medicine, New York University, New York, U.S.A

⁵Department of Psychiatry, Columbia University, and New York State Psychiatric Institute, New York, U.S.A

Summary

The paper here presented was motivated by a case study involving high-dimensional and high-frequency tidal volume traces measured during induced panic attacks. The focus was to develop a procedure to determine the significance of whether a mean curve dominates another one. The key idea of the suggested method relies on preserving the order in mean while reducing the dimension of the data. The observed data matrix is projected onto a set of lower rank matrices with a positive constraint. A multivariate testing procedure is then applied in the lower dimension. We use simulated data to illustrate the statistical properties of the proposed testing procedure. Results on the case study confirm the preliminary hypothesis of the investigators and provide critical support to their overall goal of creating an experimental model of the clinical panic attack in normal subjects.

Keywords

Dimension reduction; Follmann's test; Matrix factorization; Panic disorder Stochastic order; Tidal volume curves

© 2008, The International Biometric Society

*hanul31@stat.tamu.edu

**johanlim@yonsei.ac.kr

***marina@rice.edu

****ep120@med.nyu.edu

*****preterm@pi.cpmc.columbia.edu

*****donaldK737@aol.com

1. Introduction

High-dimensional functional data have become prominent in a number of medical and biological fields. There the units of observation are curves and the observed data consist of sets of curves, often sampled on a fine grid. More and more attention among researchers is now devoted to the development of appropriate statistical methodologies suitable for the analysis of such data. The work we present here was motivated by an ongoing collaboration with investigators at the New York State Psychiatric Institute, at Columbia University. The overall goal is to create a model of the clinical panic attack in normal human subjects, as it occurs in individuals affected by panic disorder (Preter and Klein, 2007). Here we look at data arising from a randomized study where measurements of tidal volume, that is, the volume of gas exchanged during each ventilated breath, are taken on a number of individuals subject to interventions that may induce panic attacks. Prior to the study investigators had an ordered mean hypothesis of the type $f_1(t) > f_2(t)$ for all t (and $f_1(t) > f_2(t)$ for at least one t), with f_i the mean curve for group i , $i = 1, 2$, and with groups 1 and 2 defined by two different interventions. Our task was to design a test to statistically validate this hypothesis.

The key idea of the novel procedure we suggest relies on preserving the order in mean while reducing the dimension of the data. We do this by projecting the observed data matrix onto a space of low-rank matrices, which are represented as a product of a coefficient matrix and a positively constrained basis matrix that preserves the order between curves, that is, if one curve is larger than the other one, then the coefficient vectors will preserve the same ordering. We then apply a multivariate testing procedure to the coefficient vectors in the lower dimension. Here we employ the modified Hotelling's T-statistics proposed by Follmann (1996), which has typically good power. Other testing procedures could be readily applied, such as the approximate likelihood ratio test of Tang, Gnecco, and Geller (1989). Notice that these procedures require a consistent covariance matrix estimate and cannot therefore be directly applied to a high-dimensional $n \times p$ data matrix where p is possibly larger than n .

The order-preserving matrix factorization we adopt, denoted by nonnegative basis matrix factorization (NBMF), minimizes the L_2 norm between the observed data matrix and a prespecified lower rank matrix, which imposes positive constraints on the basis vectors. In this regard, our procedure is comparable to other dimension reduction procedures, such as principal component analysis (PCA) and the nonnegative matrix factorization of Lee and Seung (1999). Unlike our procedure, however, PCA does not impose any constraint, whereas the nonnegative matrix factorization assumes positiveness of both coefficients and basis vectors. In our approach, by avoiding the additional constraint, curves data are not limited to be positive. Our NBMF problem, however, is not convex and no algorithm can guarantee convergence to a global minimum. We propose an iterative procedure that converges to a local minimum.

There are several alternative approaches that one could use to solve the testing problem at hand. A naive one is to compute the average value of each curve and then apply a one-side t -test. This global test completely ignores the pointwise nature of the data. Another approach

is a pointwise t -test. With respect to the simple global t -test, pointwise t -tests give the additional information on where the significance occurs, see Ramsay and Silverman (1997). This procedure, however, is sub-optimal and leads to large type I errors due to multiple testing. One can expect that an overall test that combines all the pointwise comparisons, such as a nonparametric procedure, would perform better. Fan (1996) and other researchers, see Serban and Wasserman (2005) and references therein, suggested nonparametric methods for curve testing problems based on representations of the curves that use basis functions, such as wavelets or Fourier bases. These methods, however, do not consider any order constraints and are therefore not applicable to our problem. Positive basis functions, such as cubic B-splines, would preserve the order constraint but are sub-optimal for dimension reduction, often resulting in many nonzero coefficients. Indeed, with B-splines the number of nonzero coefficients is proportional to the number of basis functions used and additional variable selection procedures are needed to achieve dimension reduction, see Zhang et al. (2004). An additional advantage of the method we propose is that it determines the most appropriate basis functions for any given dimension, by minimizing an L_2 approximation error.

