

Dimensionality Reduction for Binary Data through the Projection of Natural Parameters

Andrew J. Landgraf, *The Ohio State University*
Yoonkyung Lee, *The Ohio State University*

Technical Report No. 890

April, 2015

**Department of Statistics
The Ohio State University
1958 Neil Avenue
Columbus, OH 43210-1247**

Dimensionality Reduction for Binary Data through the Projection of Natural Parameters

Andrew J. Landgraf AND Yoonkyung Lee

Abstract

Principal component analysis (PCA) for binary data, known as logistic PCA, has become a popular alternative to dimensionality reduction of binary data. It is motivated as an extension of ordinary PCA by means of a matrix factorization, akin to the singular value decomposition, that maximizes the Bernoulli log-likelihood. We propose a new formulation of logistic PCA which extends Pearson's formulation of a low dimensional data representation with minimum error to binary data. Our formulation does not require a matrix factorization, as previous methods do, but instead looks for projections of the natural parameters from the saturated model. Due to this difference, the number of parameters does not grow with the number of observations and the principal component scores on new data can be computed with simple matrix multiplication. We derive explicit solutions for data matrices of special structure and provide computationally efficient algorithms for solving for the principal component loadings. Through simulation experiments and an analysis of medical diagnoses data, we compare our formulation of logistic PCA to the previous formulation as well as ordinary PCA to demonstrate its benefits.

Keywords: Binary data; Exponential family; Logistic PCA; Principal component analysis

1 Introduction

Principal component analysis (PCA) is perhaps the most popular dimensionality reduction technique (see Jolliffe, 2002, for example). It is useful for data compression, visualization, and feature discovery. PCA can be motivated either by maximizing the variance of linear combinations of the variables (Hotelling, 1933) or by minimizing the reconstruction error of a lower dimensional projection of the cases (Pearson, 1901). There is an implicit connection between standard PCA and the Gaussian distribution in Pearson's formulation. Tipping and Bishop (1999) also showed that PCA provides the maximum likelihood estimate for a factor model, where the data are assumed to be Gaussian.

Although PCA is commonly used for dimensionality reduction for various types of data in practice, the fact that PCA finds a low-rank subspace by implicitly minimizing the reconstruction error under the squared error loss renders direct application of PCA to non-Gaussian data such as binary responses or counts conceptually unappealing. Moreover, the

probabilistic interpretation of PCA with normal likelihood in Tipping and Bishop (1999) suggests the possibility of proper likelihood-based loss functions in defining the best subspace of a given rank for other types of data. With this motivation, Collins et al. (2001) proposed a generalization of PCA to exponential family data using the generalized linear model (GLM) framework, and Schein et al. (2003), Tipping (1998) and de Leeuw (2006) examined similar generalizations for binary data in particular, using the Bernoulli likelihood, which is referred to as logistic PCA. Generalized PCA estimates the natural parameters of a data matrix in a lower dimensional subspace by minimizing the negative log-likelihood under an exponential family distribution. In the Gaussian case, generalized PCA is shown to be equivalent to the truncated singular value decomposition (SVD).

In this paper, we argue that Collins et al. (2001)’s logistic PCA is more closely related to SVD than PCA because it aims at a low-rank factorization of the natural parameters matrix. Consequently, each case has its own latent factor associated with it and the number of parameters increases with the number of cases. The drawback of the formulation becomes apparent when it comes to prediction. To apply logistic PCA to new data, one needs to carry out another matrix factorization, which is prone to overfit. This is in contrast with standard PCA where the principal component scores for the new data are simply given by linear combinations of the observed values of the variables.

Retaining the structure of standard PCA, we generalize PCA in such a way that the principal component scores are linear functions of the data. This is done by interpreting Pearson (1901)’s formulation in a slightly different manner. A projection of the data with minimum reconstruction error under squared loss can be viewed alternatively as a projection of the *natural parameters* of a saturated model with minimum *deviance* for Gaussian data. This alternative interpretation allows a coherent generalization of standard PCA to exponential family distributions. When the distribution is Gaussian, this generalization simplifies to standard PCA. Due to the prevalence of binary data and for simplicity of exposition of the new generalization of PCA, we focus on logistic PCA in this paper.

Our formulation has several benefits over Collins et al. (2001)’s formulation. The number of parameters does not increase with the number of observations, the principal component scores are easily interpretable as linear functions of the data, and applying principal components to a new set of data only requires a matrix multiplication. Furthermore, while very little is known about solutions to Collins et al. (2001)’s formulation, some explicit solutions to our formulation can be derived for data matrices with special structures.

Computationally, our formulation of logistic PCA, like the previous one, leads to a non-convex problem. We derive two practical algorithms for generating solutions. The first algorithm involves iterative spectral decompositions, using the majorization-minimization (MM) algorithm (see Hunter and Lange, 2004). The second algorithm involves a convex relaxation of low-rank projection matrices. Using these algorithms, we apply our formulation of logistic PCA to several datasets, examining the advantages and trade-offs with existing methods.

The rest of the paper is organized as follows. Section 2 gives background on PCA and logistic PCA. In Section 3, we introduce our formulation of generalized PCA and qualitatively

compare it to the previous formulation. In Section 4, we derive the first-order optimality conditions for logistic PCA solutions and characterize explicit solutions that satisfy the conditions for data matrices with special structures. In Section 5, we derive two algorithms for logistic PCA. Section 6 shows the potential benefits of our formulation via data analyses with simulated and real data. Finally, Section 7 concludes the paper with a discussion and possible extensions of the proposed logistic PCA to exponential family data.

2 Background

Pearson (1901) considered a geometric problem of finding an optimal representation of multivariate data in a low dimension with respect to mean squared error. Assume that the data consist of $\mathbf{x}_i \in \mathbb{R}^d, i = 1, \dots, n$. To project the original d -dimensional data into lower dimensions, say, $k < d$, we represent each \mathbf{x} by $\boldsymbol{\mu} + \mathbf{U}\mathbf{U}^T(\mathbf{x} - \boldsymbol{\mu})$, where $\boldsymbol{\mu} \in \mathbb{R}^d$ and \mathbf{U} is a $d \times k$ matrix with orthonormal columns. Pearson (1901) showed that the minimum of the mean squared error of the k -dimensional representation,

$$\min_{\boldsymbol{\mu} \in \mathbb{R}^d, \mathbf{U}^T \mathbf{U} = \mathbf{I}_k} \sum_{i=1}^n \|\mathbf{x}_i - \boldsymbol{\mu} - \mathbf{U}\mathbf{U}^T(\mathbf{x}_i - \boldsymbol{\mu})\|^2 \quad (2.1)$$

is attained when $\boldsymbol{\mu}$ equals the sample mean and \mathbf{U} is a matrix with the first k eigenvectors of the sample covariance matrix.

The MSE criteria is closely linked to a Gaussian assumption. Borrowing the characterization of PCA in Collins et al. (2001), suppose that $\mathbf{x}_i \sim \mathcal{N}(\boldsymbol{\theta}_i, \mathbf{I}_d)$, and $\boldsymbol{\theta}_i$ are constrained to lie in a k -dimensional subspace. That is, $\boldsymbol{\theta}_i$ are in the span of a k -dimensional orthonormal basis $\{\mathbf{b}_l \in \mathbb{R}^d, l = 1, \dots, k\}$ so that $\boldsymbol{\theta}_i = \sum_{l=1}^k a_{il} \mathbf{b}_l$ for some a_{il} . In this case, the negative log-likelihood is proportional to

$$\sum_{i=1}^n \|\mathbf{x}_i - \boldsymbol{\theta}_i\|^2 = \|\mathbf{X} - \mathbf{A}\mathbf{B}^T\|_F^2, \quad (2.2)$$

where \mathbf{X} is the $n \times d$ matrix with \mathbf{x}_i^T in the i th row, \mathbf{A} is an $n \times k$ matrix with elements a_{il} , \mathbf{B} is a $d \times k$ matrix made up of the k basis vectors, and $\|\cdot\|_F$ is the Frobenius norm. It is well known (Eckart and Young, 1936) that this objective function is minimized by the rank k truncated singular value decomposition (SVD) of \mathbf{X} , where \mathbf{B} consists of the first k right singular vectors and \mathbf{A} consists of the first k left singular vectors scaled by the first k singular values or, equivalently, $\mathbf{A} = \mathbf{X}\mathbf{B}$. The fact that the columns in \mathbf{B} span the same subspace as the first k eigenvectors of $\mathbf{X}^T \mathbf{X}$ yields the equivalence in the solutions of PCA and SVD. Note that \mathbf{x}_i in (2.2) is assumed to be mean-centered. Otherwise, \mathbf{x}_i is to be replaced with $\mathbf{x}_i - \boldsymbol{\mu}$ as in (2.1).

Using the equivalence and the alternative formulation of PCA through the mean space approximation under a Gaussian distribution, Collins et al. (2001) extended PCA to exponential family data. Such an extension would be useful for handling binary, count, or

non-negative data that abound in practice since the Gaussian assumption may be inappropriate. The authors proposed to extend PCA to exponential family distributions in a similar way that generalized linear models (GLMs) (McCullagh and Nelder, 1989) extend linear regression to response variables that are non-Gaussian. Assuming that $\mathbf{x}_1, \dots, \mathbf{x}_n$ are a sample from an exponential family distribution with corresponding natural parameters $\theta_1, \dots, \theta_n$, the goal of generalized PCA is to minimize the negative log-likelihood, subject to the constraint that the estimated natural parameters belong to a k -dimensional subspace. This is done by approximating the $n \times d$ matrix of the natural parameters $\Theta = [\theta_{ij}]$ by factorization of the form $\Theta = \mathbf{A}\mathbf{B}^T$, where \mathbf{A} and \mathbf{B} are of rank k .

Based on this generalization of PCA, Schein et al. (2003) and de Leeuw (2006) have developed algorithms specifically for binary data. Schein et al. (2003) also extended the specification of the natural parameter space in Collins et al. (2001) by introducing variable main effects or biases, $\boldsymbol{\mu}$, so that $\Theta = \mathbf{1}_n\boldsymbol{\mu}^T + \mathbf{A}\mathbf{B}^T$. In an extension of probabilistic PCA (Tipping and Bishop, 1999), Tipping (1998) proposed a factor model for binary data with a similar structure, where the case-specific factors in \mathbf{A} are assumed to come from a k dimensional standard normal distribution. The marginal distribution of \mathbf{X} is then used to define a likelihood of \mathbf{B} and $\boldsymbol{\mu}$ and is maximized with respect to \mathbf{B} and $\boldsymbol{\mu}$. Like other methods, estimating principal component scores on new data requires additional computation to solve for the parameters on the new cases. More recently, Lee et al. (2010) extended sparse PCA to binary data by adding a sparsity-inducing L_1 penalty to \mathbf{B} . In recommendation systems, Johnson (2014) showed that logistic PCA can capture the latent features of a binary matrix more efficiently than matrix factorization with a squared error loss.

3 New Formulation of Generalized PCA

We propose a new formulation of generalized PCA and demonstrate its conceptual and computational advantages over the current formulation. For the new formulation, we begin with a new interpretation of standard PCA as a technique for low dimensional projection of the natural parameters of the *saturated model*, which are the same as the data under the normal likelihood.

To elaborate on this perspective, recall that for a data matrix \mathbf{X} , the $d \times k$ matrix with the first k principal component loading vectors minimizes

$$\sum_{i=1}^n \|\mathbf{x}_i - \boldsymbol{\mu} - \mathbf{U}\mathbf{U}^T(\mathbf{x}_i - \boldsymbol{\mu})\|^2 = \|\mathbf{X} - \mathbf{1}_n\boldsymbol{\mu}^T - (\mathbf{X} - \mathbf{1}_n\boldsymbol{\mu}^T)\mathbf{U}\mathbf{U}^T\|_F^2 \quad (3.1)$$

among all $d \times k$ matrices \mathbf{U} such that $\mathbf{U}^T\mathbf{U} = \mathbf{I}_k$. The principal component loadings are given by the first k eigenvectors of the sample covariance matrix. To draw the connection to the Gaussian model and GLMs, assume that x_{ij} are normally distributed with known variance. With the identity function as the canonical link, the natural parameter θ_{ij} is the mean itself in this case, and the saturated model, which is the best possible fit to the data, is the model with natural parameters $\tilde{\theta}_{ij} = x_{ij}$.

For Gaussian data, the natural parameters of the saturated model are equal to the data and minimizing the squared error in (3.1) is equivalent to minimizing the deviance. Hence, standard PCA can be viewed as a technique for projecting the natural parameters of the saturated model into a lower dimensional space by minimizing the Gaussian deviance.

3.1 Alternative formulation to logistic PCA

When the data are binary, assume instead that x_{ij} are from Bernoulli(p_{ij}). The natural parameter for the Bernoulli distribution is $\theta_{ij} = \text{logit } p_{ij}$. Let $\tilde{\theta}$ represent the natural parameters of the saturated model. The saturated model occurs when $p_{ij} = x_{ij}$, which means that

$$\tilde{\theta}_{ij} = \begin{cases} -\infty & \text{if } x_{ij} = 0 \\ \infty & \text{if } x_{ij} = 1 \end{cases}.$$

To apply an equivalent principal component analysis to binary data, we need to minimize the Bernoulli deviance by projecting the natural parameters of the saturated model onto a k -dimensional space. For convenience, define $q_{ij} = 2x_{ij} - 1$, which converts the binary variable from taking values in $\{0, 1\}$ to $\{-1, 1\}$. Let $\mathbf{Q} = 2\mathbf{X} - \mathbf{1}_n\mathbf{1}_d^T$ be the matrix with elements q_{ij} . For practical purposes, we will approximate $\tilde{\theta}_{ij}$ by $m \cdot q_{ij}$ for a large number m and show how the choice of m affects the analysis in Section 6.1.3. Therefore, $\tilde{\Theta} = m\mathbf{Q}$ approximates the matrix of natural parameters for the saturated model.

Define $D(\mathbf{X}; \tilde{\Theta})$ as the deviance of estimated natural parameter matrix $\tilde{\Theta}$ with the data matrix \mathbf{X} . As in standard PCA, the natural parameters are estimated with a matrix of the form $\tilde{\Theta} = \mathbf{1}_n\boldsymbol{\mu}^T + (\tilde{\Theta} - \mathbf{1}_n\boldsymbol{\mu}^T)\mathbf{U}\mathbf{U}^T$. The objective function to minimize is the Bernoulli deviance,

$$\begin{aligned} D(\mathbf{X}; \mathbf{1}_n\boldsymbol{\mu}^T + (\tilde{\Theta} - \mathbf{1}_n\boldsymbol{\mu}^T)\mathbf{U}\mathbf{U}^T) &= -2 \left(\log p(\mathbf{X}; \mathbf{1}_n\boldsymbol{\mu}^T + (\tilde{\Theta} - \mathbf{1}_n\boldsymbol{\mu}^T)\mathbf{U}\mathbf{U}^T) - \log p(\mathbf{X}; \tilde{\Theta}) \right) \\ &= -2 \langle \mathbf{X}, \mathbf{1}_n\boldsymbol{\mu}^T + (\tilde{\Theta} - \mathbf{1}_n\boldsymbol{\mu}^T)\mathbf{U}\mathbf{U}^T \rangle \\ &\quad + 2 \sum_{i=1}^n \sum_{j=1}^d \log \left(1 + \exp(\mu_j + [\mathbf{U}\mathbf{U}^T(\tilde{\theta}_i - \boldsymbol{\mu})]_j) \right), \end{aligned} \quad (3.2)$$

subject to $\mathbf{U}^T\mathbf{U} = \mathbf{I}_k$, where $\langle \mathbf{A}, \mathbf{B} \rangle = \text{tr}(\mathbf{A}^T\mathbf{B})$ is the trace inner product.

