# Dimension Reduction for Exponential Family Data with Applications to Text Data 

Luke Smallman

2019

Submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy

School of Mathematics
Ysgol Mathemateg

## Summary

In this manuscript, we will address the problem of dimension reduction for data modelled by an exponential family distribution, with a particular focus on text data modelled by a Poisson-count model. We are motivated to develop new methods for such data by links between principal component analysis and the Gaussian log-likelihood, which suggests both a simple way to extend PCA to the exponential family (of which the Gaussian distribution is a member), and the unsuitability of PCA when the data is appropriately modelled by a distribution which is not well-approximated by the Gaussian distribution.

We will present three novel methods for exponential family dimension reduction. The first is "Poisson Inverse Regression", a supervised method from the family of inverse regression methods. We will demonstrate that this method provides a sufficient dimension reduction. That is, the transformed data is statistically sufficient with respect to the response.

The second is Sparse Generalised Principal Component Analysis, which extends the method of Generalised Principal Component Analysis put forward by Landgraf and Lee (2015b). This method is unsupervised, as is motivated by a modification of the PCA objective function to accommodate other exponential family distributions. We demonstrate that this method performs as-well or better than other state-of-the-art methods. This work has been published as Smallman, Artemiou, et al. (2018).

The third is Sparse Simple Exponential/Poisson Principal Component Analysis. This method extends Simple Exponential Principal Component Analysis, put forward by Li and Tao (2013), enforcing sparsity in the equivalent of the loadings matrix. This method is also unsupervised, and we demonstrate its state-of-the-art performance. This work was done jointly with William Underwood from Oxford University, and is published in Smallman, Underwood, et al. (2019).

Finally, we present a new framework for analysing and synthesising dimension reduction methods, which we call "Quasi-Likelihood PCA". This is based on tensor estimating equations, which we also present as a new development. We apply this method to analyse several methods in the literature.

This would not have been possible without the love and support of so many. In particular, my deepest thanks to Scott - you have been there from the beginning, have put up with so much, and have never once doubted me.

Thank you to my parents; without your help and support I would never have been in a position to begin this, let alone finish it.

Thank you to Lynn, Glenni and Mizzy, your friendship has been invaluable.

Finally, but by no means least, thank you to Andreas - you've been a wonderful mentor and friend throughout this process.

[^0]
## Acknowledgements

I would like to acknowledge deep gratitude to Cardiff and Vale University Health Board and Cardiff University School of Mathematics for providing funding for this project.

The work of Chapter 5 was done jointly with William Underwood from Oxford University as part of a 4 week undergraduate summer project. I would like to thank William for his collaboration and enthusiasm; it was a pleasure working with him.

## Contents

Contents ..... ix
List of Figures ..... xiii
List of Tables ..... xV
1 Intro ..... 1
1.1 Dimension Reduction ..... 1
1.1.1 Sufficient Dimension Reduction ..... 2
1.1.2 Sliced Inverse Regression ..... 3
1.1.3 Principal Component Analysis ..... 3
1.2 Text Data ..... 5
1.3 Sparsity ..... 7
1.4 Exponential Families ..... 7
1.5 Related Work ..... 9
1.5.1 Distributional Dimension Reduction ..... 9
1.5.1.1 Probablistic PCA ..... 9
1.5.1.2 Bayesian PCA ..... 10
1.5.1.3 Collins et al. 2002 ..... 11
1.5.1.4 Bayesian Exponential PCA ..... 12
1.5.1.5 Simple Exponential Principal Component Analysis ..... 13
1.5.1.6 Generalised Principal Component Analysis ..... 14
1.5.1.7 Sparse Probabilistic Principal Component Analysis ..... 14
1.5.1.8 Sparse Exponential Family Principal Component Analysis ..... 14
1.5.1.9 Multinomial Inverse Regression ..... 15
1.5.1.10 Latent Dirichlet Allocation ..... 16
1.5.2 Non-Distributional Dimension Reduction ..... 17
1.5.2.1 Sparse Principal Component Analysis ..... 17
1.5.2.2 Joint Sparse Principal Component Analysis ..... 18
1.5.2.3 Robust Principal Component Analysis ..... 18
1.5.2.4 Sparse Principal Component Analysis by Rotation and Truncation ..... 19
1.5.2.5 Non-negative Matrix Factorisation ..... 20
1.6 Conclusion ..... 20
2 Poisson Inverse Regression ..... 21
2.1 Introduction ..... 21
2.2 Poisson Inverse Regression ..... 21
2.3 Estimation ..... 24
2.4 Evaluation ..... 27
2.5 Conclusion ..... 30
3 Healthcare Data ..... 33
4 Sparse Generalised Principal Component Analysis ..... 35
4.1 Introduction ..... 35
4.1.1 Text Data ..... 35
4.2 GPCA ..... 36
4.2.1 GPCA Definition ..... 36
4.3 SGPCA Definition ..... 37
4.3.1 Penalisation ..... 37
4.3.1.1 SCAD Penalty ..... 38
4.3.1.2 $\quad L^{1}$ Penalty ..... 40
4.3.1.3 Total Penalty Function ..... 40
4.3.1.4 Definition of SGPCA ..... 40
4.3.2 Estimation ..... 40
4.4 Synthetic Data Examples ..... 45
4.4.1 "Classless" Data Exploration ..... 45
4.4.2 Classed Synthetic Data Exploration ..... 46
4.4.3 Robustness Against Noise ..... 49
4.5 Dependence on Tolerance ..... 51
4.6 Healthcare Data ..... 51
4.7 Discussion ..... 52
5 Sparse Simple Exponential Family Principal Component Analysis ..... 55
5.1 SePCA ..... 55
5.2 Sparse Simple Exponential Principal Component Analysis ..... 56
5.3 Estimation ..... 57
5.4 Synthetic Data Studies ..... 59
5.4.1 Order Determination ..... 61
5.4.2 Synthetic Data with Classes ..... 61
5.5 Healthcare Data ..... 65
5.6 Discussion ..... 67
6 Quasi-Likelihood Principal Component Analysis ..... 69
6.1 Introduction ..... 69
6.2 Tensor Estimating Equations ..... 70
6.2.1 Vector Parameter Estimating Equations ..... 70
6.2.2 Tensor Preliminaries ..... 71
6.2.3 Vector Estimating Equations as Tensors ..... 73
6.2.4 Tensor Estimating Equations ..... 75
6.2.5 Asymptotic Consistency ..... 77
6.3 Categorising Generalisations of PCA ..... 80
6.3.1 Generalised PCA ..... 80
6.3.2 Collins et al. ..... 80
6.3.3 Simple Exponential PCA ..... 81
6.4 Comparisons ..... 81
6.5 Asymptotic Consistency ..... 82
6.6 Conclusions ..... 83
7 Conclusions ..... 85
Bibliography ..... 87