The remainder of the article is organized as follows. In Section 2 we describe the case study that motivated our work. In Section 3 we introduce the NBMF method and the testing procedure to test the order between two mean curves. We also provide an algorithm for its implementation. In Section 4 we illustrate performances on simulated data and in Section 5 we present results from the case study example. Section 6 concludes the article.

2. Sodium-Lactate-Induced Panic Attacks

The testing procedure we have developed in this article was motivated by an ongoing collaboration with investigators at the New York State Psychiatric Institute, at Columbia University. The overall goal is to create a model of the clinical panic attack in normal human subjects, as it occurs in individuals affected by panic disorder. Here we analyze data from an experiment that looks at high-dimensional, high-frequency measurements of tidal volume on a number of individuals subject to interventions that may induce panic attacks. Prior to the study investigators had an ordered mean hypothesis of the type $f_1 > f_2$ with groups 1 and 2 defined by two different interventions.

2.1 Experimental Study

Sodium lactate reliably produces panic attacks in patients with panic disorder (Liebowitz et al., 1985). Normals rarely have such reactivity. A distinctive feature of sodium-lactate-induced panic is a marked increase in tidal volume (Goetz et al., 1993). Klein (1993) suggested that the spontaneous panic attack may be due to a hypersensitive alarm system for the detection of signals of impending suffocation, such as rising levels of CO_2 or brain lactate. The endogenous opioid system is an important central regulator of respiratory drive. An exogenous opioid, such as morphine, blunts sensitivity to CO_2 (Fleetham et al., 1980). Conversely, naloxone, an opioid receptor antagonist, increases the ventilatory response to hypercapnic hypoxia in normal human controls (Akiyama et al., 1993). Naloxone pretreatment may make normal individuals (who putatively have an intact opioid system) vulnerable to the marked angiogenic and respiratory effects of sodium lactate. In a pilot

study Sinha, Goetz, and Klein (2007) found that lactate after naloxone, administered to normals, produced a marked increase in tidal volume that exceeded previous results from infusing only lactate. Surprisingly, lactate, despite producing a metabolic alkalosis, is a tidal volume stimulant, as has been shown in both normal humans and rats.

A randomized study with normal subjects was designed to test the investigators' hypothesis. Healthy normal male and female adult volunteers, not affected by any psychiatric or significant medical illness, were subject to different interventions. Here we focus on the two groups that received either naloxone followed by sodium lactate or saline followed by sodium lactate. The measuring and data recording device was the lifeShirt (Wilhelm, Roth, and Sackner, 2003), a garment recently developed with embedded inductive plethysmography sensors for continuous ambulatory monitoring of respiration and other physiological functions. Subjects had sensors and intravenous lines placed while supine. The experiment on each subject consisted of four phases. During phase 1, baseline measurements were recorded for 30 minutes. In phase 2, one group of subjects received naloxone, for 3–5 minutes, whereas the other group received saline. In phase 3, all subjects received sodium lactate for approximately 20 minutes. Phase 4 was the recovery stage. In total, 65 subjects completed all phases of the experiment, 37 in the “N + L” (naloxone–lactate) group and 28 in the “S + L” (saline–lactate) group. The hypothesis was that subjects receiving the naloxone–lactate sequence will have greater increases in tidal volume during the lactate phase than subjects in the other group.

3. Methods

This article focuses on the problem of testing the order between two mean curves, f_1 and f_2 , based on sampled curves. To be specific, let

$$\mathbf{Y}^{(1)} = \begin{pmatrix} Y_1^{(1)} \\ Y_2^{(1)} \\ \vdots \\ Y_{n_1}^{(1)} \end{pmatrix} = \begin{pmatrix} y_{11}^{(1)}, & y_{12}^{(1)}, & \dots, & y_{1p}^{(1)} \\ y_{21}^{(1)}, & y_{22}^{(1)}, & \dots, & y_{2p}^{(1)} \\ \vdots & \vdots & \vdots & \vdots \\ y_{n_1 1}^{(1)}, & y_{n_1 2}^{(1)}, & \dots, & y_{n_1 p}^{(1)} \end{pmatrix}$$

represent the n_1 curves for group 1 measured at p points, and $\mathbf{Y}^{(2)}$ the similarly defined $n_2 \times p$ matrix of curves for group 2. We assume that curves $Y_j^{(i)}$, for $j = 1, 2, \dots, n_i$, are independent of each other. Let now $f_i = (f_{i1}, f_{i2}, \dots, f_{ip})$, for $i = 1, 2$, indicate the p -dimensional population mean curves for groups 1 and 2, respectively. We want to test $f_1 > f_2$, that is, $f_{1k} > f_{2k}$ for every $k = 1, \dots, p$ (and $f_{1k} > f_{2k}$ for at least one k).