3.2 Generalized PCA formulation

As alluded to in the previous section, this methodology can be extended to any distribution in the exponential family. If x comes from a one-parameter exponential family distribution, then the deviance $D(x; \theta)$ is proportional to $\{-x\theta + b(\theta) + c(x)\}$, where θ is the canonical parameter, and $E(X) = b'(\theta)$. Let $g(\cdot)$ be the canonical link function, such that $g(b'(\theta)) = \theta$. The saturated model with the lowest possible deviance occurs when $\tilde{\theta} = g(x)$.

For a given distribution, using the appropriate $b(\theta)$, $g(\cdot)$, and $\tilde{\theta}$, generalized PCA can be cast as

$$\min_{\boldsymbol{\mu} \in \mathbb{R}^d, \mathbf{U}^T \mathbf{U} = \mathbf{I}_k} -\langle \mathbf{X}, \mathbf{1}_n \boldsymbol{\mu}^T + (\tilde{\boldsymbol{\Theta}} - \mathbf{1}_n \boldsymbol{\mu}^T) \mathbf{U} \mathbf{U}^T \rangle + \sum_{i=1}^n \sum_{j=1}^d b\left(\mu_j + [\mathbf{U} \mathbf{U}^T (\tilde{\boldsymbol{\theta}}_i - \boldsymbol{\mu})]_j\right).$$

For example, for the Bernoulli distribution $b(\theta) = \log(1 + \exp(\theta))$ and $\tilde{\theta} = \text{logit } x$, for Poisson $b(\theta) = \exp(\theta)$ and $\tilde{\theta} = \log x$, and for Gaussian $b(\theta) = \theta^2/2$ and $\tilde{\theta} = x$.

Further, this formulation can handle matrices with multiple types of data. Each column of the saturated natural parameter matrix will correspond to the particular member of the exponential family and $b_j(\theta)$ will change by column.

3.3 Comparison to previous techniques

The main advantage of the proposed formulation is that we only solve for the principal component loadings and not simultaneously for the principal component scores. The previous method for logistic PCA posits that the logit of the probability matrix, logit \mathbf{P} , can be represented by a low-rank matrix factorization

$$\boldsymbol{\Theta} = \mathbf{A} \mathbf{B}^T,$$

assuming $\boldsymbol{\mu} = \mathbf{0}$ here to simplify exposition. Our formulation, on the other hand, assumes the logit of the probability matrix has the form

$$\boldsymbol{\Theta} = \tilde{\boldsymbol{\Theta}} \mathbf{U} \mathbf{U}^T.$$

To highlight the difference between the two formulations, we will call the previous formulation logistic SVD (LSVD) and our formulation logistic PCA (LPCA). The $d \times k$ principal component loading matrices for LSVD and LPCA are \mathbf{B} and \mathbf{U} , respectively, and the $n \times k$ matrices of principal component scores are \mathbf{A} and $\tilde{\boldsymbol{\Theta}} \mathbf{U}$. The loading matrices are comparable, but the score matrices take different forms. The form of $\tilde{\boldsymbol{\Theta}} \mathbf{U}$, along with m , can act as an implicit regularizer. While there is no restriction on how large the elements of \mathbf{A} can be, the elements of $\tilde{\boldsymbol{\Theta}} \mathbf{U}$ are bound between $-m\sqrt{d}$ and $m\sqrt{d}$, for instance.

We illustrate a number of advantages of the alternative formulation. Conceptually, when the main effects are not included in logistic SVD, the cases and variables are treated interchangeably. That is, an analysis of \mathbf{X} will produce the same low-rank fitted matrix as an analysis of \mathbf{X}^T , and the loadings of \mathbf{X} will equal the scores of \mathbf{X}^T and vice versa. Logistic PCA, however, will very likely have a different solution for the respective loadings and scores. Since the intent of PCA is typically to explain the dependence of the variables, we believe that it is desirable to maintain the inherent difference between cases and variables in the analysis.

Another difference is that the proposed formulation allows quick and easy evaluation of principal component scores for new data. Let $\mathbf{x}^* \in \{0, 1\}^d$ be a new observation. We wish to calculate the principal component scores for this new data point assuming that the loadings

have already been estimated using another dataset. Let $\hat{\mathbf{B}}$ and $\hat{\mathbf{U}}$ be the principal component loadings estimated from the LSVD and LPCA formulations, respectively. To determine the new principal component scores, $\mathbf{a}^* \in \mathbb{R}^k$, for LSVD, one needs to find

$$\mathbf{a}^* = \arg \min_{\mathbf{a} \in \mathbb{R}^k} \sum_{j=1}^d \left[-x_j^* \mathbf{a}^T \hat{\mathbf{b}}_j + \log \left(1 + \exp(\mathbf{a}^T \hat{\mathbf{b}}_j) \right) \right], \quad (3.3)$$

where $\hat{\mathbf{b}}_j$ is the j th row of $\hat{\mathbf{B}}$. This is equivalent to a logistic regression problem, where the d -dimensional response vector is \mathbf{x}^* and the $d \times k$ design matrix is $\hat{\mathbf{B}}$. The \mathbf{a}^* can be viewed as the coefficient vector that maximizes the likelihood defined through \mathbf{x}^* .

In contrast, the LPCA formulation only requires a matrix multiplication for the new principal component scores:

$$\hat{\mathbf{U}}^T \tilde{\boldsymbol{\theta}}^*,$$

where $\tilde{\boldsymbol{\theta}}^* := m(2\mathbf{x}^* - 1) := m\mathbf{q}^*$ is taken as the approximate natural parameters for \mathbf{x}^* under the saturated model. This process is analogous to computing the principal components for new data in standard PCA, where \mathbf{x}^* itself acts as $\tilde{\boldsymbol{\theta}}^*$, a set of the natural parameters for the saturated model. Further, predicting a low-rank estimate of the natural parameters on a set of new data only requires calculating

$$\hat{\boldsymbol{\Theta}}^* = \tilde{\boldsymbol{\Theta}}^* \hat{\mathbf{U}} \hat{\mathbf{U}}^T. \quad (3.4)$$

Quick evaluation of principal component scores for new data can be particularly useful in a number of situations. For example, principal component regression (or classification) (see Hastie et al., 2009, §3.5) can be extended to logistic PCA when all the covariates are binary. If it is necessary to make predictions for a large amount of new data or predictions are required in real-time, the LSVD method may be too slow. Further, our proposed method will be much more efficient when the number of principal components to retain is chosen by cross validation (Jolliffe, 2002, §6.1.5). In this case, we select the number of components that best reconstruct the dataset on held-out observations, and the cross-validation requires applying logistic PCA to new data repeatedly.

Another major difference between this formulation and the previous one is that the alternative formulation entails much fewer parameters. In particular, LSVD has $kn - \frac{k(k-1)}{2}$ additional parameters in the \mathbf{A} matrix, if the columns are constrained to be orthogonal to each other. This additional number of parameters could potentially be very large and be ripe for over-fitting. As Welling et al. (2008) discussed, logistic SVD can be viewed as an estimation method for a factor model. Instead of marginalizing over the case-specific factors, which are latent variables in factor analysis, logistic SVD takes a degenerate approach to computing point estimates for them. Since the number of latent factors is proportional to the number of observations, overfitting can easily occur.

The alternative formulation of logistic PCA does have an additional parameter, m , that previous formulations do not. We treat m as a tuning parameter. As m gets larger, the elements of $\tilde{\boldsymbol{\Theta}}$ will get closer to $\pm\infty$ and therefore the estimated probabilities will be close

to 0 or 1. Conversely, if m is small, the probability estimates will be close to 0.5. If the user has domain knowledge of the range of the likely probabilities, they can use this to guide their choice of m . We have found that cross validation is an effective way to choose m . Simulations in Section 6.1 show the potential benefits of correctly choosing m .

Because LSVD has many more parameters than LPCA given a rank, LSVD is guaranteed to have a lower in-sample deviance. Despite this, the simulations in Section 6.1 show that LPCA can do just as well or better at estimating the true probabilities if m and k are chosen properly. Further, the loadings learned in LSVD may not generalize as well as those learned in LPCA, as the example in Section 6.2.3 exhibits.

3.4 Number of Principal Components

Selecting the appropriate dimensions for effective data representation is a common issue for dimensionality reduction techniques. There has been relatively little discussion previously in the literature of how to select the number of PCs in logistic PCA. Lee et al. (2010) derived a BIC heuristic to select the degree of sparsity for sparse logistic PCA and Li and Tao (2010) proposed a Bayesian version of exponential family PCA with a prior on the loadings that controls the number of principal components. We propose a few methods for selection of dimensionality in logistic PCA, motivated by the current practices in standard PCA and the dual interpretation of squared error as the deviance for a Gaussian model.

One common approach in standard PCA is to look at the cumulative percent of the variance explained and select the number of components such that a chosen proportion, γ , is met or exceeded. Let $\hat{\mathbf{U}}_k$ be the rank k estimate of the principal component loadings. The criteria will choose a rank k model if k is the smallest integer such that

$$1 - \|\mathbf{X} - (\mathbf{1}_n \hat{\boldsymbol{\mu}}^T + (\mathbf{X} - \mathbf{1}_n \hat{\boldsymbol{\mu}}^T) \hat{\mathbf{U}}_k \hat{\mathbf{U}}_k^T)\|_F^2 / \|\mathbf{X} - \mathbf{1}_n \hat{\boldsymbol{\mu}}^T\|_F^2 > \gamma.$$

Similarly for logistic PCA, if $D(\mathbf{X}; \hat{\boldsymbol{\Theta}}_k)$ is the Bernoulli deviance of the rank- k principal component loadings, $\hat{\boldsymbol{\Theta}}_k = \mathbf{1}_n \hat{\boldsymbol{\mu}}^T + (\tilde{\boldsymbol{\Theta}} - \mathbf{1}_n \hat{\boldsymbol{\mu}}^T) \hat{\mathbf{U}}_k \hat{\mathbf{U}}_k^T$, with the data \mathbf{X} , then we could choose the smallest integer k such that

$$1 - \frac{D(\mathbf{X}; \hat{\boldsymbol{\Theta}}_k)}{D(\mathbf{X}; \mathbf{1}_n \hat{\boldsymbol{\mu}}^T)} > \gamma.$$

This criterion has a similar interpretation as in standard PCA that at least $100\gamma\%$ of the deviance is explained by k principal component loadings. Notice that, as expected, 100% of the deviance will be explained by d components because $\hat{\mathbf{U}}_d \hat{\mathbf{U}}_d^T = \mathbf{I}_d$ and $D(\mathbf{X}; \tilde{\boldsymbol{\Theta}}) = 0$ by definition.

Another approach from standard PCA is to create a scree plot of the percent of variance explained by each component and look for an elbow in the plot. The same analogy can be made to logistic PCA, however with a modified definition of the percent of reduction in deviance for additional components. For logistic PCA, the principal component loadings matrices are not necessarily nested, meaning the first $k - 1$ columns of $\hat{\mathbf{U}}_k$ do not necessarily

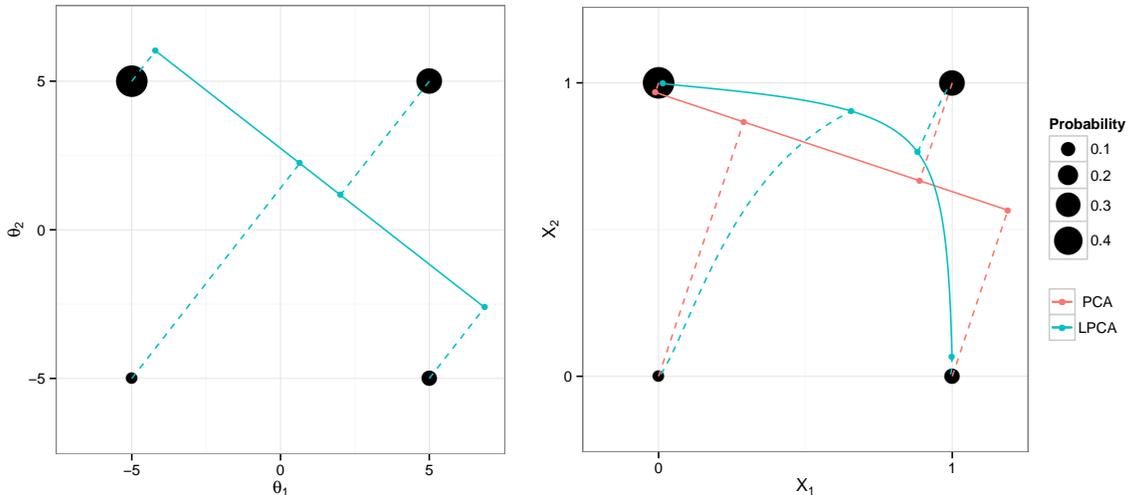


Figure 1: Logistic PCA projection of a two-dimensional Bernoulli distribution in the natural parameter space (left) and in the probability space (right) with $m = 5$ compared to the PCA projection

equal $\hat{\mathbf{U}}_{k-1}$. For the reason, it would be more appropriate to define the *marginal* percentage of deviance explained by the additional k th component as

$$\frac{D(\mathbf{X}; \hat{\Theta}_{k-1}) - D(\mathbf{X}; \hat{\Theta}_k)}{D(\mathbf{X}; \mathbf{1}_n \hat{\boldsymbol{\mu}}^T)}.$$

If d is large, these non-sequential procedures could potentially take very long. We recommend only calculating these quantities until the cumulative percent of deviance explained is fairly high, which may be much smaller than d in many situations.

3.5 Geometry of the Projection

The geometry of logistic principal components is illustrated with a distribution of two Bernoulli random variables. With two dimensions ($d = 2$), there are only four possible responses as indicated by the four points in the right panel of Figure 1. The areas of the points in the figure are proportional to the probabilities that are randomly assigned to them. The probabilities sum to one and specify the distribution, which is taken as a population version of data.

The first step of the proposed logistic PCA is to transform the data into the natural parameters of the saturated model, which is what is shown on the left panel of Figure 1. Here, $m = 5$ was used to represent the saturated model. The one-dimensional linear projection is then performed in the natural parameter space, and is chosen to minimize the Bernoulli deviance. In this example, we minimized the deviance by a grid search since there is only one free parameter for the projection. The solid line in the left panel is the rank-one

space for logit parameters, the dots on the line are the projected values, and the dashed lines show the correspondence between the saturated model parameters and their projections.

Next, we transform this linear projection to the probability space by taking the inverse logit of the values, which is shown as the blue line in the right panel of Figure 1. The one-dimensional space for the probabilities induced by logistic PCA is now non-linear and the projections are no longer orthogonal. Also shown on this plot is the one-dimensional space (red line) from standard PCA. The projections from standard PCA are not constrained to be between 0 and 1, as is seen for the projections of $(1, 0)$ and $(0, 1)$. In addition, the fitted probabilities from logistic PCA seem to be closer to the original data values than standard PCA for all except $(0, 0)$, which has the smallest probability.

4 Logistic PCA for Patterned Data

The properties of standard PCA solutions are well understood algebraically. Under certain assumptions, the solutions are explicitly known. For example, when the variables are uncorrelated, the loadings are the standard bases and the principal components are ordered from highest variance to lowest. In contrast, not much is known about the solutions of logistic PCA or logistic SVD. To obtain analogous results for patterned data with logistic PCA, we derive necessary conditions for the solutions first, and find solutions that satisfy (or nearly satisfy) these optimality conditions under different sets of assumptions on data matrices. These results help us gain a better understanding of logistic PCA.

4.1 First-Order Optimality Conditions

To enforce orthonormality $\mathbf{U}^T \mathbf{U} = \mathbf{I}_k$, we add a constraint to the objective in (3.2) via the method of Lagrange multipliers. The Lagrangian is

$$L(\mathbf{U}, \boldsymbol{\mu}, \Lambda) = D(\mathbf{X}; \mathbf{1}_n \boldsymbol{\mu}^T + (\tilde{\boldsymbol{\Theta}} - \mathbf{1}_n \boldsymbol{\mu}^T) \mathbf{U} \mathbf{U}^T) + \text{tr}(\Lambda(\mathbf{U}^T \mathbf{U} - \mathbf{I}_k)),$$

where Λ is a $k \times k$ symmetric matrix of Lagrange multipliers (Wen and Yin, 2013).