## List of Figures

1.1 PCA direction and projections for two-dimensional data ..... 4
1.2 Plate diagram for Bayesian PCA ..... 11
1.3 Plate diagram for Bayesian Exponential Family PCA ..... 13
1.4 Plate diagram for Latent Dirichlet Allocation ..... 17
2.1 Plot of two directions recovered by POIR from the we8there data. ..... 28
2.2 Plot of three directions recovered by POIR from the we8there data. ..... 29
2.3 Scatterplot of the one MNIR direction recovered from the we8there data against the rating, and density of that direction separated by rating. ..... 30
4.1 SCAD Penalty ..... 39
4.2 SCAD Penalty for different $\lambda$ ..... 39
4.3 Behaviour of all SGPCA variants as tolerance is varied. ..... 52
4.4 Plots of pairs of the first three principal components for the healthcare dataset obtained from SGPCA, GPCA and PCA. Red "+" symbols denote the "dis- charge" class; black dots represent the "follow-up" class. ..... 53
5.1 Plate diagram for SePCA ..... 56
5.2 Two directions from each algorithm for X2C, with one class shown with red squares, the other with black triangles. ..... 65
5.3 Two directions from each algorithm for X3C ..... 66
5.4 The resulting principal components from applying SPPCA, SSPPCA, GPCA and SGPCA to the healthcare data. The "discharge" class is plotted as red squares, the "follow-up" class is shown as black triangles. ..... 67

## List of Tables

1.1 Univariate EF distributions ..... 8
1.2 SPPCA prior-penalty equivalences ..... 15
4.1 Synthetic Loading 1 ..... 47
4.2 Synthetic Loading 2 ..... 48
4.3 Classed Synthetic Loadings ..... 49
4.4 Investigations of the performance of all three SGPCA variants across varying levels of noise. ..... 50
5.1 Loadings for X1D ..... 62
5.2 Two loadings from X2D ..... 63
5.3 Percentage of correct identification of $d$ for SPPCA and SSPPCA ..... 63
5.4 Average (Euclidean) silhouettes ..... 65
5.5 Average silhouettes first the healthcare data. ..... 68

## — Chapter 1 -

## Introduction

With the explosion of "big data", methods for dimension reduction have become increasingly important; recently, there has been considerable research into different ways to perform dimension reduction for a wide variety of types of data. Text data is one of those types of data - it is plentiful, its analysis is frequently very impactful, and it has certain statistical peculiarities. In particular, text data is not Gaussian nor even symmetrically distributed which negatively impacts the performance of many of the standard methods for dimension reduction. In this work, we will develop a Poisson-model-based method for text dimension reduction in Chapter 2, sparse extensions of two exponential family methods of dimension reduction in Chapter 4 and Chapter 5, and finally a framework for constructing, comparing and deriving asymptotic results for a class of estimators known as tensor estimating equations in Chapter 6 which we use to analyse several exponential family dimension reduction methods.

In this chapter, we will begin by introducing the fundamental notions of dimension reduction in Section 1.1 and text data in Section 1.2. We then motivate the desire to include a notion of sparsity in our dimension reduction techniques in Section 1.3. For our focus on text data, we will be working extensively with the exponential family of distributions, which we introduce in Section 1.4. Finally, we examine related work in the literature in Section 1.5. The methods in this section will be divided into those with distributional assumptions in Section 1.5.1 and those without in Section 1.5.2.