3.1 Order-Preserving Dimension Reduction

With p possibly larger than n , the key idea behind our proposal is to represent $\mathbf{Y}^{(1)}$ and $\mathbf{Y}^{(2)}$ via lower dimensional vectors of coefficients that preserve the order between the two mean curves, so that the hypothesis $f_1 > f_2$ can be tested in a lower dimension. Specifically, let $\mathbf{W}^{(1)}$ (and $\mathbf{W}^{(2)}$) be the lower dimensional approximation of $\mathbf{Y}^{(1)}$ (and $\mathbf{Y}^{(2)}$) and μ_1 (and μ_2)

be its population mean vector. Thus, we test the order between f_1 and f_2 by testing the order between μ_1 and μ_2 .

Let $\|\mathbf{A}\|_F$ be the standard L_2 matrix norm, also known as Frobenius norm, of $\mathbf{A}$, that is,

$\|\mathbf{A}\|_F = (\sum_{ij} A_{ij}^2)^{1/2}$. We seek to do dimension reduction by finding a small number of local features of the curves, each defined as a positive linear combination of the p time points. More precisely, we find a low-rank approximation to the data matrix $\mathbf{Y}$ that consists of a $n \times r$ coefficient matrix $\mathbf{W}$ and an $r \times p$ nonnegative basis matrix $\mathbf{H}$, by solving the following nonlinear optimization problem

$$\text{minimize } \|\mathbf{Y} - \mathbf{WH}\|_F^2, \text{ subject to } \mathbf{H} \geq 0, \|H_k\|_F^2 = 1 \quad (1)$$

for $k = 1, 2, \dots, r$, where H_k is the k th row vector of $\mathbf{H}$ and where the dimension r of $\mathbf{W}$ is a parameter to be set by the user.

3.2 Testing the Order Between Two Mean Curves

Let us now consider the testing problem. We first apply a common basis matrix $\mathbf{H}$ to both data matrices, $\mathbf{Y}^{(1)}$ and $\mathbf{Y}^{(2)}$, and find lower-dimensional approximations. To be specific, we set

$$\mathbf{Y} = \begin{pmatrix} \mathbf{Y}^{(1)} \\ \mathbf{Y}^{(2)} \end{pmatrix} \quad (2)$$

and find r -rank approximations with positive basis vectors of the type

$$\mathbf{Y} = \begin{pmatrix} \mathbf{Y}^{(1)} \\ \mathbf{Y}^{(2)} \end{pmatrix} \approx \begin{pmatrix} \mathbf{W}^{(1)} \\ \mathbf{W}^{(2)} \end{pmatrix} \mathbf{H}. \quad (3)$$

We can now apply a multivariate statistical testing procedure to the lower-dimensional approximations of the data. Testing the order between the mean curves f_1 and f_2 is equivalent to testing the order between the mean vectors of their lower-dimensional approximations, μ_1 and μ_2 . From (3) it is clear that the following relations between f_1 and f_2 , and μ_1 and μ_2 hold: $f_1 \approx \mu_1 \mathbf{H}$ and $f_2 \approx \mu_2 \mathbf{H}$, with μ_1 and μ_2 , being the true coefficients of the finite basis approximation. This implies

$$f_1 \geq f_2 \iff \mu_1 \mathbf{H} - \mu_2 \mathbf{H} \geq 0 \iff (\mu_1 - \mu_2) \mathbf{H} \geq 0 \iff (\mu_1 - \mu_2) \mathbf{H} \mathbf{H}^T \geq 0 \iff \mu_1 - \mu_2 \geq 0$$

and we are left to test $H_0: \mu_1 = \mu_2$ against $H_1: \mu_1 > \mu_2$ componentwise (and $\mu_1 > \mu_2$ at least at one point), based on the lower-dimensional approximations $\mathbf{W}^{(1)}$ and $\mathbf{W}^{(2)}$.

Here we use the Follmann's multivariate procedure to test the order between μ_1 and μ_2 . Other procedures, such as the approximate likelihood ratio test by Tang et al. (1989), may be used here to test the same hypothesis. In order to test a one-sided alternative hypothesis, Follmann (1996) suggested the use of the following modified Hotellings' T-statistics:

$$T = (\bar{\mathbf{W}}^{(1)} - \bar{\mathbf{W}}^{(2)}) (n_1^{-1} \mathbf{S}_1 + n_2^{-1} \mathbf{S}_2)^{-1} (\bar{\mathbf{W}}^{(1)} - \bar{\mathbf{W}}^{(2)})^T \times I \left(\sum_{k=1}^r (\bar{W}_k^{(1)} - \bar{W}_k^{(2)}) > 0 \right), \quad (4)$$

where $\mathbf{S}_i = \frac{1}{n_i - 1} \sum_{j=1}^{n_i} (\mathbf{W}_j^{(i)} - \bar{\mathbf{W}}^{(i)}) (\mathbf{W}_j^{(i)} - \bar{\mathbf{W}}^{(i)})^T$ and $\bar{\mathbf{W}}^{(i)} = \frac{1}{n_i} \sum_{j=1}^{n_i} \mathbf{W}_j^{(i)}$, for $i = 1, 2$, with $\bar{W}_k^{(i)}$ being the k th component of $\bar{\mathbf{W}}^{(i)}$ and $I(\cdot)$ the indicator function. Follmann's test has a good power for the alternative hypothesis of a positive mean vector of differences. The test rejects if a quadratic form of the sample mean vector exceeds its 2α critical value and the sum of the elements of the mean vector exceeds zero. This test is shown to have type I error rate equal to α for both cases of known and unknown covariance matrix. It can also be shown that (4) converges in distribution to a half χ^2 distributed random variable with r degrees of freedom.