Taking the gradient of the Lagrangian with respect to \mathbf{U} , $\boldsymbol{\mu}$, and Λ and setting them equal to $\mathbf{0}$, we obtain the first-order optimality conditions for the solution of logistic PCA:

$$\left[(\mathbf{X} - \hat{\mathbf{P}})^T (\tilde{\boldsymbol{\Theta}} - \mathbf{1}_n \boldsymbol{\mu}^T) + (\tilde{\boldsymbol{\Theta}} - \mathbf{1}_n \boldsymbol{\mu}^T)^T (\mathbf{X} - \hat{\mathbf{P}}) \right] \mathbf{U} = \mathbf{U} \Lambda \quad (4.1)$$

$$(\mathbf{I}_d - \mathbf{U} \mathbf{U}^T) (\mathbf{X} - \hat{\mathbf{P}})^T \mathbf{1}_n = \mathbf{0}_d, \text{ and} \quad (4.2)$$

$$\mathbf{U}^T \mathbf{U} = \mathbf{I}_k. \quad (4.3)$$

The matrix $\hat{\mathbf{P}}$ has the estimate of p_{ij} at \mathbf{U} and $\boldsymbol{\mu}$, $\hat{p}_{ij} = \sigma(\mu_j + [\mathbf{U} \mathbf{U}^T (\tilde{\boldsymbol{\Theta}}_i - \boldsymbol{\mu})]_j)$, as its ij th element, where $\sigma(\theta) = \frac{\exp(\theta)}{1 + \exp(\theta)}$ is the inverse logit (or sigmoid) function. The details of

the calculation can be found in Appendix A.1. For the following sections, let

$$\mathbf{C}^m := \left[(\mathbf{X} - \hat{\mathbf{P}})^T (\tilde{\Theta} - \mathbf{1}_n \boldsymbol{\mu}^T) + (\tilde{\Theta} - \mathbf{1}_n \boldsymbol{\mu}^T)^T (\mathbf{X} - \hat{\mathbf{P}}) \right],$$

which is labeled to explicitly state the dependence of $\tilde{\Theta}$ and $\hat{\mathbf{P}}$ on m .

Applying the Lagrangian method of multipliers to the standard PCA formulation (2.1) yields $[(\mathbf{X} - \mathbf{1}_n \boldsymbol{\mu}^T)^T (\mathbf{X} - \mathbf{1}_n \boldsymbol{\mu}^T)] \mathbf{U} = \mathbf{U} \Lambda$ as part of the first-order optimality conditions, which is very similar to the form of equation (4.1). Unlike standard PCA, (4.1) is nonlinear in \mathbf{U} and the solution is not known in closed form through an eigen-decomposition because the matrix in the left hand side of (4.1) depends on \mathbf{U} through $\hat{\mathbf{P}}$. However, we can derive some explicit results for special cases using the optimality conditions.

4.2 Independence

There are a few natural extremes of data dependence that a given binary dataset can exhibit. On one end of the spectrum, all of the columns can be in the span of a single vector. In this case, all the columns are equal to each other, and the dataset can be reconstructed with arbitrary precision with a rank-one approximation.

The opposite end of the spectrum is when the variables are independent of each other. In standard PCA, this implies that the covariance matrix is diagonal, so the principal component loadings are the standard basis vectors. Analogous to standard PCA results, we show below that, if the l th column of a dataset is uncorrelated with the other $d - 1$ columns and its column mean is $\frac{1}{2}$ then the l th standard basis vector, \mathbf{e}_l , satisfies the first-order optimality conditions of logistic PCA with $k = 1$. If its column mean is not equal to $\frac{1}{2}$, \mathbf{e}_l can nearly satisfy the optimality conditions with large enough m .

Let X_j be the length n vector of the j th column of \mathbf{X} and $\bar{X}_j = \mathbf{1}_n^T X_j / n$ be the corresponding column mean.

Theorem 1. *Assume that $X_l^T X_j / n = \bar{X}_l \bar{X}_j$, for all $j \neq l$, i.e. the l th variable is uncorrelated with all other variables.*

- (i) *If $\bar{X}_l = \frac{1}{2}$, then $\mathbf{u} = \mathbf{e}_l$, the l th standard basis vector, satisfies the first-order optimality conditions of logistic PCA, regardless of m . That is,*

$$\mathbf{C}^m \mathbf{e}_l - \lambda_m \mathbf{e}_l = \mathbf{0},$$

for some λ_m .

- (ii) *If $\bar{X}_l \neq \frac{1}{2}$, then the first-order optimality conditions can be satisfied as close as desired with $\mathbf{u} = \mathbf{e}_l$ for m large enough. Formally, for any $\epsilon > 0$, there exists m_0 such that, for all $m > m_0$,*

$$\|\mathbf{C}^m \mathbf{e}_l - \lambda_m \mathbf{e}_l\|^2 < \epsilon,$$

with $\lambda_m = 0$.

The proof is given in Appendix A.2.1. If multiple columns are uncorrelated with the remaining columns, this result easily generalizes to larger k . For example, if k columns are uncorrelated with all other columns, then a rank k solution comprising of the the corresponding k standard basis vectors can be made arbitrarily close to (or exactly if the column means equal $\frac{1}{2}$) satisfying the necessary conditions (4.1)–(4.3).

This leads to a natural question: when there are multiple candidate solutions, which one decreases the deviance the most?

Theorem 2. *For logistic PCA with $k = 1$, the standard basis vector which decreases deviance the most is the one corresponding to column with mean closest to $\frac{1}{2}$.*

The proof is given in Appendix A.2.2. This result corresponds to the ordering of variables by variance in standard PCA. The variables with means closest to $\frac{1}{2}$ have the largest variance. If the variables are independent, the variance explained will be largest with a standard basis vector corresponding to the variable with largest variance. Similar to the previous theorem, this theorem can be easily extended to k larger than 1. In this case, the loading matrix made up of the the standard basis vectors corresponding to the k columns that are closest to $\frac{1}{2}$ will decrease the deviance the most out of all loading matrices that comprise of standard basis vectors.

4.3 Compound symmetry

With independence and perfect correlation being two extremes of the structure of the data, somewhere in the middle is compound symmetry. A covariance matrix Σ is compound symmetric if the diagonals are constant ($\Sigma_{jj} = c_1$ for all j) and the off-diagonals are all equal to each other ($\Sigma_{jk} = c_2$ for all $j \neq k$). The compound symmetry of Σ implies equal correlations among the variables. If this is the case, $\frac{1}{\sqrt{d}}\mathbf{1}_d$ is an eigenvector of the covariance matrix. We show that, under more limiting conditions, $\frac{1}{\sqrt{d}}\mathbf{1}_d$ satisfies the optimality conditions for logistic PCA when $\mathbf{Q}^T\mathbf{Q}$ is compound symmetric.

$\mathbf{Q}^T\mathbf{Q}$ has a natural interpretation. The diagonals always equal n and the jk th off-diagonal is the number of records in which the j th and k th variables are the same minus the number of records in which the j th and k th variables differ. It measures how much the j th and k th variables agree with each other, and can range from $-n$ (total disagreement) to n (total agreement). $\mathbf{Q}^T\mathbf{Q}$ is compound symmetric if all the bivariate agreements are the same.

Theorem 3. *Consider logistic PCA without main effects, where $\boldsymbol{\mu} = \mathbf{0}$. Assume that $\mathbf{Q}^T\mathbf{Q}$ is compound symmetric. If β 's exist such that the following condition is satisfied,*

$$\sigma \left(\frac{m}{d} \sum_{l \notin \{j,k\}} q_{il} \right) = \frac{1}{2} + \sum_{l \notin \{j,k\}} q_{il} \beta_{jk,l}, \text{ for all } j \neq k, j, k = 1, \dots, d, \text{ and } i = 1, \dots, n \quad (4.4)$$

then $\mathbf{u} = \frac{1}{\sqrt{d}}\mathbf{1}_d$ satisfies the first-order optimality conditions for logistic PCA, as characterized by equations (4.1)–(4.3).

The proof is given in Appendix A.2.3. Equation (4.4) states that the fitted probabilities can be represented as an affine function of the q_{ij} s. When there are fewer columns, the condition is more likely to be met.

Corollary 1. *If $d \leq 4$ and $\mathbf{Q}^T \mathbf{Q}$ is compound symmetric, then (4.4) is satisfied and therefore $\mathbf{u} = \frac{1}{\sqrt{d}} \mathbf{1}_d$ satisfies the first-order optimality conditions. When $d = 2$, $\mathbf{Q}^T \mathbf{Q}$ is always compound symmetric.*

Proof. For $d = 4$, without loss of generality, let $j = 1$ and $k = 2$.

$$\sigma\left(\frac{m}{4}(q_{i3} + q_{i4})\right) = \begin{cases} \sigma(m/2) & \text{if } q_{i3} = q_{i4} = 1 \\ \frac{1}{2} & \text{if } q_{i3} \neq q_{i4} \\ \sigma(-m/2) & \text{if } q_{i3} = q_{i4} = -1 \end{cases}.$$

Since $\sigma(-m/2) = 1 - \sigma(m/2)$,

$$\sigma\left(\frac{m}{4}(q_{i3} + q_{i4})\right) = \frac{1}{2} + q_{i3}\beta_{12,3} + q_{i4}\beta_{12,4}$$

where

$$\beta_{12,3} = \beta_{12,4} = \frac{\sigma(m/2) - 1/2}{2}.$$

The same can be shown for $d = 2$ or $d = 3$. □

5 Computation

Optimizing for the principal component loadings of logistic PCA is difficult because of the non-convex objective function and the orthonormality constraint. We derive two algorithms for generating solutions. The first is guaranteed to find a local solution and the second finds a global solution to a relaxed version of logistic PCA.

5.1 Majorization-minimization (MM) algorithm

One approach to minimizing the deviance is to iteratively minimize simpler objectives. Majorization-minimization (Hunter and Lange, 2004) seeks to solve difficult optimization problems by majorizing the objective function with a simpler objective, and minimizing the majorizing function. The majorization function must be equal to or greater than the original objective for all inputs, and equal to it at the current input value.

The deviance of a single estimated natural parameter θ is quadratically approximated at $\theta^{(t)}$ by

$$\begin{aligned} -2 \log p(x; \theta) &= -2x\theta + 2 \log(1 + \exp(\theta)) \\ &\approx -2x\theta^{(t)} + 2 \log(1 + \exp(\theta^{(t)})) + 2(\hat{p}^{(t)} - x)(\theta - \theta^{(t)}) + \hat{p}^{(t)}(1 - \hat{p}^{(t)})(\theta - \theta^{(t)})^2 \\ &\leq -2x\theta^{(t)} + 2 \log(1 + \exp(\theta^{(t)})) + 2(\hat{p}^{(t)} - x)(\theta - \theta^{(t)}) + \frac{1}{4}(\theta - \theta^{(t)})^2, \end{aligned}$$

where $\hat{p}^{(t)} = \sigma(\theta^{(t)})$. The inequality is due to the variance of a Bernoulli random variable being bounded above by 1/4. de Leeuw (2006) showed that the deviance itself is majorized by this same function.

Therefore, the deviance for the whole matrix is majorized by

$$\begin{aligned} & \sum_{i,j} \left\{ -2x_{ij}\theta_{ij}^{(t)} + 2\log(1 + \exp(\theta_{ij}^{(t)})) + 2(\hat{p}_{ij}^{(t)} - x_{ij})(\theta_{ij} - \theta_{ij}^{(t)}) + \frac{1}{4}(\theta_{ij} - \theta_{ij}^{(t)})^2 \right\} \\ &= \frac{1}{4} \sum_{i,j} (\theta_{ij} - z_{ij}^{(t)})^2 + C \end{aligned}$$

where C is a constant that does not depend on θ_{ij} and

$$z_{ij}^{(t)} = \theta_{ij}^{(t)} + 4[x_{ij} - \sigma(\theta_{ij}^{(t)})] \quad (5.1)$$

are the working variables in the t th iteration. Further, let $\mathbf{Z}^{(t)}$ be a matrix whose ij th element equals $z_{ij}^{(t)}$. The working variables have a similar form to the so-called adjusted response used in the iteratively reweighted least squares algorithm for generalized linear models (McCullagh and Nelder, 1989). Instead of having weights equal to the estimated variance at the current estimates, we use the upper bound, which allows for minimization of the majorization function.

The logistic PCA objective function can be majorized around estimates of $\mathbf{U}^{(t)}$ and $\boldsymbol{\mu}^{(t)}$ as

$$D(\mathbf{X}; \mathbf{1}_n \boldsymbol{\mu}^T + (\tilde{\boldsymbol{\Theta}} - \mathbf{1}_n \boldsymbol{\mu}^T) \mathbf{U} \mathbf{U}^T) \leq \frac{1}{4} \sum_{i,j} \left(\mu_j + [\mathbf{U} \mathbf{U}^T (\tilde{\boldsymbol{\theta}}_i - \boldsymbol{\mu})]_j - z_{ij}^{(t)} \right)^2 + C,$$

and hence, the next iterates of $\mathbf{U}^{(t+1)}$ and $\boldsymbol{\mu}^{(t+1)}$ can be obtained by minimizing

$$\sum_{i,j} \left(\mu_j + [\mathbf{U} \mathbf{U}^T (\tilde{\boldsymbol{\theta}}_i - \boldsymbol{\mu})]_j - z_{ij}^{(t)} \right)^2 = \|\mathbf{1}_n \boldsymbol{\mu}^T + (\tilde{\boldsymbol{\Theta}} - \mathbf{1}_n \boldsymbol{\mu}^T) \mathbf{U} \mathbf{U}^T - \mathbf{Z}^{(t)}\|_F^2. \quad (5.2)$$

Given initial estimates for \mathbf{U} and $\boldsymbol{\mu}$, a solution can be found by iteratively minimizing equation (5.2), subject to orthonormality constraint $\mathbf{U}^T \mathbf{U} = \mathbf{I}_k$. With fixed $\boldsymbol{\mu}$, the minimizer of the majorizing function can be found by expanding equation (5.2). Letting $\tilde{\boldsymbol{\Theta}}_c := \tilde{\boldsymbol{\Theta}} - \mathbf{1}_n \boldsymbol{\mu}^T$ and $\mathbf{Z}_c^{(t)} := \mathbf{Z}^{(t)} - \mathbf{1}_n \boldsymbol{\mu}^T$, then

$$\begin{aligned} \arg \min_{\mathbf{U}^T \mathbf{U} = \mathbf{I}_k} \|\tilde{\boldsymbol{\Theta}}_c \mathbf{U} \mathbf{U}^T - \mathbf{Z}_c^{(t)}\|_F^2 &= \arg \min tr \left(\mathbf{U} \mathbf{U}^T \tilde{\boldsymbol{\Theta}}_c^T \tilde{\boldsymbol{\Theta}}_c \mathbf{U} \mathbf{U}^T \right) - tr \left(\mathbf{U} \mathbf{U}^T \tilde{\boldsymbol{\Theta}}_c^T \mathbf{Z}_c^{(t)} \right) - tr \left((\mathbf{Z}_c^{(t)})^T \tilde{\boldsymbol{\Theta}}_c \mathbf{U} \mathbf{U}^T \right) \\ &= \arg \min tr \left[\mathbf{U}^T \left(\tilde{\boldsymbol{\Theta}}_c^T \tilde{\boldsymbol{\Theta}}_c - \tilde{\boldsymbol{\Theta}}_c^T \mathbf{Z}_c^{(t)} - (\mathbf{Z}_c^{(t)})^T \tilde{\boldsymbol{\Theta}}_c \right) \mathbf{U} \right] \\ &= \arg \max tr \left[\mathbf{U}^T \left(\tilde{\boldsymbol{\Theta}}_c^T \mathbf{Z}_c^{(t)} + (\mathbf{Z}_c^{(t)})^T \tilde{\boldsymbol{\Theta}}_c - \tilde{\boldsymbol{\Theta}}_c^T \tilde{\boldsymbol{\Theta}}_c \right) \mathbf{U} \right]. \end{aligned}$$