### 1.1 Dimension Reduction

Dimension reduction methods can broadly be divided into supervised and unsupervised methods; i.e. those which take into account the values of some "response" variable, and those which do not. For the former, we will take as our canonical example the method of sliced inverse regression (SIR), which is within the family of "sufficient dimension reduction" techniques. For the latter, we will use Principal Component Analysis (PCA).

Before exploring the details of SIR, we will take a detour into the general idea of sufficient dimension reduction, as this will provide context both for SIR and for the method of Poisson Inverse Regression (PoIR) in Chapter 2.

### 1.1.1 Sufficient Dimension Reduction

The family of sufficient dimension reduction methods is a family of supervised dimension reduction techniques with the explicit aim of finding a projection of the original data into a lower-dimensional space such that the projected data is statistically sufficient. Formally, we give the definition:

Definition 1.1.1. Let $\gamma: \mathscr{X} \rightarrow \mathscr{Z}$, where $\mathscr{X}$ is the domain of a random vector $\mathcal{X}$ and $\mathscr{Z}$ is the domain of the random vector $\gamma(\mathcal{X})$ satisfying $|\mathscr{X}|<|\mathscr{Z}|$. Then $\gamma$ is a sufficient dimension reduction mapping with respect to a random variable $y$ if

$$
y \Perp X \mid \gamma(X)
$$

There is a strong advantage to this family in the application of predictive influence; by preserving all information that the observed data has about the response with our dimension reduction, we are able to perform the same quality of predictive influence with the lower-dimensional data.

Given this definition, it can be seen that such transformations are not unique; if $\gamma$ is a sufficient dimension reduction, then so is $a \gamma$, for any $a \in \mathbb{R}$ $\{0\}$. In the more restricted set of linear sufficient dimension reductions (which are projection matrices), we can define the notion of a "dimension reduction subspace". This encapsulates all full-rank linear transformations of the dimension reduction, bringing us closer to identifiability.

Definition 1.1.2. A "dimension reduction subspace" is the columnspace of a sufficient linear dimension reduction matrix.

A less obvious problem still remains; a suitably "small" sufficient dimension reduction can be made larger by adding superfluous components. This leads us to question if there is a "smallest" sufficient dimension reduction. This is addressed by the concept of a "central dimension reduction subspace", defined below.

Definition 1.1.3. The "central dimension reduction subspace" is the intersection of all dimension reduction subspaces.

If the central subspace exists, then it is a dimension reduction subspace with the minimal needed dimension. To see this, note that if a lower dimensional dimension
reduction subspace exists, then it is necessarily included in the intersection forming the central dimension reduction subspace, thus the central space cannot be of larger dimension. The dimension of the central subspace is often referred to as the "structural" dimension, as the central subspace is the minimal needed structure for the data (with respect to the response). Estimating this central subspace is the goal of a number of dimension reduction methods, notably Sliced Inverse Regression (Li 1991) which popularised the idea of sufficient dimension reduction and which we will now explore.

### 1.1.2 Sliced Inverse Regression

Sliced inverse regression Li (1991) begins with the assumption that the relationship between the predictors $\mathcal{X}$ and the response $y$ has the form

$$
\begin{equation*}
y=f\left(\beta_{1} \mathcal{X}, \ldots, \beta_{k} \mathcal{X}, \boldsymbol{\epsilon}\right) \tag{1.1}
\end{equation*}
$$

where $\varepsilon$ represents noise. They also the additional restriction that

$$
\begin{equation*}
\mathbb{E}\left(\boldsymbol{X} \mid \beta_{1} \mathcal{X}, \ldots, \boldsymbol{\beta}_{k} \mathcal{X}\right)=c_{0}+c_{1} \boldsymbol{\beta} \mathcal{X}+\ldots+c_{k} \boldsymbol{\beta}_{k} \mathcal{X} \tag{1.2}
\end{equation*}
$$

for all $\beta \in \mathbb{R}^{p}$ with $c_{0}, \ldots, c_{k}$ real constants. This assumption is often referred to in the literature as the "linear conditional expectation" condition; it was shown in Eaton (1986) that this is condition is satisfied if and only if $\mathcal{X}$ has an elliptically symmetric distribution. The central result of sliced inverse regression is then the following theorem.

Theorem 1.1.1. Under (1.1) and (1.2), the centred inverse regression curve $\mathbb{E}(X \mid y)$ $\mathbb{E}(X)$ lies within the central dimension reduction subspace.

Remark. A brief history note: the notion of dimension reduction subspaces and the central dimension reduction subspace was developed after the publication of sliced inverse regression. As such, Theorem 1.1.1 is stated differently in its original published form.

Remark. The word "sliced" in sliced inverse regression refers to part of the estimation procedure, in which the response $y$ is sliced into $H$ bands and the sample mean of the observed predictors (after standardisation) is calculated within each slice. These sample conditional means are then formed into a covariance matrix, weighted by the proportion of observations in each slice. Finally, the directions $\boldsymbol{\beta}_{1}, \ldots, \boldsymbol{\beta}_{k}$ are then estimated from the eigenvectors with the largest eigenvalues.