3.3 Probabilistic View

A probabilistic interpretation of our NBMF method can help in setting practical guidelines for the choice of r , the dimension of $\mathbf{H}$. We start by explaining the model we assume for the observation matrix $\mathbf{Y}$. Suppose that each row vector Y_j of $\mathbf{Y}$ has a mean vector f , which belongs to the space spanned by the column vectors of the $r^* \times p$ matrix $\mathbf{H}^*$. Suppose that the Y_j 's have covariance matrix Ω for $j = 1, 2, \dots, n$ and that they can be represented as

$$\mathbf{Y} = \begin{pmatrix} \mu \\ \vdots \\ \mu \end{pmatrix} \mathbf{H}^* + \mathbf{E}, \quad (5)$$

with μ a $1 \times r^*$ vector and where each column vector of $\mathbf{E}$ has mean vector 0 and covariance matrix Ω .

Suppose we mistakenly choose a dimension r , which is higher than the true dimension r^* . Let $\mathbf{H}$ be the basis matrix constructed by adding to $\mathbf{H}^*$ the extra basis vectors $H_{r^*+1}, \dots, H_r$, and let $\mathbf{H}^\perp$ be the matrix of basis vectors $H_{r+1}, \dots, H_p$, orthogonal to those of $\mathbf{H}$. We rewrite (5) as $\mathbf{Y} = \eta \mathbf{H} + \mathbf{E}_1 \mathbf{H} + \mathbf{E}_2 \mathbf{H}^\perp$, where

$$\eta = \begin{pmatrix} \mu & 0 \\ \mu & 0 \\ \vdots & \vdots \\ \mu & 0 \end{pmatrix}.$$

The optimization problem (1) becomes

$$\|\mathbf{Y} - \mathbf{W}\mathbf{H}\|_F^2 = \|(\eta + \mathbf{E}_1 - \mathbf{W})\mathbf{H} + \mathbf{E}_2 \mathbf{H}^\perp\|_F^2 = \|(\eta + \mathbf{E}_1 - \mathbf{W})\mathbf{H}\|_F^2 + \|\mathbf{E}_2 \mathbf{H}^\perp\|_F^2$$

because of the orthogonality between $\mathbf{H}$ and $\mathbf{H}^\perp$.

Thus, the solution to the problem, given $\mathbf{H}$, is $\mathbf{W} = \boldsymbol{\eta} + \mathbf{E}_1$ and we have

$$\min_{\mathbf{W}} \|\mathbf{Y} - \mathbf{WH}\|_{\text{F}}^2 = \|\mathbf{E}_2 \mathbf{H}^\perp\|_{\text{F}}^2 \approx n \|\mathbf{H}^{\perp T} \Omega \mathbf{H}^\perp\| = n \sum_{k=r+1}^p \|H_k^T \Omega H_k\|, \quad (6)$$

where $\|A\| = \sum_{ij} |A_{ij}|$. This follows from the fact that the variance of a column vector of $\mathbf{E}_2 \mathbf{H}^\perp$ is $\mathbf{H}^{\perp T} \Omega \mathbf{H}^\perp$, because of the orthogonality between $\mathbf{H}$ and $\mathbf{H}^\perp$ and the relationship $\mathbf{E} = \mathbf{E}_1 \mathbf{H} + \mathbf{E}_2 \mathbf{H}^\perp$, and from the orthogonality among $H_{r+1}, \dots, H_p$. Note that the Frobenius norm of $\|\mathbf{E}_2 \mathbf{H}^\perp\|_{\text{F}}^2$ is defined as the sum of squares of each component in the matrix. Hence, it is reasonable to expect it to be close to $n \|\mathbf{H}^{\perp T} \Omega \mathbf{H}^\perp\|$.

A difficulty arises from the fact that Ω is unknown. Here we assume Ω to be isotropic, that is, invariant to any rotation of the orthonormal axes. This assumption is reasonable for our case study, where individual curves are observed at a set of equally spaced points and where the error process is assumed to be stationary and identically distributed across samples. With this assumption, the following statistics, which we denote by mean square error (MSE),

$$\min_{\mathbf{W}} \|\mathbf{Y} - \mathbf{WH}\|_{\text{F}}^2 / \{(p-r)n\} = \|\mathbf{E}_2 \mathbf{H}^\perp\|_{\text{F}}^2 / \{(p-r)n\} \approx \sum_{k=r+1}^p \|H_k^T \Omega H_k\| / (p-r) \quad (7)$$

is expected to be constant for every $r \geq r^*$. In the special case $\Omega = \sigma^2 \mathbf{I}$, we can expect the MSE to be approximately equal to σ^2 for $r \geq r^*$. Thus, if the MSE does not markedly decrease after a certain r , we can choose r as the dimension of the reduced space. Later in the examples we will demonstrate that this approach provides a good reference to select an appropriate r .