The trace is maximized when \mathbf{U} consists of the first k eigenvectors of the symmetric matrix $\tilde{\boldsymbol{\Theta}}_c^T \mathbf{Z}_c^{(t)} + (\mathbf{Z}_c^{(t)})^T \tilde{\boldsymbol{\Theta}}_c - \tilde{\boldsymbol{\Theta}}_c^T \tilde{\boldsymbol{\Theta}}_c$ (Fan, 1949).

input : Binary data matrix (\mathbf{X}), m , number of principal components (k), convergence criteria (ϵ)
output: $d \times k$ orthonormal matrix of principal component loadings (\mathbf{U}) and column main effects ($\boldsymbol{\mu}$)

Set $t = 0$ and initialize $\boldsymbol{\mu}^{(0)}$ and $\mathbf{U}^{(0)}$. We recommend setting $\mu_j^{(0)} = \text{logit}\bar{X}_j$ and setting $\mathbf{U}^{(0)}$ to the first k right singular vectors of \mathbf{Q}

repeat

1. $t \leftarrow t + 1$
2. Set the working variables
 $z_{ij}^{(t)} = \theta_{ij} + 4[x_{ij} - \sigma(\theta_{ij})]$ where $\theta_{ij} = \mu_j^{(t-1)} + [\mathbf{U}^{(t-1)}(\mathbf{U}^{(t-1)})^T(\tilde{\boldsymbol{\theta}}_i - \boldsymbol{\mu}^{(t-1)})]_j$
3. $\boldsymbol{\mu}^{(t)} = \frac{1}{n}(\mathbf{Z}^{(t)} - \tilde{\boldsymbol{\Theta}}\mathbf{U}^{(t-1)}(\mathbf{U}^{(t-1)})^T)\mathbf{1}_n$
4. Carry out the eigen decomposition of $(\tilde{\boldsymbol{\Theta}} - \mathbf{1}_n(\boldsymbol{\mu}^{(t)})^T)^T(\mathbf{Z}^{(t)} - \mathbf{1}_n(\boldsymbol{\mu}^{(t)})^T) + (\mathbf{Z}^{(t)} - \mathbf{1}_n(\boldsymbol{\mu}^{(t)})^T)^T(\tilde{\boldsymbol{\Theta}} - \mathbf{1}_n(\boldsymbol{\mu}^{(t)})^T) - (\tilde{\boldsymbol{\Theta}} - \mathbf{1}_n(\boldsymbol{\mu}^{(t)})^T)^T(\tilde{\boldsymbol{\Theta}} - \mathbf{1}_n(\boldsymbol{\mu}^{(t)})^T) = E\Lambda E^T$ and set $\mathbf{U}^{(t)}$ to the first k eigenvectors in E

until *Deviance converges*;

Algorithm 1: Majorization-minimization algorithm for logistic PCA

With fixed \mathbf{U} , the minimization of the majorizing function with respect to $\boldsymbol{\mu}$ is a least squares problem, and the majorizing function is minimized by

$$\boldsymbol{\mu} = \frac{1}{n}(\mathbf{Z}^{(t)} - \tilde{\boldsymbol{\Theta}}\mathbf{U}\mathbf{U}^T)\mathbf{1}_n,$$

which can be interpreted as the average differences between the projection of the uncentered saturated natural parameters and the current working variables. Algorithm 1 presents the MM algorithm for logistic PCA.

Since the optimization problem is non-convex, finding the global minimum is difficult in general. Due to the fact that the majorization function is above the deviance and tangent to it at the previous iteration, minimizing the majorization function at each iteration must either decrease the deviance or cause no change at each iteration. While the majorization-minimization does not guarantee a global minimum generally, the quadratic majorization function and the smooth objective does guarantee finding a local minimum (Lange, 2013, Chapter 12).

5.2 Convex relaxation

Another approach to solving for the loadings of logistic PCA is to relax the non-convex domain of the problem to a convex set. Instead of solving over rank- k projection matrices $\mathbf{U}\mathbf{U}^T$, we can optimize over the convex hull of rank k projection matrices, called the Fantope

(Dattorro, 2005), which is defined as

$$\mathcal{F}^k = \text{conv}\{\mathbf{U}\mathbf{U}^T \mid \mathbf{U} \in \mathbb{R}^{d \times k}, \mathbf{U}^T\mathbf{U} = \mathbf{I}_k\} = \{\mathbf{H} \mid \mathbf{0} \preceq \mathbf{H} \preceq \mathbf{I}_d, \text{tr}(\mathbf{H}) = k\}.$$

With the Fantope, k no longer needs to be an integer, but can be any positive number.

The logistic PCA problem then becomes

$$\begin{aligned} & \text{minimize} && D(\mathbf{X}; \mathbf{1}_n \boldsymbol{\mu}^T + (\tilde{\boldsymbol{\Theta}} - \mathbf{1}_n \boldsymbol{\mu}^T) \mathbf{H}) \\ & \text{subject to} && \boldsymbol{\mu} \in \mathbb{R}^d \text{ and } \mathbf{H} \in \mathcal{F}^k. \end{aligned}$$

For fixed $\boldsymbol{\mu}$, this is a convex problem in \mathbf{H} since the objective and the constraints are both convex. Similarly, for fixed \mathbf{H} , this is convex in $\boldsymbol{\mu}$. Hence, the problem is bi-convex.

With fixed \mathbf{H} , minimizing with respect to $\boldsymbol{\mu}$ is similar to a GLM problem and many strategies can be used. We choose to solve for $\boldsymbol{\mu}$ first and leave it fixed while \mathbf{H} is solved for. With $\mathbf{H} = \mathbf{0}$, the deviance is minimized with respect to $\boldsymbol{\mu}$ when $\mu_j = \text{logit} \bar{X}_j$. We do not update $\boldsymbol{\mu}$ again in our implementation, but other strategies, such as an alternating gradient descent, could easily be implemented if desired.

Once $\boldsymbol{\mu}$ is fixed, to solve the relaxed problem with respect to \mathbf{H} , we can use the projected gradient method (Boyd et al., 2003), which involves updating \mathbf{H} at each iteration by

$$\mathbf{H}^{(t+1)} = \Pi_{\mathcal{F}^k} \left(\mathbf{H}^{(t)} - \alpha_t \mathbf{G}^{(t)} \right).$$

$\mathbf{G}^{(t)}$ is the gradient of the deviance with respect to \mathbf{H} evaluated at $\mathbf{H}^{(t)}$, α_t is the step size for the t th iteration, and $\Pi_{\mathcal{F}^k}(\cdot)$ is the Euclidean projection operator onto the convex set \mathcal{F}^k .

This method is guaranteed to produce a solution arbitrarily close to the global minimum as the number of iterations goes to ∞ if the step size, α_t , decreases in an appropriate manner. One way to ensure convergence is to have the step sizes be square summable, but not summable (Boyd et al., 2003), which can be achieved with $\alpha_t = a_0/t$, for instance, where a_0 is the starting step size.

Without taking into account the symmetry of \mathbf{H} , the gradient of the deviance with respect to \mathbf{H} is

$$\frac{\partial D}{\partial \mathbf{H}} = 2 \left(\hat{\mathbf{P}} - \mathbf{X} \right)^T \left(\tilde{\boldsymbol{\Theta}} - \mathbf{1}_n \boldsymbol{\mu}^T \right), \quad (5.3)$$

which is derived in a similar way to (4.1). Since \mathbf{H} is symmetric, the gradient needs to be symmetrized as follows (see Petersen and Pedersen (2012)):

$$\mathbf{G} = \left[\frac{\partial D}{\partial \mathbf{H}} \right] + \left[\frac{\partial D}{\partial \mathbf{H}} \right]^T - \text{diag} \left[\frac{\partial D}{\partial \mathbf{H}} \right]. \quad (5.4)$$

To update \mathbf{H} in the negative gradient direction, the projection operator $\Pi_{\mathcal{F}^k}$ has to be defined explicitly. Vu et al. (2013) showed that the Euclidean projection of a positive definite matrix \mathbf{M} with a spectral decomposition $\mathbf{M} = \sum_{j=1}^d \lambda_j \mathbf{u}_j \mathbf{u}_j^T$ onto \mathcal{F}^k is given by

$$\Pi_{\mathcal{F}^k}(\mathbf{M}) = \sum_{j=1}^d \lambda_j^+(\nu) \mathbf{u}_j \mathbf{u}_j^T, \quad (5.5)$$

input : Binary data matrix (\mathbf{X}), m , rank (k), starting step size (a_0)

output: $d \times d$ rank k Fantope matrix (\mathbf{H})

Set $\mu_j = \text{logit} \bar{X}_j$, $j = 1, \dots, d$

Set $t = 0$ and initialize $\mathbf{H}^{(0)}$

repeat

1. $t \leftarrow t + 1$. Set $\alpha_t = a_0 \frac{1}{t}$
2. Calculate the derivative, $\mathbf{G}^{(t-1)}$, of the deviance at $\mathbf{H}^{(t-1)}$ using equations (5.3) and (5.4)
3. Move in the negative direction of the derivative and project onto the rank k Fantope using equation (5.5), $\mathbf{H}^{(t)} = \Pi_{\mathcal{F}^k} \left(\mathbf{H}^{(t-1)} - \alpha_t \mathbf{G}^{(t-1)} \right)$

until *Deviance converges*;

Algorithm 2: Fantope algorithm for logistic PCA

where $\lambda_j^+(\nu) = \min(\max(\lambda_j - \nu, 0), 1)$ with ν that satisfies $\sum_j \lambda_j^+(\nu) = k$. A quick and easy way to solve for ν is by bisection search, with an initial lower bound of $\lambda_d - k/d$ and an initial upper bound of λ_1 , the largest eigenvalue. Algorithm 2 presents the complete algorithm for the Fantope solution.

Even though this algorithm is guaranteed to find a global minimum of the deviance over the Fantope space, it is of practical interest how to translate the estimated $\hat{\mathbf{H}}$ into an estimated loadings matrix $\hat{\mathbf{U}}$ because we may be more interested in principal component loadings and the scores they induce. For that reason, once the algorithm has converged, we may convert the Fantope matrix $\hat{\mathbf{H}}$ into a projection matrix by setting $\hat{\mathbf{U}}$ equal to the matrix with the first k eigenvectors of $\hat{\mathbf{H}}$. We will discuss the impact of this conversion in Section 6.1.2.

Both the Fantope solution and the projection solution have the previously-stated advantage that applying the principal components analysis to new data only involves matrix multiplications. There are some practical differences however. If the goal of the analysis is to have a small number of principal components which can be interpreted or used for subsequent analyses, the Fantope solution may not be preferred. Even for $k = 1$, it is possible for the solution to be of full rank, no matter d . However, the Fantope algorithm is useful if viewed as a constrained method of approximating a binary matrix, when interpretation is not desired. This type of approximation can be useful when accurate estimates of the fitted probabilities are required, for example, in collaborative filtering. Further, since the problem is convex, it can also be solved more quickly and reliably than the formulation with a projection matrix.

An R (R Core Team, 2015) implementation of both algorithms for logistic PCA can be found at github.com/andland/logisticPCA.

6 Numerical Examples

In this section, we demonstrate how the logistic PCA formulation works on both simulated and real data. In particular, we numerically examine the differences between the previous formulation and the proposed formulation, the effect of m , and the effectiveness of the algorithms.

6.1 Simulation

For comparison of our formulation of logistic PCA (LPCA) with the previous formulation (LSVD), we simulate binary matrices from a family of multivariate Bernoulli mixtures that induce a low-rank structure in both the true probability matrix and the logit matrix. The components or clusters of the Bernoulli mixtures are determined by cluster-specific probability vectors.

6.1.1 Simulation setup

Let k indicate the number of mixture components or clusters of the probability vectors for d -variate binary random vectors making up an $n \times d$ data matrix \mathbf{X} . For a specified n , d , and k , let C_i be the cluster assignment for the i th observation, $i = 1, \dots, n$, which takes one of k values $\{1, \dots, k\}$ with equal probability. Further, we specify k true probability vectors, $\mathbf{p}_c \in [0, 1]^d$, for $c = 1, \dots, k$. In our simulations, the probabilities are independently generated from a Beta(α, β) distribution and X_{ij} given $C_i = c$ are independently generated from Bernoulli(p_{cj}). For a Beta distribution, instead of the shape parameters α and β , we vary the mean $\bar{p} = \frac{\alpha}{\alpha + \beta}$ and concentration parameter $\phi = \frac{\alpha + \beta}{2}$, which is inversely related to the variance of $\bar{p}(1 - \bar{p})/(2\phi + 1)$. If we let \mathbf{A} be an $n \times k$ indicator matrix with $A_{ic} = 1_{\{C_i=c\}}$ and \mathbf{B}_* be a $d \times k$ matrix with the c th column equal to \mathbf{p}_c , then the true probability matrix is

$$\mathbf{P} = [p_{ij}] = \mathbf{A}\mathbf{B}_*^T,$$

or equivalently, the logit matrix is $\Theta = \text{logit } \mathbf{P} = \mathbf{A}\mathbf{B}^T$ with $\mathbf{B} = \text{logit } \mathbf{B}_*$. The accuracy of the approximation of \mathbf{P} through logistic PCA depends on the number of principal components considered and the value of parameter m . To reduce confusion between the true number of clusters k and the number of principal components considered, we will adopt the notation \hat{k} for the number of principal components considered. For all simulations in this section, we set $n = 100$ and $d = 50$. For simulations with $\bar{p} = 0.5$, we did not include the main effects $\boldsymbol{\mu}$ for either LPCA or LSVD. We did this to minimize the differences in implementations and the main effects are likely close to $\mathbf{0}$ when $\bar{p} = 0.5$.

6.1.2 Fantope versus MM

To compare the performance of the MM and Fantope algorithms, a dataset was simulated using a two-component ($k = 2$) mixture of probabilities from a beta distribution with $\bar{p} = 0.5$ and $\phi = 3$. We then estimated the logistic principal component loadings with $\hat{k} = 2$ and

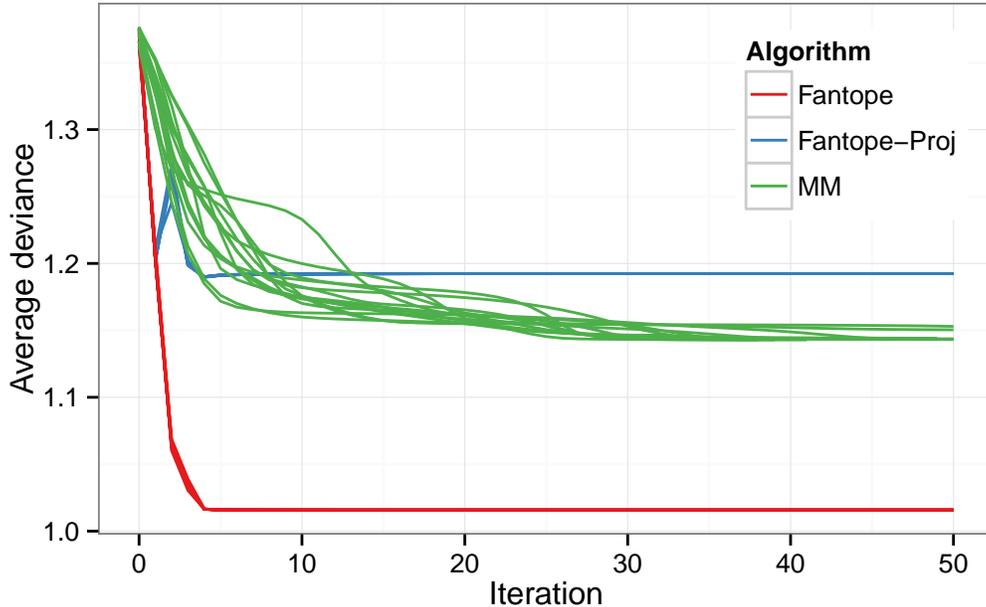


Figure 2: Average deviance by iteration number for solutions from Fantope and MM algorithms with 15 different initial values on a simulated dataset

$m = 4$ fifteen times with both the Fantope algorithm and the MM algorithm, each with a random initialization for $\mathbf{U}^{(0)}$ or $\mathbf{H}^{(0)} = \mathbf{U}^{(0)}(\mathbf{U}^{(0)})^T$. For the Fantope, the starting step size was determined by averaging the inverse of the element-wise second derivatives evaluated at the initial value of $\mathbf{H}^{(0)}$. For each iteration, we kept track of the average deviance of the estimates, which is defined as $D(\mathbf{X}; \hat{\Theta})/(nd)$. Figure 2 displays the average deviance of a sequence of the solutions from the algorithms, each line representing a different initialization. We also kept track of the deviance of the projection matrix estimate from the Fantope estimate for each iteration, which we labeled “Fantope-Proj” in the plot. For this estimate, we use the first \hat{k} eigenvectors of \mathbf{H} to estimate \mathbf{U} .