### 1.1.3 Principal Component Analysis

PCA was first formulated in Pearson (1901) and (perhaps more importantly) reformulated in Hotelling (1933). It is often formulated in terms of finding a successive series
of orthogonal directions which maximise variance in turn - the first direction is in the direction of greatest variance, the second direction in the direction of greatest variance which is orthogonal to the first, and so on. However, there is an equivalent definition which is more suited to our purposes. Formally, we define the PCA loadings matrix $\mathbf{U}$ (and the associated mean vector $\mu$ ) by

$$
\begin{equation*}
\mathbf{U}, \mu:=\underset{\mathbf{U} \in \mathbb{R}^{p \times k}, \mathbf{U}^{\mathrm{T}} \mathbf{U}=\mathbf{I}, \mu \in \mathbb{R}^{p}}{\arg \min } \sum_{i=1}^{n}\left\|\mathbf{x}_{i}-\mu-\mathbf{U U}^{\mathrm{T}}\left(\mathbf{x}_{i}-\mu\right)\right\|_{2}^{2} \tag{1.3}
\end{equation*}
$$

where $\mathbf{x}_{i} \in \mathbb{R}^{p}$ for $i=1, \ldots, n$ are our observed data. This form was considered by Pearson for the optimal projection of the $p$-dimensional data into a $k$ dimensional subspace under squared error loss. In this formulation, the optimal choice for $\mu$ is $\overline{\mathbf{x}}$ (the sample mean), and for $\mathbf{U}$ the first $k$ eigenvectors of the sample covariance matrix ordered descending by eigenvalue. Later, we will relate this form to the deviance of a particular Gaussian model for the data, and use that connection to extend PCA to all exponential family distributions. For now, we show in Figure 1.1 the result of applying PCA to some two-dimensional data (shown as black points). The line shows the first direction found by PCA, and the red points on the line are the projections of the data onto the line.


Figure 1.1: PCA direction and projections for two-dimensional data

It is worth noting that (1.3) is not the "typical" definition. The more usual definition is as the directions of maximal variance, which can be derived from the variance of the
standardised data. To be precise, we can define the first principal component by the weight vector $\mathbf{w}_{1}$ such that

$$
\begin{equation*}
\mathbf{w}_{1}=\underset{\|w\|=1}{\arg \max } \sum_{i=1}^{n}\left(\mathbf{w} \cdot \mathbf{z}_{i}\right)^{2}=\underset{\|w\|=1}{\arg \max } \mathbf{w}^{\mathrm{T}} \mathbf{Z}^{\mathrm{T}} \mathbf{Z} \mathbf{w} \tag{1.4}
\end{equation*}
$$

where $\mathbf{Z}$ is the matrix with each row an observation $\mathbf{z}_{i}^{\mathrm{T}}$ and each $v e c z_{i}$ is defined by $\mathbf{z}_{i}=\mathbf{x}_{i}-\frac{1}{n} \sum_{i=1}^{n} \mathbf{x}_{i}$. It can then be shown that the maximal value is the largest eigenvalue of $\mathbf{Z}^{\mathrm{T}} \mathbf{Z}$ which is obtained at the corresponding eigenvector.

The $k^{\text {th }}$ component is then defined as the first principal component of the matrix

$$
\hat{\mathbf{z}}_{k}=\mathbf{Z}-\sum_{i=1}^{k-1} \mathbf{Z w}_{i} \mathbf{w}_{i}^{\mathrm{T}}
$$

and it can be further shown that this is precisely the $k^{\text {th }}$ eigenvector of $\mathbf{Z}^{\mathrm{T}} \mathbf{Z}$ when sorted by decreasing eigenvalue.
Remark. It is worth noting that the usual computational method for calculating the principal components is via singular value decomposition of $\mathbf{Z}$, as this avoids the need to calculate $\mathbf{Z}^{\mathrm{T}} \mathbf{Z}$ which can be computationally expensive when the number of observations and/or dimension of them is high.

Remark. It is also important to note that the principal components are not generally unique. Trivially, one can note that if $\mathbf{w}$ is a principal component, then $-\mathbf{w}$ could replace it.

In general, dimension reduction methods are transformations from a high-dimensional feature space to a lower-dimensional feature space. We require an important restriction, however: these transformations must have some form of data-driven optimality. In the case of PCA, this is the minimised squared reconstruction error (or the maximised successive variances). Here, the reconstruction error means the $L^{2}$ distance between the original data and the reconstruction from the lower dimensional approximation. Our hope is that given data and a dimension reduction method with some appropriate measure of optimality, we can find a lower-dimensional representation of the data which is suitable for further statistical analysis or methodology which would have been infeasible or unsatisfactory with the original data.

### 1.2 Text Data

The majority of this work (specifically Chapters 2, 4 and 4) will be motivated and demonstrated using text data, under a Poisson counts model. For edification, we will now intro-
duce some key information about text data and how it can be studied from a statistical perspective.