3.4 Iterative Algorithm

Finally, we describe our optimization algorithm to solve equation (1). The function to be minimized in (1) is convex either in $\mathbf{W}$ or $\mathbf{H}$, but not in both. Using this fact we propose the following iterative procedure to solve the least square (LS) problem and find a local minimum:

- i. Given the current estimate of $\mathbf{W}$, solve the constrained LS problem:

$$\{\text{minimize } \|\mathbf{Y} - \mathbf{WH}\|_{\text{F}}^2, \text{ subject to } \mathbf{H} \geq 0, \|H_k\|_{\text{F}}^2 = 1, k=1, 2, \dots, r\}.$$
- ii. Given the current estimate of $\mathbf{H}$, update $\mathbf{W}$ to $\mathbf{W} = \mathbf{YH}^T (\mathbf{H} \mathbf{H}^T)^{-1}$, as the solution to the unconstrained LS problem: $\{\text{minimize } \|\mathbf{Y} - \mathbf{WH}\|_{\text{F}}^2\}$.

Step (i) requires us to solve a quadratic program (QP) with linear inequality constraints and quadratic equality constraints. We address this step with a two-stage procedure: First, we solve the QP with linear inequalities,

$$\text{minimize } \|\mathbf{Y} - \mathbf{WH}\|_{\text{F}}^2, \text{ subject to } \mathbf{H} \geq 0, \quad (8)$$

and then we normalize the resulting estimates $\mathbf{H}$ as $H_k = H_k / \|H_k\|_F$, for $k = 1, 2, \dots, r$. Simple algebra can show the equivalence between the two optimization procedures, that is step (i) and the two-stage procedure we use.

The Karush–Kuhn–Tucker (KKT) conditions of a QP with linear equality constraints is a set of linear equations that can be solved analytically, see Boyd and Vandenberghe (2004). Two most common ways to solve a QP, or its KKT conditions, are the interior point method and the simplex method. The interior point method solves the QP with linear inequality constraints by reducing it to a sequence of linear equality constrained problems. The simplex method solves the KKT conditions by reformulating the problem into a linear programming problem. In this article we use the MOSEK optimization toolbox, available for free on the internet at www.mosek.com.

The proposed iterative LS procedure converges to a local minimum because each step finds a new estimate that improves the Frobenius norm. Let $(\mathbf{W}^{(m)}, \mathbf{H}^{(m)})$ indicate the current state at iteration m and let $(\mathbf{W}^{(m+1)}, \mathbf{H}^{(m+1)})$ be the subsequent estimates from steps (i) and (ii). Then,

$$\begin{aligned} \|\mathbf{Y} - \mathbf{W}^{(m)}\mathbf{H}^{(m)}\|_F^2 &\geq \min_{\mathbf{H}} \|\mathbf{Y} - \mathbf{W}^{(m)}\mathbf{H}\|_F^2 \\ &= \|\mathbf{Y} - \mathbf{W}^{(m)}\mathbf{H}^{(m+1)}\|_F^2 \geq \min_{\mathbf{W}} \|\mathbf{Y} - \mathbf{W}\mathbf{H}^{(m+1)}\|_F^2 \\ &= \|\mathbf{Y} - \mathbf{W}^{(m+1)}\mathbf{H}^{(m+1)}\|_F^2. \end{aligned}$$

The final estimates of $\mathbf{H}$ and $\mathbf{W}$ are obtained as the solutions of steps (i) and (ii) at the last iteration. In the data examples we will indicate these estimates as $\hat{\mathbf{H}}$ and $\hat{\mathbf{W}}$, respectively.

4. Numerical Examples

We used simulated data to investigate the statistical properties of our order-preserving dimension reduction testing procedure.

4.1 Selection of r

We start with a simple example that illustrates how the MSE in (7) decreases for $r = r^*$ and is constant for every $r = r^*$. Here, r^* is the true dimension of the lower-dimensional representation of the data, whereas r is the dimension prespecified by the user. This simple example supports our choice of the MSE as a good reference for selecting an appropriate r .