There are several notable facts in this figure. The deviance for the Fantope solution fairly quickly converges to the same minimum value, regardless of initialization. As expected, the Fantope solution provides a deviance lower bound, while the deviance for the projection matrix resulting from the Fantope solution does not necessarily decrease in the same manner. In fact, the MM algorithm eventually finds solutions better than the “projected” Fantope solution for every initialization. The MM algorithm is more sensitive to initialization, but the range of the deviance values of solutions is fairly narrow in this example.

This numerical comparison suggests that when the projection matrix $\mathbf{U}\mathbf{U}^T$ is concerned, the MM algorithm performs better than the Fantope algorithm followed by projection. For this reason, we use the MM algorithm only in the subsequent analyses. However, it could potentially be useful to run both the MM and Fantope algorithms. The Fantope solution gives a lower bound of the deviance for the global solution with a projection matrix while the

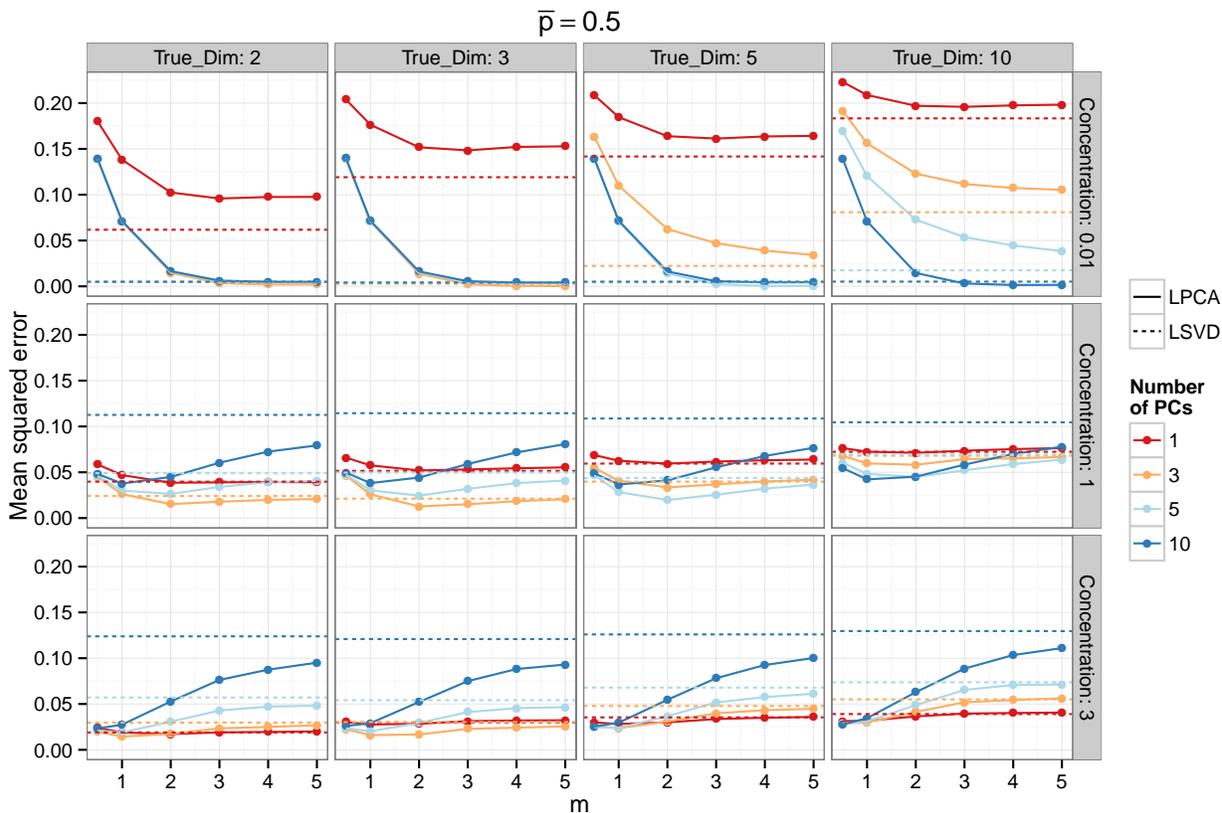


Figure 3: Mean squared error of probability estimates derived from logistic PCA with varying m and logistic SVD in a simulation experiment where the rank of the true probability matrix ranges from 2 to 10 and the extent of concentration of the true probabilities varies from low to high

MM solution and Fantope-Proj solution provide upper bounds. The gap between the lower and upper bounds could give some indication of our confidence in the projection solution. If there is a small gap, we can be fairly sure that the projection solution is close to the global minimum.

6.1.3 Estimate of true probabilities

We compare our formulation of logistic PCA (LPCA) with previous formulations (LSVD) and study the effect of the additional m parameter on the probability estimates. For the numerical study, we simulated 12 binary matrices from a variety of different cluster models with $\bar{p} = 0.5$. We varied the true number of clusters ($k \in \{2, 3, 5, 10\}$) and the concentration parameter for the probabilities ($\phi \in \{0.01, 1, 3\}$). For each of the 12 data matrices, we performed dimensionality reduction with LSVD and LPCA. For both, we varied the estimated number of components ($\hat{k} \in \{1, 3, 5, 10\}$) and for LPCA we also varied m from 0.5 to 5. Again, for

LPCA, we used the MM algorithm and for LSVD we used the iterative SVD algorithm from de Leeuw (2006).

After estimating the probability matrix as $\hat{\mathbf{P}}$, we compared it to the true probability matrix \mathbf{P} by taking the element-wise mean squared error $\|\hat{\mathbf{P}} - \mathbf{P}\|_F^2/(nd)$. Figure 3 shows the results. The first thing to notice is the effect of the concentration parameter. When ϕ is low (the top row of Figure 3), most of the true probabilities are close to 0 or 1 with not much in between. In this situation, having the correct estimated dimension \hat{k} is crucial for both methods and, in fact, having $\hat{k} > k$ does no harm. Further, for LPCA, higher m is better for estimating the probability matrix in this situation. Both of these results are in line with our expectations. When the rank of the estimate is allowed to be higher, the estimates are able to fit the data more closely in general. In doing so, the resulting estimated probabilities will be close to 0 and 1. As stated in Section 3.3, higher values of m enable estimates that are closer to 0 and 1 as well.

On the other hand, if the concentration is high, most of the true probabilities will be close to $\bar{p} = 0.5$. The bottom row of Figure 3 shows the results when this is the case. The opposite conclusions are reached from this scenario. In general, having a lower rank or a lower m is better. This is again as expected.

When the concentration is moderate ($\phi = 1$, the middle row of Figure 3), the true probabilities are generated from a uniform distribution. For LPCA, having the correct estimated dimension is important, but so is m . For each of the different true dimensions, the lowest mean squared error is achieved when the estimated dimension matches the true dimension. This is not true for all values of m , as when $k = 10$, the estimate with $\hat{k} = 10$ is poor for large m . For each of the estimated ranks, there is a local minimum of MSE for choosing m . In contrast to the high and low concentration cases, the results for LSVD do not mirror those of LPCA for the same \hat{k} . Here, an estimated rank of $\hat{k} = 3$ is best for all of the true ranks.

Finally, for any of the given dataset simulations, there is a combination of m and \hat{k} for which LPCA had a mean squared error as small as or smaller than any LSVD estimate. The challenge obviously is to find the best settings and to determine the optimal combination data-adaptively in applications. We show that, for each \hat{k} , using five-fold cross validation is an effective way to choose m . To accomplish this, we randomly split the rows into five groups. For each of the groups, we perform logistic PCA with a given m on all of the data except the rows in that group. We then predict the natural parameters on the held-out group of rows with the given m and the fitted orthonormal matrix using (3.4) and record the predictive deviance. After looping through all five groups, the m with the lowest predictive deviance is used for logistic PCA with all the rows.

The resulting mean squared errors using the data-adaptively chosen m 's are reported in Figure 4. Using this strategy, the low-rank estimates from logistic PCA give more accurate estimates for the high and middle concentration scenarios, while logistic SVD has more accurate estimates for the low concentration scenario, unless $\hat{k} \geq k$. Standard techniques, such as those mentioned in Section 3.4, can be used to choose k for both logistic PCA and logistic SVD.

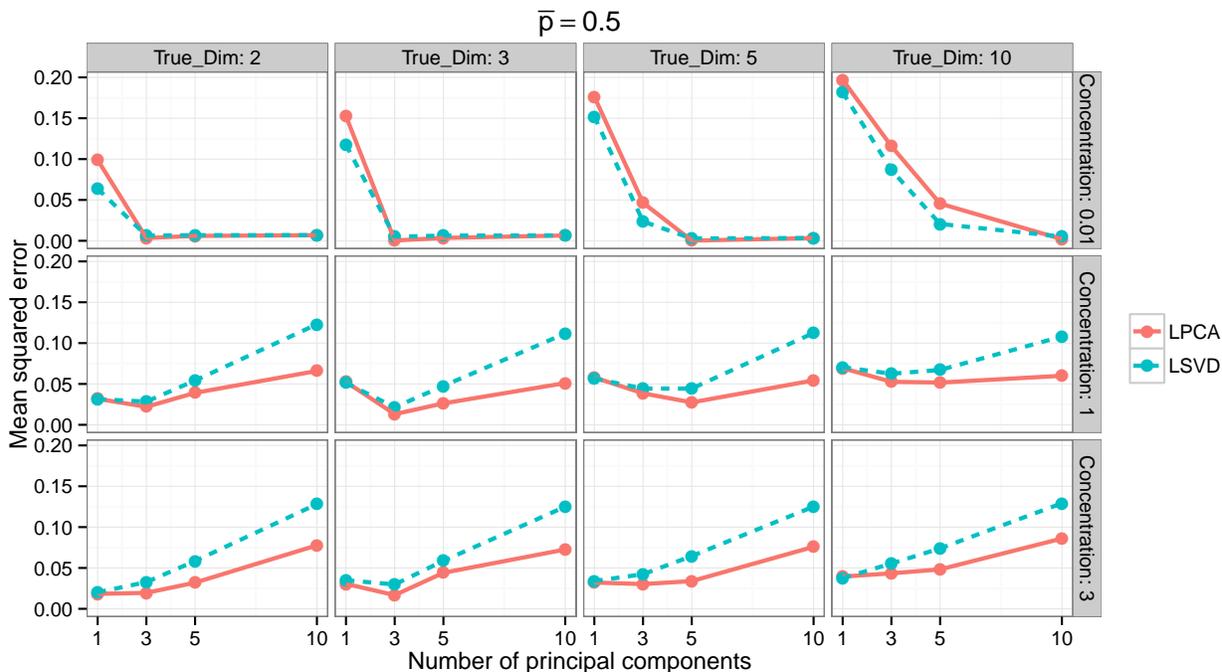


Figure 4: Mean squared error of probability estimates derived from logistic PCA and logistic SVD in the simulation experiment where m in logistic PCA is chosen by cross validation

We have also run another simulation experiment for a more sparse situation with $\bar{p} = 0.1$ with all the other factors kept the same. The plot of the resulting estimated errors is in Figure 11 in Appendix A.3.2. Since the column means are far away from 0.5, we have also included main effects μ in the model. The results are largely the same as described in this section. The only major difference is that estimates of a lower rank than the true rank perform better overall, probably due to the fact that predicting all zeros gives a decent prediction in this case.

6.2 Data Analysis

We present an application of logistic PCA to patient-diagnosis data, which are part of the electronic health records data on 12,000 adult patients admitted to the intensive care units at Ohio State University’s Medical Center from 2007 to 2010. Patients can be admitted to an intensive care unit (ICU) for a wide variety of reasons, some of them frequently co-occurring. While in the ICU, patients are diagnosed with one or more medical conditions of over 800 disease categories from the International Classification of Diseases (ICD-9). It is often of interest to practitioners to study the comorbidity in patients, that is, what medical conditions patients are diagnosed with simultaneously. The latent factor view of principal components seems to be appropriate for describing the concept of comorbidity through a lower dimensional approximation of the disease probabilities or their transformations. Such a

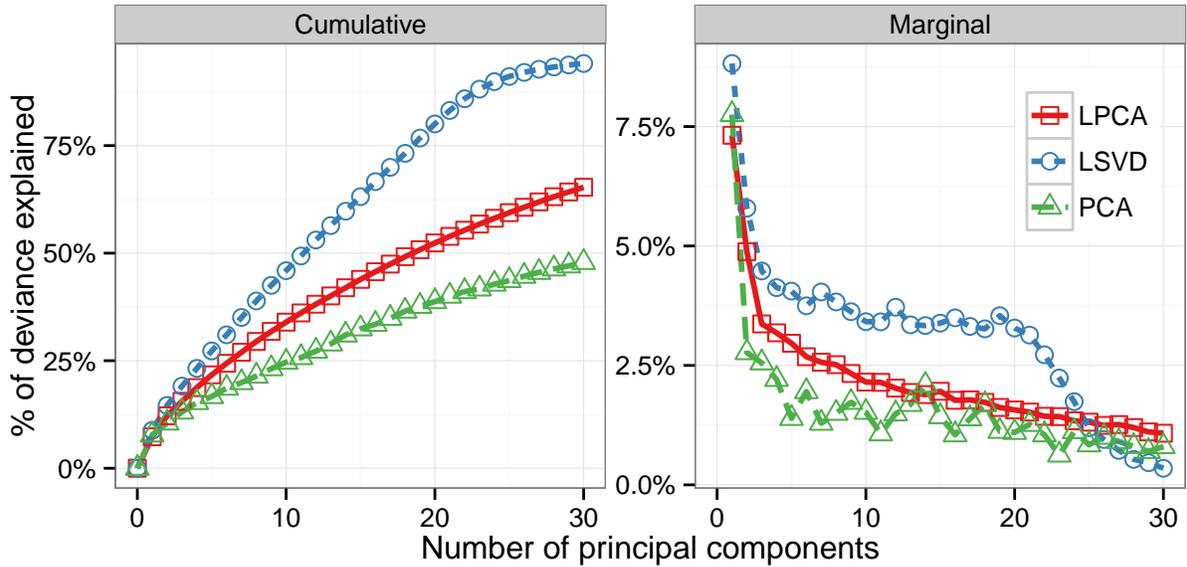


Figure 5: Cumulative and marginal percent of deviance explained by principal components of LPCA, LSVD, and PCA for the patient-diagnosis data.

representation could reveal a common underlying structure capturing simultaneous existence of multiple medical conditions.