We will use, as is standard in the literature and especially in practice, the vector space model" of text data, where documents are represented by vectors in $\mathbb{N}^{d}$. Each element of the vector corresponds to a stem, with the value being the number of times that stem appears in the document. A stem is effectively a word reduced down to an immutable part which does not change under conjugation, pluralisation, etc. Before vector representation, documents undergo the process of stemming, where each word is reduced to the appropriate stem. For example, the words "go", "going" and "gone" would typically all be reduced to the stem "go". This helps reduce the size of the vectors significantly, whilst preserving (by and large) meaning. Documents also typically have "stopwords" removed before stemming. Stopwords usually include "and", "the', etc. These words usually do not convey any of the sentiment or meaning of the document.

As a consequence of removing such stopwords, the remaining stems typically have quite low frequencies of occurrence. This is a key point to note for this work; with a sufficiently large $\lambda$, the $\operatorname{Po}(\lambda)$ distribution is well-approximated by a normal distribution with mean and variance both equal to $\lambda$. This approximation performs much more poorly with small $\lambda$, which is our experience of the typical case with text data after stemming and stopword removal. Empirically, we have found that infrequent but highly informative words can have $\lambda<1$, necessitating an alternative treatment to the normal approximation.

This poor approximation by the normal distribution is precisely why we, and other authors, have devoted time to studying extensions of PCA which are derived from an assumption of alternative exponential-family distributions. It is our expectation that when a normal approximation is inappropriate, so too is the standard PCA which, while not imposing any formal distributional assumptions, is deeply connected to the normal distribution's log-likelihood.

It is worth noting that instead of stems of individual words, it is also possible to consider stems of $n$ words, called $n$-grams. However, the same ideas for modelling these apply, it is merely a matter of which one chooses to be the base unit of meaning. It is also possible to work without stemming under the premise that stemming could throw away important information which ought to be preserved. The main danger of not stemming is possible diffusing the importance that would be assigned to one stem across multiple words to the extent that they appear unimportant, whilst the stem would be important.

This representation of text data does not lend itself to modelling by continuous distributions, so we will need to develop techniques for dimension reduction which are appropriate for it. In particular, PCA (for reasons which will be discussed later) is gen-
erally unsuitable.

### 1.3 Sparsity

We will now address the main motivation of Chapters 4 and 5, the notion of sparse loadings. Consider the example of PCA - each new variable is a linear combination of the original variables. In general, the coefficients of those linear combinations will be almost entirely non-zero. Our assumption is that this is typically not necessary, but is an artefact of overfitting to the observed data, and that reducing the number of nonzero coefficients can be used to improve the generalisation performance of PCA on new data. In fact, this is the premise of Zou, Hastie, and Tibshirani (2006) which proposes a method for Sparse Principal Component Analysis (SPCA).

To elaborate, our assumption that sparsity will improve generalisation performance arises from the complications of working with finite sample sizes. Typically we work with estimators which are unbiased, but in the finite sample case we should expect that they will be influenced by the peculiarities of the observed data. Imposing sparsity constraints is an attempt to counteract these finite sample peculiarities by requiring a simple structure to our estimands unless there is sufficient evidence in the data to outweigh the preference for sparsity.

There is an additional benefit to sparse loadings, beyond that of improved performance; sparse loadings are generally more interpretable. In a typical setup, having thousands of non-zero components makes it difficult to interpret how the original variables are being used in the new (dimension reduced) variables. Reducing the number of non-zero components means fewer original variables are involved in the construction of each new one; this can make it possible to give a subject-specific explanation for what the new variables represent.

Generally, sparse loadings can be achieved by adding a penalty to the optimisation procedure. Popular choices are the $L^{1}$ and Smoothly Clipped Absolute Deviation (SCAD) penalties which penalise the total magnitude of the components in slightly different ways and the $L^{0}$ penalty which penalises the number of non-zero components. We will revisit the first two of these penalties in Chapter 4, and the latter in Chapter 5.

### 1.4 Exponential Families

Much of the following work will be centred around the exponential family of distributions. This family of distributions encompasses the Gaussian and Poisson distributions amongst others; given our interest in applying dimension reduction to text data (which
we model with the Poisson) and the link between classical PCA and the Gaussian distribution, the exponential family seems a sensible choice.

We define an exponential family (EF) distribution by the form of its probability density (or mass) function. Let $\mathcal{X}$ have an EF distribution with parameter $\theta$ (known as the natural parameter). Then its density function $f(\mathbf{x} \mid \theta)$ has the form

$$
h(\mathbf{x}) \exp (\theta \cdot \mathbf{t}(\mathbf{x})-b(\theta))
$$

The function $\mathbf{t}(\mathbf{x})$ is a sufficient statistic for the data; in many cases this will simply be the identity function. The function $b$ is called the $\log$-partition function, and it normalises the distribution to have integral 1 over the whole domain. The function $h(\mathbf{x})$ is usually referred to as a base measure. It can be shown that $\mathbb{E}(X)=\nabla b(\theta)$ if $\mathbf{t}$ is the identity. In fact, more generally it is true that $\nabla b(\theta)=\mathbb{E}(\mathbf{t}(\mathcal{X}))$, and the higher order moments of the sufficient statistic can also be derived from the log-partition function. We define $\mathbf{g}$ as the left inverse of $\nabla b$ (that is, $g(\nabla b(\theta))=\theta$ ) and call it the canonical link function. This function is of central important when studying EF distributions, as it tells us how to map between the expectation and the natural parameters. This has an obvious application, mapping the sample mean of the sufficient statistic to an estimate of the natural parameters.