The numerical study was set up as follows: We fixed the true column rank r^* of $\mathbf{H}$ to 5, the dimension p of each curve to 10 (and 100), and the number n of subjects to 20. We randomly generated a coefficient matrix $\mathbf{W}$ of dimensions 20×5 from uniform distributions in the interval $[-10, 10]$ and a basis matrix $\mathbf{H}$ of dimensions 5×10 (and 5×100) from uniform distributions in $[0, 10]$. We also randomly generated an error matrix $\mathbf{E}$ using independent normal distributions with mean 0 and variance 1 and constructed an observation matrix as $\mathbf{Y} = \mathbf{W}\mathbf{H} + \mathbf{E}$. With this setting we generated 100 observational matrices. We solved the NBMF problem with each of the generated observational matrices and denoted

the solution for the l th matrix as $\hat{\mathbf{W}}_r^{(l)}$ and $\hat{\mathbf{H}}_r^{(l)}$, for $r = 1, \dots, 9$ and $l = 1, \dots, 100$. Figure 1 shows boxplots of the MSEs computed for each $r = 1, 2, \dots, 9$. As expected, the MSE values decrease as r increases toward r^* whereas they stay constant for $r > r^*$.

4.2 Power and Size of the NBMF Method

Next we report results on the power and test size of our NBMF method. We chose true mean functions that would resemble the case study data, that is,

$$f_1(t_k) = 224 + 0.26t_k + u_1(t_k), \quad f_2(t_k) = 200 + 0.3t_k - 0.0001t_k^2 + u_2(t_k), \quad (9)$$

where $u_1(t) = 20 \sqrt{t/2^{10}} \sin(50\pi/(t/2^{10} + 0.05))$, $u_2(t) = 10 \sin(\pi t/2^4)$ for $k = 1, 2, \dots, p$. A Doppler function and a sine function were added to f_1 and f_2 , respectively. We then set $r = 2$ and found the lower-dimensional approximations by NBMF, that is, the $1 \times r$ vectors μ_1, μ_2 , and the lower-dimension $r \times p$ basis matrix $\mathbf{H}$. We generated the sample coefficients W_{ij} in the lower-dimension space from a multivariate normal distribution with mean μ_i and covariance $\sigma_i^{*2} \mathbf{I}_{r \times r}$, $j = 1, 2, \dots, n_i$ for each $i = 1, 2$. Data were obtained by adding noise, that is, $\mathbf{Y} = \mathbf{WH} + \mathbf{E}$, with $\mathbf{W} = (W_{ij})$, $n \times r$, and $n = n_1 + n_2$, and where the noise term $\mathbf{E}$, $n \times p$, was generated from a multivariate normal with mean 0 and covariance $\sigma^2 \mathbf{I}_{n \times n}$.

We report on the performance of the test procedure for different values of σ^{*2} and σ^2 , which imply different signal-to-noise ratios (SNR; the ratio of the two standard deviations). We also looked into varying the distance between the two mean functions that we measured as

$\Delta^* = \frac{1}{T} \int (f_1(t) - f_2(t))^2 dt$, that is, as a standardized L_2 norm. We chose two values of σ^{*2} , that is, $\sigma^{*2} = 10^2, 10^4$. We chose σ^2 to obtain $SNR = 1, 1/3$. We generated 25 sample curves from each of the two groups, each curve with $p = 1024$ time points. We varied Δ^* by setting $f_1(t) = f_2(t)$ for $t \in [1, \dots, t_k^*]$ and by choosing various t_k^* . We also looked at differences in the opposite direction, that is, $f_1(t) = f_2(t)$ for $t \in [t_k^*, \dots, 1024]$ and various t_k^* .

The analytic form of the power of the NBMF test is not available. We therefore calculated power and size empirically by generating random sample curves under the null or alternative hypotheses, for fixed Δ^* and SNR. For computational convenience we fixed the lower-dimensional approximation at $r = 2$. This procedure was repeated m times and the proportion of rejections was computed. We used $m = 1000$. Under the null hypothesis the test size with $\sigma^{*2} = 10^2$ was 0.0510 for $SNR = 1$ and 0.0630 for $SNR = 1/3$, whereas with $\sigma^{*2} = 10^4$ it was 0.021 for $SNR = 1$ and 0.0580 for $SNR = 1/3$. As expected, results slightly worsen when σ^{*2} or σ^2 increases. Table 1 shows the power for various t_k^* , SNR, and σ^{*2} . The test shows good power.

5. Case Study

Let us now describe the application of the proposed testing procedure to the case study on sodium-lactate-induced panic attacks.

5.1 Preprocessing of Tidal Volume Traces

We considered data spanning over a time window covering the sodium lactate infusion. Based on their previous experience with sodium lactate infusions investigators did not expect a quick onset of effect. We therefore chose a window of approximately 17 minutes before the end of the infusion. Each subject has a different tidal volume (V_t) baseline. We performed baseline adjustment by calculating the median V_t of the baseline measurements and subtracting it from the V_t trace of each subject while under sodium lactate infusion. Data are massive. During the experiment tidal volume measurements were automatically saved 50 times per second. We thinned the data by considering traces obtained taking one every k th data points. We examined plots of several reduced traces to make sure we were preserving important features of the data and decided on $k = 25$ as a safe choice. This gave us two measurements per second.