We analyzed a random sample of 1,000 patients. There were 584 ICD-9 codes that had at least one of the randomly-selected patients assigned to them. These patient-diagnosis data were organized in a binary matrix, \mathbf{X} , where x_{ij} is 1 if patient i has disease j and 0 otherwise. The proportions of disease occurrences ranged from 0.001 to 0.488 (the maximum corresponding to acute respiratory failure) with a median of 0.005 and the third quartile equals 0.017, meaning that most of the disease categories were rare. For comparison, we also applied logistic SVD and standard PCA to the data.

6.2.1 Selecting number of principal components

As described in Section 3.4, to decide on the number of principal components to include in the model or approximation of the data matrix, we calculated the cumulative percent of deviance and the marginal percent of deviance explained by both LSVD and LPCA, as well as standard PCA. Figure 5 illustrates the change in the percent of deviance explained as the number of components increases. For logistic PCA, m was chosen by five-fold cross validation. In order to have the same m for all k considered, we performed cross validation with $k = 15$, for which $m = 8$ had the lowest cross validation deviance. For standard PCA, we calculated the Bernoulli deviance using the reconstructed values as probability estimates. Since the reconstructed entries in the matrix could be outside the range of 0 to 1, we truncated them to be within the range $[10^{-10}, 1 - 10^{-10}]$.

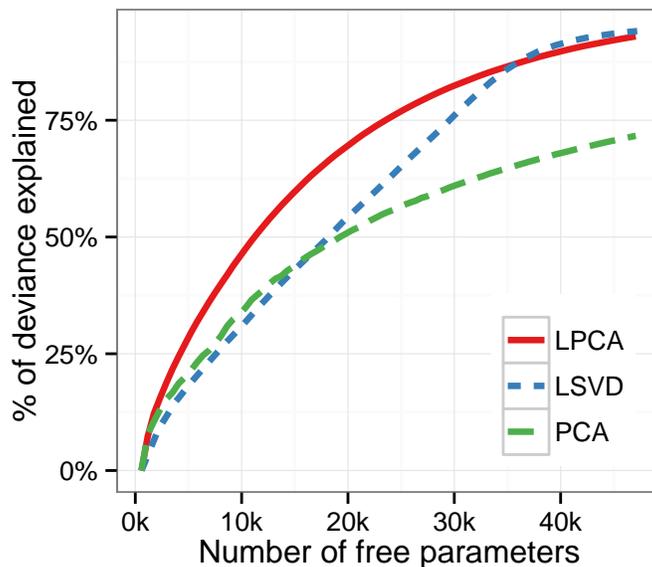


Figure 6: Cumulative percent of deviance explained by principal components of LPCA, LSVD, and PCA plotted against the number of free parameters for the patient-diagnosis data.

If we were using LSVD, one reasonable choice for the number of components would be 24 because that is where the cumulative percent of deviance explained begins to level off and the cumulative percent of deviance explained is quite high, at 90%. The marginal percent of deviance plot suggests that the marginal contributions level off after the second component for both LPCA and LSVD. To have a manageable number of components to analyze, we may use two components.

6.2.2 Quality of fit

For this dataset, it is clear that LPCA fits the data significantly better than PCA, even though they have the same number of parameters per component added. Also, LSVD has higher percent of deviance explained than LPCA after the first few components. However, it is not a completely fair comparison between LSVD and LPCA, because LSVD has extra parameters in \mathbf{A} to better fit the data. To illustrate this, we also plotted the percent of deviance explained as a function of the number of free parameters in Figure 6. When indexed by the number of free parameters, LPCA looks quite favorable compared to LSVD. In fact, there isn't a large difference between standard PCA and LSVD for smaller numbers of parameters.

To further show the advantages of LPCA, we used the loadings learned from this data to construct a low-rank estimate of the logit matrix using (3.4) for a different set of 1000 randomly selected patients. The percent of *predictive* deviance is plotted as a function of rank in Figure 7. Also plotted on this figure is the percent of predictive deviance by standard

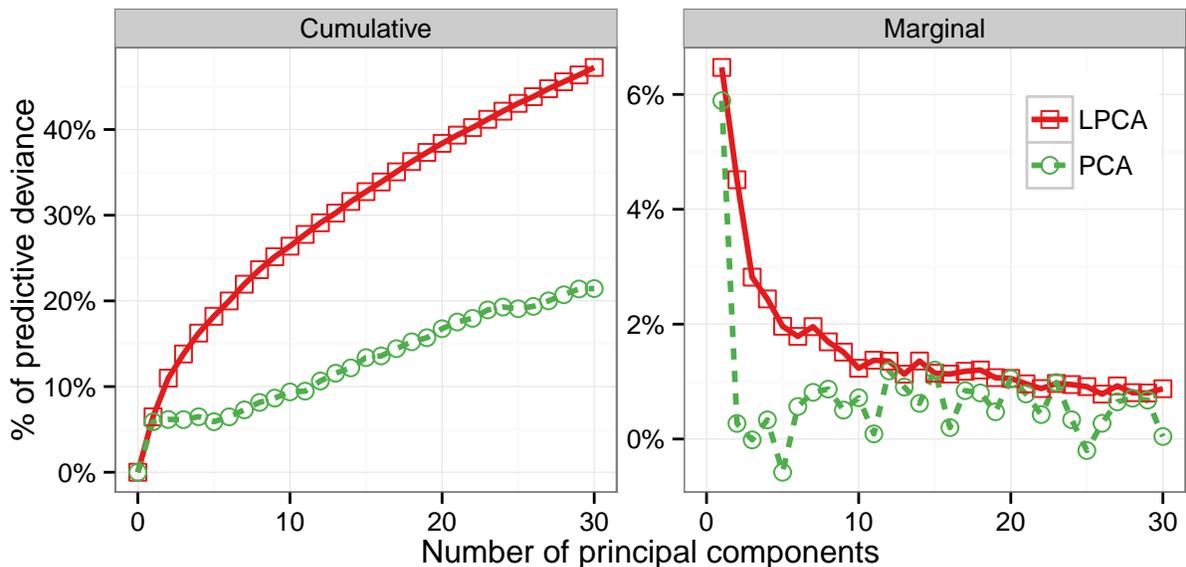


Figure 7: Cumulative and marginal percent of *predictive* deviance by the principal components of LPCA and PCA for the patient-diagnosis data.

PCA. We did not perform LSVD on this data since we were trying to recreate a scenario in which one wouldn't have the ability to solve for principal component scores on the new data. Just as in the explained deviance, the loadings learned from LPCA are able to reconstruct the new data much better than those learned from PCA.

6.2.3 Evidence of LSVD overfitting through principal component regression

To explore how LSVD might overfit, we created a simulation experiment using the same ICU data. We randomly generated a d -dimensional coefficient parameter, β , from a standard normal distribution. For each patient, we simulated a response variable, y_i , from a normal distribution with mean equal to $\mathbf{x}_i^T \beta$ and the variance, σ^2 . Three different variances were chosen in order for the signal-to-noise (SNR) ratio, $\text{var}(\mathbf{x}_i^T \beta) / \sigma^2$, to be either 1, 5, or 10. These represent weak, moderate, and strong signals, respectively.

For each k , we used the k principal components learned from PCA, LPCA, or LSVD as covariates in a linear regression. On a different set of 1000 randomly selected patients, we first derived the principal component scores using the appropriate procedure for each method and then predicted the response for the new patients. As stated in Section 3.3, obtaining principal component scores on new data with LSVD requires fitting a logistic regression for each new observation, while LPCA and PCA only require matrix multiplications.

The left column in Figure 8 shows the in-sample mean squared error (MSE) and the right column shows the out-of-sample MSE. The three rows of Figure 8 correspond to the three signal-to-noise ratios. For reference, we have also shown the intercept only model (labeled

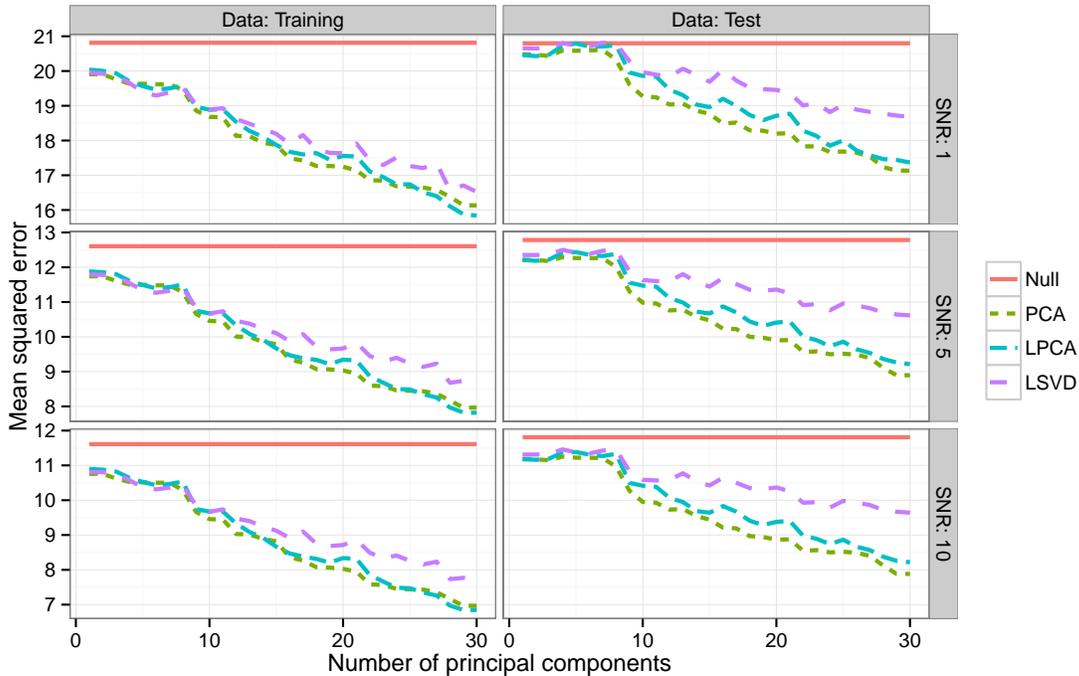


Figure 8: The in-sample and out-of-sample MSE of models for principal component regression based on the patient-diagnosis data

Null). For all three dimensionality reduction methods, the in-sample MSE decreases as the number of components increase, as we would expect. In general, PCA and LPCA track fairly close together. LSVD, on the other hand does markedly worse on both the training and test data, for all three signal levels. Further, the difference between the MSE from LPCA and LSVD is larger on the test data than on the training data. This simulation was run several times with similar results.

The point of this simulation is to show that, even though LSVD may explain more deviance than the other methods due to the extra parameters, that may also cause it to learn loadings that do not generalize well to other settings. This is very important for dimensionality reduction techniques, since they are not usually done in isolation, but rather are part of a larger data analysis. There are no guarantees that LPCA will generalize more effectively than LSVD in every setting, however we have shown in a controlled setting that LSVD produced loadings that do not generalize as well.

6.2.4 Interpretation of loadings

We look at the principal component loadings of LPCA with two components, as chosen by the scree plot, to attempt to interpret the comorbidity in patients. While there is no guarantee that the principal components will be interpretable, insight can occasionally be gained from the analysis. As mentioned in Section 3.4, the two components do not necessarily have an



Figure 9: The loadings of two components from logistic PCA of the patient-diagnosis data

order as in standard PCA, since they are solved for jointly. Also, it is not guaranteed that either of the two components analyzed here are represented in the analysis with one or three components.

The resulting loadings are plotted in Figure 9. We have highlighted the disease categories for the loadings with the highest absolute values. The points are also color-coded according to the broader category that the disease belongs to, which is also related to the number before the label.

A version of this plot was shown to subject-matter experts in the department of biomedical informatics at Ohio State University and they were able to identify meaning in the components (Hyun and Newton, 2013). The first component has high loadings for acute kidney failure and acute respiratory failure, among others, which are common serious conditions that cause a patient to be admitted into the ICU. Three of the large loadings from the second component are diseases of the circulatory system (green): systolic heart failure, myocardial infarction, aneurysm of coronary vessels, and hypertensive renal disease. It also includes two endocrine, nutritional and metabolic diseases, and immunity disorders (brown). The subject-matter experts stated that the diseases with high loadings were ones that they have observed co-occurring relatively often. Based on these findings, it seems that the prin-

principal components have a meaningful interpretation related to the diseases patients have at the ICU.

7 Discussion

Previous formulations of logistic PCA have extended the singular value decomposition to binary data. Our formulation more consistently extends PCA to binary data by finding projections of the natural parameters of the saturated model. Our method produces principal components which are linear combinations of the data and can be quickly calculated on new data. We have given two algorithms for minimizing the deviance and shown how they perform on both simulated and real data.

Further, the formulation proposed in this paper can be extended to other members of the exponential family. Using the appropriate deviance and natural parameters from the saturated model, the formulation can naturally be applied to many types of data.

When $d \gg n$, standard PCA can be inconsistent and it has been shown that adding sparsity constraints to PCA can induce consistency (Johnstone and Lu, 2009). Sparse loadings have the additional benefit of easier interpretation. Lee et al. (2010) extended LSVD by adding an L_1 penalty to the loadings matrix \mathbf{B} . Our formulation can be extended in the same way, but further research is needed to find the best way to solve for the loadings. However, it would be straightforward to incorporate the sparsity constraint with the Fantope solution, similar to how it was done in Vu et al. (2013) with standard PCA.

Acknowledgements

We would like to thank Vince Vu for his feedback, especially on the convex formulation and algorithm. We also thank Sookyung Hyun and Cheryl Newton at College of Nursing for providing the medical diagnoses data and valuable comments on a preliminary data analysis. This research was supported in part by National Science Foundation grant DMS-12-09194.

A Appendix

A.1 Calculation of gradient for logistic PCA

The gradient of the deviance in (3.2) with respect to \mathbf{U} can be seen from the steps below.

$$\begin{aligned} \frac{1}{2} \frac{\partial D}{\partial \mathbf{U}} &= - \frac{\partial}{\partial \mathbf{U}} \text{tr} \left(\mathbf{X}^T \left(\mathbf{1}_n \boldsymbol{\mu}^T + (\tilde{\boldsymbol{\Theta}} - \mathbf{1}_n \boldsymbol{\mu}^T) \mathbf{U} \mathbf{U}^T \right) \right) \\ &\quad + \frac{\partial}{\partial \mathbf{U}} \sum_{i,j} \log \left(1 + \exp(\mu_j + [\mathbf{U} \mathbf{U}^T (\tilde{\boldsymbol{\theta}}_i - \boldsymbol{\mu})]_j) \right) \end{aligned}$$

By standard matrix derivative rules (see, for example, Petersen and Pedersen (2012)),

$$\frac{\partial}{\partial \mathbf{U}} \text{tr}(\mathbf{X}^T(\tilde{\boldsymbol{\theta}} - \mathbf{1}_n \boldsymbol{\mu}^T) \mathbf{U} \mathbf{U}^T) = \left(\mathbf{X}^T(\tilde{\boldsymbol{\theta}} - \mathbf{1}_n \boldsymbol{\mu}^T) + (\tilde{\boldsymbol{\theta}} - \mathbf{1}_n \boldsymbol{\mu}^T)^T \mathbf{X} \right) \mathbf{U}.$$

To take the derivative of the second piece, letting $\hat{\theta}_{ij} = \mu_j + [\mathbf{U} \mathbf{U}^T(\tilde{\boldsymbol{\theta}}_i - \boldsymbol{\mu})]_j$, note that

$$\frac{\partial}{\partial u_{kl}} \sum_{i,j} \log \left(1 + \exp(\hat{\theta}_{ij}) \right) = \sum_{i,j} \frac{\exp(\hat{\theta}_{ij})}{1 + \exp(\hat{\theta}_{ij})} \frac{\partial \hat{\theta}_{ij}}{\partial u_{kl}} = \sum_{i,j} \hat{p}_{ij} \frac{\partial \hat{\theta}_{ij}}{\partial u_{kl}}.$$

Since

$$\frac{\partial [\mathbf{U} \mathbf{U}^T(\tilde{\boldsymbol{\theta}}_i - \boldsymbol{\mu})]_j}{\partial u_{kl}} = \begin{cases} (\tilde{\theta}_{ik} - \mu_k) u_{jl} & \text{if } k \neq j \\ (\tilde{\theta}_{ik} - \mu_k) u_{jl} + (\tilde{\boldsymbol{\theta}}_i - \boldsymbol{\mu})^T \mathbf{U}_l & \text{if } k = j, \end{cases}$$

that makes

$$\begin{aligned} \frac{\partial}{\partial u_{kl}} \sum_{i,j} \log \left(1 + \exp(\hat{\theta}_{ij}) \right) &= \sum_{i,j} \hat{p}_{ij} (\tilde{\theta}_{ik} - \mu_k) u_{jl} + \sum_i \hat{p}_{ik} (\tilde{\boldsymbol{\theta}}_i - \boldsymbol{\mu})^T \mathbf{U}_l \\ &= (\tilde{\boldsymbol{\theta}}_k - \mathbf{1}_n \mu_k)^T \hat{\mathbf{P}} \mathbf{U}_l + \hat{P}_k^T (\tilde{\boldsymbol{\theta}} - \mathbf{1}_n \boldsymbol{\mu}^T) \mathbf{U}_l. \end{aligned}$$

In matrix notation,

$$\frac{\partial}{\partial \mathbf{U}} \sum_{i,j} \log \left(1 + \exp(\mu_j + [\mathbf{U} \mathbf{U}^T(\tilde{\boldsymbol{\theta}}_i - \boldsymbol{\mu})]_j) \right) = \left(\hat{\mathbf{P}}^T (\tilde{\boldsymbol{\theta}} - \mathbf{1}_n \boldsymbol{\mu}^T) + (\tilde{\boldsymbol{\theta}} - \mathbf{1}_n \boldsymbol{\mu}^T)^T \hat{\mathbf{P}} \right) \mathbf{U},$$

and the result in Equation (4.1) follows.