In Table 1.1 we will now give a summary of a number of common EF distributions along with their corresponding log-partition functions, base measures, sufficient statistics and canonical link functions. In the table, we use the symbol $\theta$ to denote the natural parameter as a function of the typical parameter(s).

Table 1.1: Usual parameter, natural parameter, sufficient statistic, log-partition, and canonical link functions for some common univariate exponential family distributions.

| Distribution $^{\text {Parameter }^{a}}$ | $\theta$ | $h(x)$ | $t(x)$ | $b(\theta)$ | $g(\theta)$ |  |
| :--- | :---: | :---: | :---: | :---: | :---: | :---: |
| Poisson $^{\text {Pormal }^{b}}$ | $\lambda$ | $\log (\lambda)$ | $1 / x!$ | $x$ | $\mathrm{e}^{\theta}$ | $\log (\theta)$ |
| Ninomial $^{c}$ | $\mu$ | $\mu / \sigma$ | $\frac{\exp \left(-x^{2} / 2 \sigma^{2}\right)}{\sigma \sqrt{2 \pi}}$ | $x / \sigma$ | $\theta^{2} / 2$ | $\theta$ |
| Exponential $^{\text {Bin }^{2}}$ | $\lambda$ | $\log \frac{p}{1-p}$ | $\binom{n}{x}$ | $x$ | $n \log \left(1+\mathrm{e}^{\theta}\right)$ | $\log \left(\frac{n}{\theta-n}\right)$ |

${ }^{a}$ the "usual" parameter
${ }^{b}$ with known variance $\sigma^{2}$
${ }^{c}$ with known number of trials $n$

With the canonical link function, we can define the so-called saturated model; given
a matrix of observations $\mathbf{X} \in \mathbb{R}^{n \times p}$, the saturated model assumes that each component $x_{i j}$ (for $i=1, \ldots n$ and $j=1, \ldots p$ ) is an observation of a random variable $X_{i j}$ from some EF distribution with natural parameter $\theta_{i j}=g_{i j}\left(x_{i j}\right)$. Finally, we can define the deviance, a measure of goodness of fit for models. Let $\ell_{S}(\mathbf{X})$ denote the log-likelihood of observations $\mathbf{X}$ under the saturated model and let $\ell_{\mathrm{A}}(\mathbf{X})$ denote the log-likelihood under some alternative model within the same exponential family. Then the deviance of that alternative model is

$$
-2\left(\ell_{\mathrm{A}}(\mathrm{X})-\ell_{\mathrm{S}}(\mathrm{X})\right)
$$

Remark. Although we have discussed exponential families in a fair deal of generality, for the rest of this thesis we will work under the restriction that the sufficient statistic function $\mathbf{t}$ is the identity function. Furthermore, we will usually (but not always) be referring to univariate distributions.

### 1.5 Related Work

In this section, we will look at related work in the literature, beginning with methods that make assumptions about the distribution of the data in Section 1.5.1, then at those which do not make distributional assumptions in Section 1.5.2. Of the two methods we have already considered, Sliced Inverse Regression (SIR) falls mostly into the non-distributional category as its primary assumption is on the relationship between the observed predictors and the response (also one could argue that the linear conditional expectation assumption is distributional, due to its equivalence to the assumption of an elliptically symmetric distribution); PCA also most naturally fits into the nondistributional category as despite its close relationship with the normal log-likelihood, it is typically motivated and derived from a non-distributional standpoint.

### 1.5.1 Distributional Dimension Reduction

### 1.5.1.1 Probabilistic Principal Component Analysis

The first distributional method we will look at is "Probabilistic Principal Component Analysis", proposed in Tipping and Bishop (1999). Here, the authors relate PCA to factor analysis. They model the observed data by

$$
\begin{equation*}
X \mid \mathbf{t} \sim \mathrm{N}\left(\mathrm{~W} \mathbf{t}+\mu, \sigma^{2} \mathbf{I}\right) \tag{1.5}
\end{equation*}
$$

where $\mathbf{I}$ is the identity matrix of appropriate dimension, $\mathbf{t}$ is a $k$ dimensional latent factor (i.e. it is unobserved) and $\mathbf{W}$ is a $p \times k$ dimensional matrix (corresponding to
the dimension $p$ of the observed variable $\mathcal{X}$ ). This is often referred to as a Gaussian model with isotropic noise. They also make the assumption that $\mathbf{t} \sim \mathrm{N}(0, \mathbf{I})$, where $\operatorname{dim}(\mathrm{t})<\operatorname{dim}(\mathcal{X})$. Marginalising, this gives that

$$
\begin{equation*}
x \sim \mathrm{~N}\left(\mu, \mathrm{wW}^{\mathrm{T}}+\sigma^{2} \mathbf{I}\right) \tag{1.6}
\end{equation*}
$$

which allows $\mu$ to be easily estimated by maximum likelihood as the sample mean of the observed data. Using Bayes rule, the authors show that the reversed conditional distribution is