A smoothing procedure was necessary in order to reveal the breathing patterns of interest to the investigators. The method we investigated uses wavelet decompositions to filter out high-to-medium frequency components of the data that are unrelated to the breathing frequencies, that is they constitute irrelevant information. The details of this preprocessing step are described elsewhere (Preter et al., 2007). Here we apply our proposed NBMF procedure to two different preprocessed sampled curves, that is, a smoothed version of the data and the extracted trends. Figure 2 shows, on the first row, two sample V_t curves after baseline adjustment, an $N + L$ sample on the left panel, and an $S + L$ on the right panel. The second row shows the same two curves after the wavelet smoothing procedure and the third row shows the corresponding trends. Figure 3 displays the mean curves for the $N + L$ (line) and $S + L$ (dotted line) groups, smoothed traces in the left panel, and trends in the right panel. Working with trends has the advantage of essentially avoiding complications with registration issues. Curve registration, see, for example, Ramsay and Silverman (1997), is a process according to which curves are “calibrated” across time, that is, aligned with respect to some common feature. Registration procedures for respiratory flows and tidal volumes, however, are not trivial. In our study subjects have very different breathing cycles and basically do not exhibit any global common feature.

5.2 Results

The investigators’ claim can be formalized as a hypothesis testing problem with alternative hypothesis of the type $(N + L)(t) > (S + L)(t)$ for all t (and $(N + L)(t) > (S + L)(t)$ for at least one t). We had available a total of 65 curves with 2048 observed points, 37 curves belong to the $N + L$ group, and 28 to the $S + L$ group. We tested the hypothesis on the smoothed curves and also on the trends. A very small r , that is, a large reduction, can be used for the trends. Using the MSE criterion we selected $r = 4$ for the trends and $r = 25$ for the smoothed data. Figure 4 shows original and approximated curves by $\hat{W}\hat{H}$ for six subjects.

In order to start our iterative procedure, $\mathbf{W}$ was randomly generated from the uniform distribution in the interval $[-100, 100]$. Convergence was achieved in only 10 iterations. We first applied the test to the entire time lag of the data (17 minutes). The test was significant (see Table 2), confirming the intuition of the investigators that subjects receiving the

naloxone–lactate sequence have greater increases in tidal volume. Because investigators were also interested in an indication of the time at which the significance occurs, we applied the testing procedure to shorter time lags of the data. We computed the p -values both from the half chi-square distribution of the Follmann's test statistics for large samples and from a permutation test. Results are summarized in Table 2 and indicate that the dominance of the $N + L$ mean curve over the $S + L$ one becomes more and more significant starting at approximately 10 minutes before the end of the sodium lactate infusion.

These results confirm the hypothesis of the investigators according to which naloxone pretreatment causes normal individuals, who presumably have an intact opioid system and are invulnerable to lactate panicogenesis, to manifest a lactate-provoked tidal volume increment similar to the clinical panic attack. This is an important result that confirms the preliminary findings of the investigators in support of the fact that naloxone–lactate interaction may provide an experimental model of the clinical panic attack in normal subjects. This finding will need to be confirmed by an experiment double-blindly blocking the $N + L$ effect by antipanic drugs but not by panic irrelevant drugs. If this antipanic agent-specific blockage were found positive, it would afford two useful advances. Currently, there is no specific screening method for testing putative antipanic drugs except by experimental treatment of panic disorder patients. Probably of more ultimate importance: supporting the hypothesis that an opioidergic dysfunction may be the pathophysiological mechanism underlying panic disorder allows new theoretical and practical approaches. If opioidergic dysfunction underlies panic pathophysiology, then the appropriateness of a new class of therapeutic agents comes into question (Preter and Klein, 2007).

6. Conclusion

We have proposed a dimension reduction procedure to test the significance of whether a mean curve dominates another one. The key idea of the suggested method relies on preserving the order in mean, while reducing the dimension of the data, and then applying a multivariate testing procedure to the reduced data. In addition, we have proposed an iterative algorithm to solve the projection problem. Our work was motivated by a study that looks at high-dimensional, high-frequency measurements of tidal volume on a number of individuals subject to interventions that may induce panic attacks. Our results have confirmed the hypothesis of the investigators according to which subjects receiving sodium lactate after naloxone have greater increases in tidal volume than subjects who do not receive the prior infusion of naloxone.

Acknowledgements

Vannucci is supported by NSF, DMS-0605001, and by NIH, R01 HG003319-01; Klein by NIMH, MH-30906, and MH-067749. Donald Klein is a member of the Scientific Advisory Board of Vivometrics, Inc.