The gradient of the deviance with respect to $\boldsymbol{\mu}$ in (4.2) is derived as follows.

$$\begin{aligned} \frac{1}{2} \frac{\partial D}{\partial \boldsymbol{\mu}} &= - \frac{\partial}{\partial \boldsymbol{\mu}} \text{tr} \left(\mathbf{X}^T \mathbf{1}_n \boldsymbol{\mu}^T (\mathbf{I}_d - \mathbf{U} \mathbf{U}^T) \right) \\ &\quad + \frac{\partial}{\partial \boldsymbol{\mu}} \sum_{i,j} \log \left(1 + \exp(\mu_j + [\mathbf{U} \mathbf{U}^T(\tilde{\boldsymbol{\theta}}_i - \boldsymbol{\mu})]_j) \right) \end{aligned}$$

Using standard vector differentiation,

$$\frac{\partial}{\partial \boldsymbol{\mu}} \text{tr} \left(\mathbf{X}^T \mathbf{1}_n \boldsymbol{\mu}^T (\mathbf{I}_d - \mathbf{U} \mathbf{U}^T) \right) = (\mathbf{I}_d - \mathbf{U} \mathbf{U}^T) \mathbf{X}^T \mathbf{1}_n$$

and

$$\frac{\partial}{\partial \boldsymbol{\mu}} \sum_{i,j} \log \left(1 + \exp(\mu_j + [\mathbf{U} \mathbf{U}^T(\tilde{\boldsymbol{\theta}}_i - \boldsymbol{\mu})]_j) \right) = \sum_{i,j} \hat{p}_{ij} (\mathbf{e}_j - \mathbf{u}_j^T \mathbf{U}^T) = (\mathbf{I}_d - \mathbf{U} \mathbf{U}^T) \hat{\mathbf{P}}^T \mathbf{1}_n,$$

where \mathbf{e}_j is a length d standard basis vector with 1 in the j th position.

A.2 Proof of Theorems

A.2.1 Theorem 1

Proof.

- (i) When $\mathbf{u} = \mathbf{e}_l$, the solutions for the main effects $\boldsymbol{\mu}$ are known analytically. For $j \neq l$, $\hat{\mu}_j = \text{logit}\bar{X}_j$. $\hat{\mu}_l$ is undefined because μ_l has no effect on the deviance. We will let $\hat{\mu}_l = \text{logit}\bar{X}_l$ for simplicity, although any constant would work. In this case, the estimated natural parameters are

$$\hat{\theta}_{ij} = \hat{\mu}_j + \delta_{jl}(\tilde{\boldsymbol{\theta}}_i - \hat{\boldsymbol{\mu}})^T \mathbf{e}_l = \begin{cases} \hat{\mu}_j & \text{if } j \neq l \\ \tilde{\theta}_{il} & \text{if } j = l \end{cases},$$

where δ_{jl} is the Kronecker delta.

Since \mathbf{e}_l is a standard basis vector,

$$[\mathbf{C}^m \mathbf{e}_l]_j = c_{jl}^m,$$

where

$$c_{jl}^m = (X_j - \hat{P}_j)^T (\tilde{\boldsymbol{\theta}}_l - \mathbf{1}_n \hat{\mu}_l) + (X_l - \hat{P}_j)^T (\tilde{\boldsymbol{\theta}}_j - \mathbf{1}_n \hat{\mu}_j). \quad (\text{A.1})$$

We will show that c_{jl}^m is equal to 0 when $j \neq l$ and $\bar{X}_l = \frac{1}{2}$. Looking at the first part of the summation in (A.1),

$$\begin{aligned} (X_j - \hat{P}_j)^T (\tilde{\boldsymbol{\theta}}_l - \mathbf{1}_n \hat{\mu}_l) &= (X_j - \mathbf{1}_n \sigma(\hat{\mu}_j))^T (mQ_l - \mathbf{1}_n \hat{\mu}_l) \\ &= (X_j - \mathbf{1}_n \bar{X}_j)^T (m(2X_l - \mathbf{1}_n) - \mathbf{1}_n \hat{\mu}_l) \\ &= m [2X_j^T X_l - n\bar{X}_j - n\bar{X}_j \hat{\mu}_l/m - 2n\bar{X}_j \bar{X}_l + n\bar{X}_j + n\bar{X}_j \hat{\mu}_l/m] \\ &= m [2X_j^T X_l - 2n\bar{X}_j \bar{X}_l] \\ &= m [2n\bar{X}_j \bar{X}_l - 2n\bar{X}_j \bar{X}_l] = 0. \end{aligned}$$

The last line is due to the assumption that X_j and X_l are uncorrelated.

From the fact that $(x_{il} - \sigma(mq_{il}))(mq_{ij} - \hat{\mu}_j) = (mq_{ij}q_{il} - q_{il}\hat{\mu}_j)/(1 + e^m)$, the second part of c_{jl}^m in (A.1) is

$$\begin{aligned} (X_l - \hat{P}_l)^T (\tilde{\boldsymbol{\theta}}_j - \mathbf{1}_n \hat{\mu}_j) &= (X_l - \sigma(mQ_l))^T (mQ_j - \mathbf{1}_n \hat{\mu}_j) \\ &= \sum_{i=1}^n (x_{il} - \sigma(mq_{il}))(mq_{ij} - \hat{\mu}_j) \\ &= \sum_{i=1}^n \frac{mq_{ij}q_{il} - q_{il}\hat{\mu}_j}{1 + e^m} \\ &= n\bar{Q}_l \frac{m\bar{Q}_j - \hat{\mu}_j}{1 + e^m}. \end{aligned}$$

If $\bar{Q}_l = 0$, or equivalently $\bar{X}_l = \frac{1}{2}$, then this is exactly zero, for all m .

When $j = l$, also using the fact that $q_{il}^2 = 1$,

$$\begin{aligned} c_{il}^m &= 2(X_l - \hat{P}_l)^T (\tilde{\Theta}_l - \mathbf{1}_n \hat{\mu}_l) \\ &= 2 \sum_{i=1}^n \frac{mq_{il}^2 - q_{il} \hat{\mu}_l}{1 + e^m} \\ &= 2n \frac{m - \hat{\mu}_l \bar{Q}_l}{1 + e^m}. \end{aligned}$$

When $\bar{X}_l = \frac{1}{2}$,

$$\mathbf{C}^m \mathbf{e}_l = \lambda_m \mathbf{e}_l,$$

for all m , where $\lambda_m = \frac{2nm}{1+e^m}$. With $\boldsymbol{\mu} = \hat{\boldsymbol{\mu}}$ and $\mathbf{u} = \mathbf{e}_l$, the first-order optimality conditions (4.1–4.3) are exactly satisfied.

(ii) When $\bar{X}_l \neq \frac{1}{2}$, the first part of (A.1) still equals zero, but the other part does not. The squared norm of equation (A.1) equals

$$\|\mathbf{C}^m \mathbf{e}_l\|^2 = \left(2n \frac{m - \hat{\mu}_l \bar{Q}_l}{1 + e^m}\right)^2 + \sum_{j:j \neq l} \left(n \bar{Q}_l \frac{m \bar{Q}_j - \hat{\mu}_j}{1 + e^m}\right)^2,$$

which can be made as small as we desire by increasing m . □

A.2.2 Theorem 2

Proof. If $\mathbf{u} = \mathbf{e}_l$ and $\hat{\mu}_j = \text{logit} \bar{X}_j$ for $j \neq l$, then the deviance of the l th column does not depend on the column mean \bar{X}_l and is given by

$$-2 \sum_{i=1}^n \log \sigma(q_{il} \tilde{\theta}_{il}) = -2n \log \sigma(m).$$

Thus, the deviance depends on the other $d - 1$ columns of the dataset. The deviance of the j th column ($j \neq l$) is

$$-2n (\bar{X}_j \log \bar{X}_j + (1 - \bar{X}_j) \log(1 - \bar{X}_j)),$$

which is maximized at $\bar{X}_j = \frac{1}{2}$ and decreases as \bar{X}_j moves away from $\frac{1}{2}$. Therefore, choosing l with the column mean closest to $\frac{1}{2}$ will result in a fit with the lowest deviance. □

A.2.3 Theorem 3

Proof. The first-order optimality condition for the loading vector with $k = 1$ and no main effects is

$$\mathbf{C}^m \mathbf{u} = \lambda_m \mathbf{u},$$

where

$$\mathbf{C}^m := (\mathbf{X} - \hat{\mathbf{P}})^T \tilde{\boldsymbol{\Theta}} + \tilde{\boldsymbol{\Theta}}^T (\mathbf{X} - \hat{\mathbf{P}}).$$

We will show that, with $\mathbf{u} = \frac{1}{\sqrt{d}} \mathbf{1}_d$ and the conditions listed in the theorem, \mathbf{C}^m is compound symmetric, which in turn implies that \mathbf{u} satisfies the first-order optimality conditions.

One useful implication of $\mathbf{u} \propto \mathbf{1}_d$ is that $\hat{p}_{ij} = \hat{p}_{ik}$ for all i and j, k because $\hat{\theta}_{ij} = u_j(\mathbf{u}^T \tilde{\boldsymbol{\theta}}_i) = \frac{m}{d} \sum_{l=1}^d q_{il}$. We will therefore let $\hat{\mathbf{p}}$ be the column vector with i th element equal to \hat{p}_{ij} for all j .

First, we show that $c_{jj}^m = c_{kk}^m$, for all j, k . $c_{jj}^m = c_{kk}^m$ if and only if $X_j^T Q_j - \hat{P}_j^T Q_j = X_k^T Q_k - \hat{P}_k^T Q_k$. This, in turn, is equivalent to

$$\hat{\mathbf{p}}^T (Q_j - Q_k) = \frac{1}{2} \mathbf{1}_n^T (Q_j - Q_k).$$

Focusing on the left hand side,

$$\begin{aligned} \hat{\mathbf{p}}^T (Q_j - Q_k) &= \sum_{i=1}^n \hat{p}_i (q_{ij} - q_{ik}) \\ &= \sum_{i: q_{ij} \neq q_{ik}} \hat{p}_i (q_{ij} - q_{ik}) \end{aligned}$$

The second line is due to the summation only being non-zero when $q_{ij} \neq q_{ik}$. When this is true, $\sum_{l=1}^d q_{il} = \sum_{l \notin \{j, k\}} q_{il}$. If $q_{ij} = q_{ik}$, $\hat{p}_i (q_{ij} - q_{ik})$ will equal 0 regardless of \hat{p}_i . Therefore, we can state

$$\sum_{i: q_{ij} \neq q_{ik}} \hat{p}_i (q_{ij} - q_{ik}) = \sum_{i=1}^n \sigma \left(\frac{m}{d} \sum_{l \notin \{j, k\}} q_{il} \right) (q_{ij} - q_{ik}),$$

and from (4.4),

$$\begin{aligned} \hat{\mathbf{p}}^T (Q_j - Q_k) &= \left(\frac{1}{2} \mathbf{1}_n + \sum_{l \notin \{j, k\}} Q_l \beta_{jk, l} \right)^T (Q_j - Q_k) \\ &= \frac{1}{2} \mathbf{1}_n^T (Q_j - Q_k). \end{aligned}$$

The last line is due to the compound symmetry of $\mathbf{Q}^T \mathbf{Q}$, since all the off-diagonal elements are equal to each other. This proves that $c_{jj}^m = c_{kk}^m$.

We will now show that $c_{jk}^m = c_{lr}^m$, as long as $j \neq k$ and $l \neq r$. An off-diagonal element of \mathbf{C}^m is

$$\begin{aligned} c_{jk}^m &= m(X_j - \hat{\mathbf{p}})^T Q_k + m(X_k - \hat{\mathbf{p}})^T Q_j \\ &= m \left[Q_j^T Q_k + \frac{1}{2} \mathbf{1}_n^T (Q_j + Q_k) - \hat{\mathbf{p}}^T (Q_j + Q_k) \right]. \end{aligned}$$

Showing $c_{jk}^m = c_{lr}^m$ is equivalent to showing

$$\begin{aligned} Q_j^T Q_k + \frac{1}{2} \mathbf{1}_n^T (Q_j + Q_k) - \hat{\mathbf{p}}^T (Q_j + Q_k) &= Q_l^T Q_r + \frac{1}{2} \mathbf{1}_n^T (Q_l + Q_r) - \hat{\mathbf{p}}^T (Q_l + Q_r) \\ \frac{1}{2} \mathbf{1}_n^T (Q_j + Q_k) - \hat{\mathbf{p}}^T (Q_j + Q_k) &= \frac{1}{2} \mathbf{1}_n^T (Q_l + Q_r) - \hat{\mathbf{p}}^T (Q_l + Q_r), \end{aligned}$$

where the terms cancel because $Q_j^T Q_k = Q_l^T Q_r$. Rearranging the terms,

$$\hat{\mathbf{p}}^T (Q_j - Q_r) - \frac{1}{2} \mathbf{1}_n^T (Q_j - Q_r) = \hat{\mathbf{p}}^T (Q_l - Q_k) - \frac{1}{2} \mathbf{1}_n^T (Q_l - Q_k),$$

where we can see that both sides equal 0 because, as we have already proven, $c_{jj}^m = c_{rr}^m$ and $c_{ll}^m = c_{kk}^m$.

Therefore, \mathbf{C}^m is compound symmetric for all m . This implies that $\mathbf{u} = \frac{1}{\sqrt{d}} \mathbf{1}_d$ is an eigenvector of \mathbf{C}^m and (in conjunction with the fact that $\mathbf{u}^T \mathbf{u} = 1$) \mathbf{u} satisfies the first-order optimality conditions (4.1, 4.3). \square

A.3 Further simulation results

A.3.1 Number of iterations to convergence

For the same simulation setup as in Figure 3, we kept track of the number of iterations required until convergence for both the LSVD algorithm and the LPCA algorithm. The initial value of \mathbf{U} or \mathbf{B} was the first k right singular vectors of \mathbf{Q} for both. Each iteration for the two algorithms is comparable in that they involve computing the working variables of the majorization function and finding either an SVD or an eigen-decomposition. Both of the algorithms terminate if the difference of the average deviance between successive iterations is less than 10^{-5} or if the maximum number of iterations (1000) is reached.