$$
\begin{equation*}
\mathbf{t} \mid \boldsymbol{X} \sim \mathrm{N}\left(\mathbf{M}^{-1} \mathbf{W}^{\mathrm{T}}(X-\mu), \sigma^{2} \mathbf{M}^{-1}\right) \tag{1.7}
\end{equation*}
$$

where $\mathbf{M}=\mathbf{W}^{\mathbf{T}} \mathbf{W}+\sigma^{2} \mathbf{I}$. It is then shown that the maximum likelihood estimator for $\mathbf{W}$ can be given in explicit form as

$$
\begin{equation*}
\mathbf{W}_{\mathrm{ML}}=\mathbf{U}_{k}\left(\boldsymbol{\Lambda}_{k}-\sigma^{2} \mathbf{I}\right)^{1 / 2} \mathbf{R} \tag{1.8}
\end{equation*}
$$

where $\mathbf{U}_{k}$ is the matrix formed from the $k$ principal eigenvectors of the sample covariance of $\mathcal{X}, \boldsymbol{\Lambda}_{k}$ is a diagonal matrix consisting of the corresponding eigenvalues, and $\mathbf{R}$ is an arbitrary $k \times k$ orthogonal rotation matrix. In practice, the authors recommend choosing R as the identity matrix.

To relate this to the classical PCA, note that given (1.7) we can summarise $\mathbf{t} \mid \boldsymbol{X}$ by its mean

$$
\begin{equation*}
\mathbb{E}(\mathbf{t} \mid \boldsymbol{X})=\mathbf{M}^{-1} \mathbf{W}_{\mathrm{ML}}^{\mathrm{T}}(X-\mu) \tag{1.9}
\end{equation*}
$$

which, as $\sigma^{2} \rightarrow 0$ (and thus $\mathbf{M} \rightarrow\left(\mathbf{W}_{\mathrm{ML}}^{\mathrm{T}} \mathbf{W}_{\mathrm{ML}}\right)$ ), becomes an orthogonal projection into latent space. This, then, is the recovery of the traditional PCA from the probabilistic version. It is worth noting that the model proposed by Tipping and Bishop (1999) does become singular (and thus undefined) as $\sigma^{2} \rightarrow 0$ however.

### 1.5.1.2 Bayesian Principal Component Analysis

Fundamentally, Bayesian Principal Component Analysis (BPCA) (proposed by Bishop (1999)) is an extension of Probabilistic Principal Component Analysis (PPCA) which specifies a full Bayesian model for the latent variable model, shown in Figure 1.2. Here, $\sigma^{2}, \mu$ and $\alpha$ are treated as constants to be estimated for simplicity; the main difference between BPCA and PPCA then is the distribution specified over $\mathbf{W}$. This distribution is given as

$$
\begin{equation*}
\mathbb{P}(\mathbf{W} \mid \alpha)=\prod_{i=1}^{k-1}\left(\frac{\alpha_{i}}{2 \pi}\right)^{d / 2} \exp \left(-\frac{1}{2} \alpha_{i}\left\|\mathbf{w}_{i}\right\|^{2}\right) \tag{1.10}
\end{equation*}
$$

where $\mathbf{W}$ is once again a $p \times k$ dimensional matrix as in PPCA, $\mathbf{w}_{i}$ is the $i^{\text {th }}$ column of $\mathbf{W}$ and $\alpha_{i}$ is the $i^{\text {th }}$ component of $\boldsymbol{\alpha}$. The motivation for this form of prior distribution is


Figure 1.2: Plate diagram for Bayesian PCA
in Automatic Relevance Determination (ARD), proposed in Mackay (1995) initially for neural networks. It can be seen that $\alpha_{i}$ controls the precision (inverse variance) of the $i^{\text {th }}$ column of $\mathbf{W}$, with an increasingly large value of $\alpha_{i}$ indicating that the probability mass is concentrated around the 0 vector. ARD proposes that sufficiently large values of $\alpha_{i}$ indicate that the $i^{\text {th }}$ column of $\mathbf{W}$ is irrelevant, and should thus be removed from the model. This provides a significant advantage to BPCA, as it means in practice that one can set an initial value of $k$ to $p-1$ (one less than the observed dimension of $\mathcal{X}$ ) and allow ARD to remove columns from $\mathbf{W}$ to determine the appropriate value of $k$. This final $k$ value is referred to by the authors as the "effective dimension".

### 1.5.1.3 Collins et al. 2002

The work of Collins et al. (2002) begins with the observation that PCA is equivalent to finding a set of vector parameters $\theta_{1}, \ldots, \theta_{n}$ which lie in a low-dimensional subspace from some observed values $x_{1}, \ldots, x_{n}$ which are corrupted by Gaussian noise. Using this, they extend the notion of PCA to the exponential family using the notion of Bregman divergences. They do this by noting that the conditional log-likelihood of an observation $x$ given the natural parameter $\theta$ is

$$
\begin{equation*}
\log \mathbb{P}(x \mid \theta)=\log h(x)+x \theta-b(\theta) \tag{1.11}
\end{equation*}
$$

using the notation from Section 1.4 and assuming that the sufficient statistic $t$ is the identity. When maximising this log-likelihood with respect to $\theta$, the term $\log h(x)$ is constant, and so can be disregarded. Therefore, the primary difference between exponential families is in the function $b(\theta)$. In order to relate this to Bregman divergences, we will first define what a Bregman divergence is.