References

- Akiyama Y, Nishimura M, Kobayashi S, Yoshioka A, Yamamoto M, Miyamoto K, Kawakami Y. Effects of naloxone on the sensation of dyspnea during acute respiratory stress in normal adults. *Journal of Applied Physiology*. 1993; 74:590–595. [PubMed: 8458774]

- Boyd, S.; Vandenberghe, L. *Convex Optimization*. Cambridge, U.K.: Cambridge University Press; 2004.
- Fan J. Tests of significance based on wavelet thresholding and Neyman's truncation. *Journal of the American Statistical Association*. 1996; 91:674–699.
- Fleetham J, Clarke H, Dhingra S, Chernick V, Anthiosen N. Endogenous opiates and chemical control of breathing in humans. *American Review of Respiratory Disease*. 1980; 121:1045–1049. [PubMed: 7416592]
- Follmann D. A simple multivariate test for one-sided alternatives. *Journal of the American Statistical Association*. 1996; 91:854–861.
- Goetz R, Klein DF, Gully D, Kahn J, Liebowitz M, Fyer A, Gorman J. Panic attacks during placebo procedures in the laboratory: Physiology and symptomatology. *Archives of General Psychiatry*. 1993; 50:280–285. [PubMed: 8466389]
- Klein DF. False suffocation alarms, spontaneous panics and related conditions: An integrative hypothesis. *Archives of General Psychiatry*. 1993; 50:306–317. [PubMed: 8466392]
- Lee DD, Seung HS. Learning the parts of objects by non-negative matrix factorization. *Nature*. 1999; 401:788–793. [PubMed: 10548103]
- Liebowitz M, Gorman J, Fyer A, Levitt M, Dillon D, Levy G, Appleby I, Anderson S, Palij M, Davies S, Klein D. Lactate provocation of panic attacks: Clinical and behavioral findings. *Archives of General Psychiatry*. 1985; 41:764–770. [PubMed: 6742978]
- Preter M, Klein DF. Panic, suffocation false alarms, separation anxiety and endogenous opioids. *Progress in Neuro-Psychopharmacology and Biological Psychiatry*. 2007 Aug. epub.
- Preter M, Lee SH, Vannucci M, Petkova E, Kim S, Klein DF. Ventilatory responses to a naloxone-antecedent-lactate challenge in normal subjects. Submitted to *Archives of Psychiatry*. 2007
- Ramsay, JO.; Silverman, BW. *Functional Data Analysis*. New York: Springer-Verlag; 1997.
- Serban N, Wasserman L. CATS: Clustering after transformation and smoothing. *Journal of the American Statistical Association*. 2005; 100:990–999.
- Sinha S, Goetz R, Klein D. Physiological and behavioral effects of naloxone and lactate in normals with relevance to the pathophysiology of panic disorder. *Psychiatry Research*. 2007; 149:309–314. [PubMed: 17156855]
- Tang D-I, Gnecco C, Geller NL. An approximate likelihood ratio test for a normal mean vector with nonnegative components with application to clinical trials. *Biometrika*. 1989; 76:577–583.
- Wilhelm FH, Roth WT, Sackner MA. The lifeShirt: An advanced system for ambulatory measurement of respiratory and cardiac function. *Behaviour Modification*. 2003; 27:671–691.
- Zhang HH, Wahba G, Lin Y, Voelker M, Ferris M, Klein R, Klein B. Variable selection and model building via likelihood basis pursuit. *Journal of the American Statistical Association*. 2004; 99:659–672.

Figure 1. Boxplots of MSEs for different values of r , based on 100 generated data sets. The left panel shows the boxplots for $p = 10$ and the right panel those for $p = 100$.

Figure 2.

The first row shows two sample VT curves after baseline adjustment. The trace in the left panel is an $N + L$ sample, the one in the right panel an $S + L$. The second row shows the same two curves after wavelet smoothing and the third row shows the corresponding trends.

Figure 3. Mean curves for N + L (line) and S + L (dotted line) after preprocessing. The left panel is for traces after the wavelet smoothing, the right panel is for the trends.

Figure 4.

Original (line) and approximated (dotted line) traces for 6 randomly selected subjects. We used $r = 25, 4$, for smoothed data and trends, respectively. The left panel is for traces after the wavelet smoothing, the right panel is for the trends.

Table 1Power of the NBMF testing procedure ($p = 1024$)

$(p - t_k^*)_{\text{for } f_1 = f_2}$ in $(1, \dots, t_k^*)$	SNR for $\sigma^{*2} = 10^2$	
	1	1/3
10	1.0	0.690
20	1.0	0.909
30	1.0	0.999
40	1.0	1.0
t_k^* for $f_1 = f_2$ in $(t_k^*, \dots, p)$	SNR for $\sigma^{*2} = 10^4$	
	1	1/3
50	0.801	0.468
75	0.998	0.729
100	1.0	1.0

Table 2

P-values of test results with smoothed data and trends. The half χ^2 distribution of the Follmann's test statistics for large samples was used. P-values from a permutation test are reported in parentheses.

Time lag	Smoothed	Trend
First 5 minutes	0.1435 (0.1795)	0.4172 (0.4295)
First 10 minutes	0.0564 (0.0405)	0.3641 (0.3710)
First 15 minutes	0.0031 (0.0000)	0.0738 (0.0815)
Last 5 minutes	0.0180 (0.0070)	0.0457 (0.0445)
Last 10 minutes	0.0082 (0.0010)	0.0481 (0.0615)
Whole time	0.0030 (0.0000)	0.0483 (0.0535)