It is apparent from Figure 10 that the number of iterations required for convergence is a lot less for LPCA than for LSVD, in general, and in many cases LSVD reaches the maximum. (Note that the y-axis is on the log scale.) For $\phi \in \{1, 3\}$, the number of iterations for LSVD increases as the number of estimated components increases. Also, as m increases, the number of iterations for LPCA increases, in general.

We believe that LSVD requires more iterations to converge because there are more parameters for it to optimize for, so there are many little updates that can be made to the parameters which improve the deviance enough. One explanation for why it takes LPCA with large m longer to converge is that small changes in \mathbf{U} can have a relatively large effect in the estimated natural parameters with large m .

A.3.2 Simulation results for more sparse situation

Just like in Section 6.1.3, we simulated from a variety of scenarios, only this time $\bar{p} = 0.1$, for a situation where the binary matrix consists of mostly 0's. The estimated LSVD and LPCA included main effects. The results in Figure 11 are mostly the same as with $\bar{p} = 0.5$

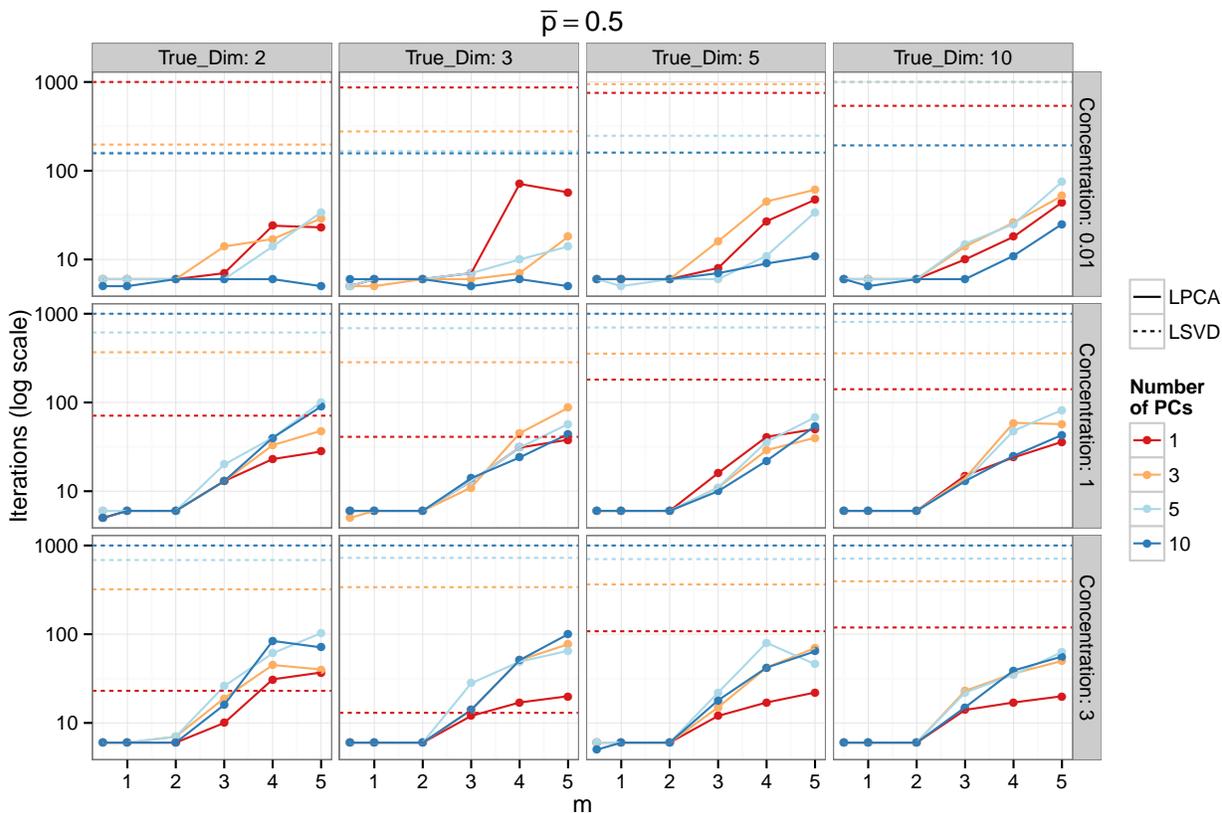


Figure 10: Number of iterations required for convergence of the logistic PCA and logistic SVD algorithms for a variety of scenarios

except that an estimate of the probability matrix with a rank lower than the true rank does not degrade performance as badly.

A.3.3 Deviance comparison of LPCA and LSVD

Finally, under the same setup as before with $\bar{p} = 0.5$, Figure 12 shows the deviance of the estimates. As expected, LSVD has lower deviance or equivalent deviance for all scenarios. The improvement from LSVD is smallest when $\phi = 0.01$. When the probabilities are more concentrated, however, LSVD clearly has a lower deviance than LPCA. It also appears that the higher the estimated rank is, the greater the decrease in the deviance is from LSVD.

References

Boyd, S., L. Xiao, and A. Mutapcic (2003). Subgradient methods. *Lecture notes of EE392o, Stanford University, Autumn Quarter 2004*.

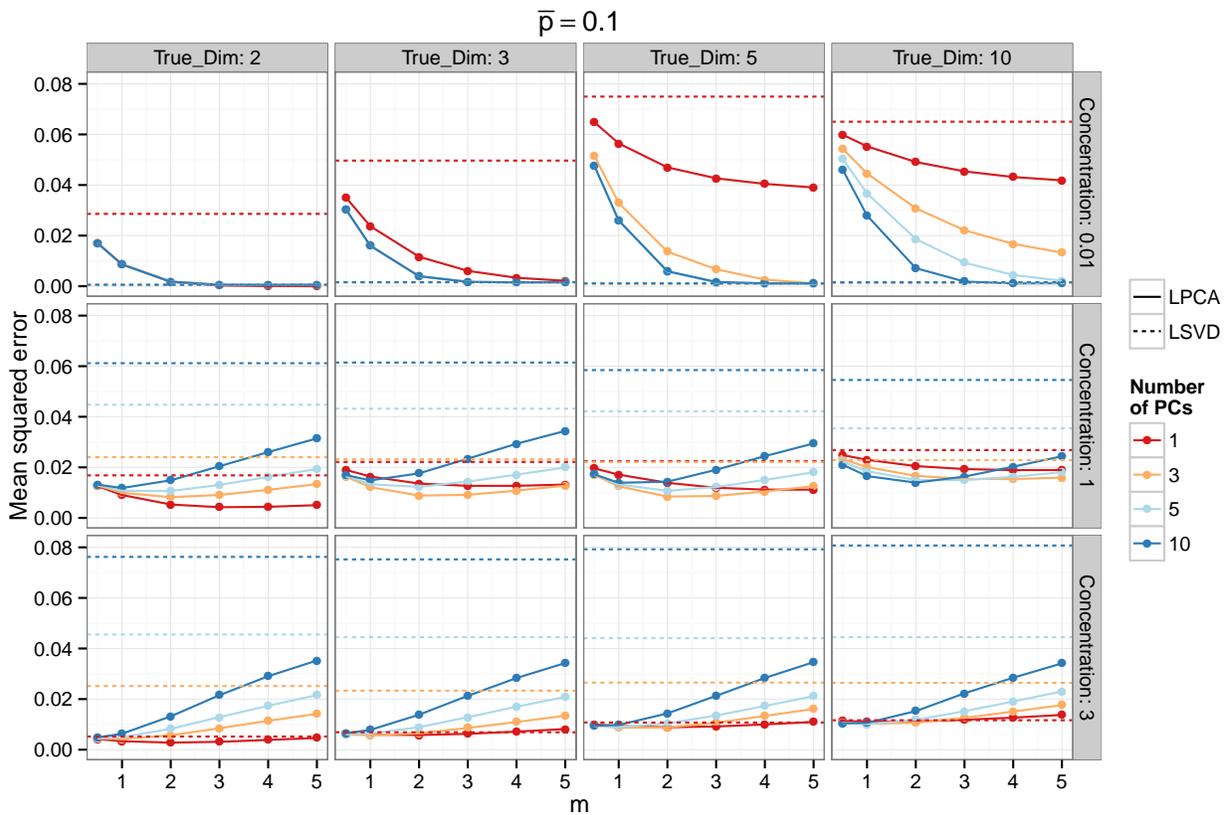


Figure 11: Mean squared error of probability estimates when $\bar{p} = 0.1$ for a variety of scenarios

Collins, M., S. Dasgupta, and R. E. Schapire (2001). A generalization of principal components analysis to the exponential family. In T. Dietterich, S. Becker, and Z. Ghahramani (Eds.), *Advances in Neural Information Processing Systems 14*, pp. 617–624. MIT Press.

Dattorro, J. (2005). *Convex optimization & Euclidean distance geometry*. Meboo Publishing USA.

de Leeuw, J. (2006). Principal component analysis of binary data by iterated singular value decomposition. *Computational Statistics and Data Analysis* 50(1), 21–39.

Eckart, C. and G. Young (1936). The approximation of one matrix by another of lower rank. *Psychometrika* 1(3), 211–218.

Fan, K. (1949). On a theorem of Weyl concerning eigenvalues of linear transformations I. *Proceedings of the National Academy of Sciences of the United States of America* 35(11), 652–655.

Hastie, T., R. Tibshirani, and J. Friedman (2009). *The Elements of Statistical Learning*, Volume 2. Springer.

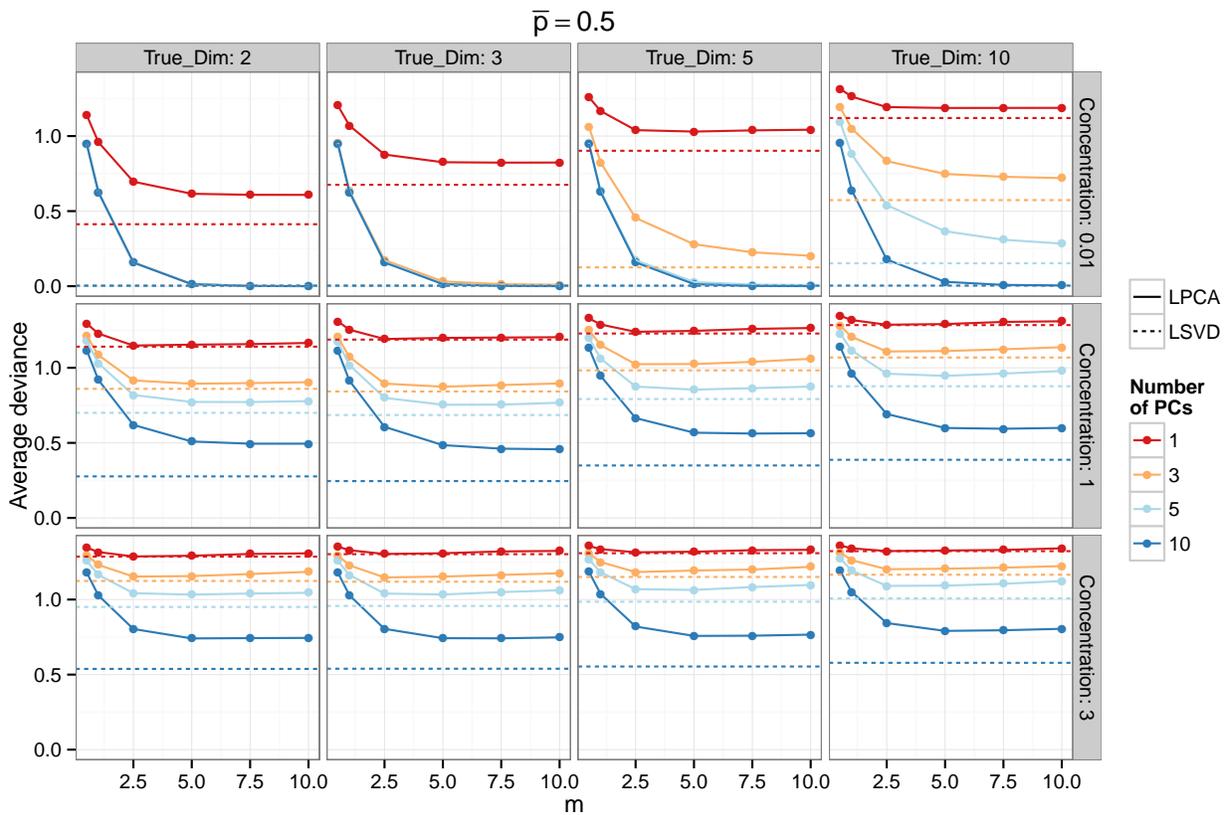


Figure 12: Average deviance of probability estimates when $\bar{p} = 0.5$ for a variety of scenarios

Hotelling, H. (1933). Analysis of a complex of statistical variables into principal components. *Journal of Educational Psychology* 24(6), 417–441.

Hunter, D. R. and K. Lange (2004). A tutorial on MM algorithms. *The American Statistician* 58(1), 30–37.

Hyun, S. and C. Newton (2013, December). Personal communication.

Johnson, C. C. (2014). Logistic matrix factorization for implicit feedback data. In *Advances in Neural Information Processing Systems 27: Distributed Machine Learning and Matrix Computations Workshop*.

Johnstone, I. M. and A. Y. Lu (2009). On consistency and sparsity for principal components analysis in high dimensions. *Journal of the American Statistical Association* 104(486), 682–693.

Jolliffe, I. (2002). *Principal component analysis*. Springer.

Lange, K. (2013). *Optimization*. Springer.

- Lee, S., J. Z. Huang, and J. Hu (2010). Sparse logistic principal components analysis for binary data. *The Annals of Applied Statistics* 4(3), 1579–1601.
- Li, J. and D. Tao (2010). Simple exponential family PCA. In *International Conference on Artificial Intelligence and Statistics*, pp. 453–460.
- McCullagh, P. and J. A. Nelder (1989). *Generalized linear models (Second edition)*. London: Chapman & Hall.
- Pearson, K. (1901). On lines and planes of closest fit to systems of points in space. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* 2(11), 559–572.
- Petersen, K. B. and M. S. Pedersen (2012, November). The matrix cookbook. Version 20121115.
- R Core Team (2015). *R: A Language and Environment for Statistical Computing*. Vienna, Austria: R Foundation for Statistical Computing.
- Schein, A. I., L. K. Saul, and L. H. Ungar (2003). A generalized linear model for principal component analysis of binary data. In *Proceedings of the Ninth International Workshop on Artificial Intelligence and Statistics*, Volume 38.
- Tipping, M. E. (1998). Probabilistic visualisation of high-dimensional binary data. In M. Kearns, S. Solla, and D. Cohn (Eds.), *Advances in Neural Information Processing Systems 11*, pp. 592–598. MIT Press.
- Tipping, M. E. and C. M. Bishop (1999). Probabilistic principal component analysis. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 61(3), 611–622.
- Vu, V. Q., J. Cho, J. Lei, and K. Rohe (2013). Fantope projection and selection: A near-optimal convex relaxation of sparse pca. In C. Burges, L. Bottou, M. Welling, Z. Ghahramani, and K. Weinberger (Eds.), *Advances in Neural Information Processing Systems 26*, pp. 2670–2678. Curran Associates, Inc.
- Welling, M., C. Chemudugunta, and N. Sutter (2008). Deterministic latent variable models and their pitfalls. In *SIAM International Conference on Data Mining*, pp. 196–207.
- Wen, Z. and W. Yin (2013). A feasible method for optimization with orthogonality constraints. *Mathematical Programming* 142(1-2), 397–434.