Definition 1.5.1. Let $F: \Delta \rightarrow \mathbb{R}$ be a differentiable and strictly convex function on a closed, convex set $\Delta \subseteq \mathbb{R}$. Then for any $p, q \in$ real, the Bregman divergence associated with $F$ is

$$
\begin{equation*}
B_{F}(p \| q)=F(p)-F(q)-f(q)(p-q) \tag{1.12}
\end{equation*}
$$

where $f(x)=F^{\prime}(x)$.
We then relate $b(\theta)$ to a "dual" Bregman divergence by the equation

$$
\begin{equation*}
F\left(b^{\prime}(\theta)\right)+b(\theta)=b^{\prime}(\theta) \theta \tag{1.13}
\end{equation*}
$$

Collins et al. (2002) note that, under some fairly general conditions, $f(x)=\left[b^{\prime}\right]^{-1}(x)$ which, as noted in Section 1.4 is precisely the "canonical link function". Using this, we can see that the log-likelihood (1.11) can be rewritten as

$$
\begin{equation*}
\log \mathbb{P}(x \mid \theta)=\log h(x)+F(x)-B_{F}(x \| \theta) \tag{1.14}
\end{equation*}
$$

thus maximising the likelihood with respect to $\theta$ can be achieved simply by minimising the Bregman divergence $B_{F}(x \| \theta)$. The authors overload the Bregman divergence notation for vectors by

$$
\begin{equation*}
B_{F}(\mathrm{v} \| \mathbf{w})=\sum_{i=1}^{p} B_{F}\left(v_{i} \| w_{i}\right) \tag{1.15}
\end{equation*}
$$

for all $\mathbf{v}, \mathbf{w} \in \mathbb{R}^{p}$, and for matrices by

$$
\begin{equation*}
B_{F}(\mathbf{V} \| \mathbf{W})=\sum_{i=1}^{p} \sum_{j=1}^{q} B_{F}\left(V_{i j} \| W_{i j}\right) \tag{1.16}
\end{equation*}
$$

for all $\mathbf{V}, \mathbf{W} \in \mathbb{R}^{p \times q}$. Finally, we define $\boldsymbol{\Theta}=\mathbf{A V}$ as the $n \times p$ matrix of natural parameters, such that the $i^{\text {th }}$ row is the vector of natural parameters associated with the $i^{\text {th }}$ observation $\mathbf{x}_{i}$. Here, we have $\mathbf{A} \in \mathbb{R}^{n \times k}$ and $\mathbf{V} \in \mathbb{R}^{k \times p}$, with $k<p$, and $p$ the dimension of the observed variables $\mathbf{x}_{i}, i=1, \ldots, n$. Then the log-likelihood can be written as

$$
\begin{equation*}
\sum_{i=1}^{n} B_{F}\left(\mathbf{x}_{i} \| g\left(\theta_{i}\right)\right) \tag{1.17}
\end{equation*}
$$

where $\boldsymbol{\theta}_{i}$ is the $i^{\text {th }}$ row of $\boldsymbol{\Theta}$. One can then think of $\mathbf{V}$ defining a lower-dimensional basis for the surface $Q(\mathbb{V})=\left\{g(\mathbf{a V}) \mid \mathbf{a} \in \mathbb{R}^{k}\right\}$, such that this surface passes close to each observation $\mathbf{x}_{i}$, and each row of A giving the coefficients with respect to that basis of the closest point on $Q(\mathbf{V})$ to $\mathbf{x}_{i}$. Thus, A gives us the lower dimensional "projection" of the observed data.

### 1.5.1.4 Bayesian Exponential Principal Component Analysis

The method of Mohamed et al. (2009), which they call Bayesian Exponential Principal Component Analysis (BXPCA), shares a similar approach to that of Collins et al. (2002) but from a Bayesian perspective. In Figure 1.3 we show the plate diagram for the model.

To elaborate on the model, we have


Figure 1.3: Plate diagram for Bayesian Exponential Family PCA

- $\mu$ is drawn from a normal distribution with mean $\mathbf{m}$ and covariance matrix $\mathbf{S}$
- $\Sigma$ is a diagonal matrix with the $i^{\text {th }}$ diagonal element drawn from an inverse gamma distribution parametrised by $\alpha$ and $\beta$
- Each of the $n$ vectors $\mathbf{v}_{i}(i=1, \ldots, n)$ are drawn from a normal distribution with mean $\boldsymbol{\mu}$ and covariance matrix $\Sigma$
- Each observation $\mathbf{x}_{i}$ of the random variable $\mathcal{X}$ is drawn from an exponential family distribution with natural parameter $\sum_{j=1}^{k} v_{i j} \boldsymbol{\eta}_{j}$, where $v_{i j}$ is the $j^{\text {th }}$ element of $\mathbf{v}_{i}$
- Each $\boldsymbol{\eta}_{i}(i=1, \ldots, k)$ is drawn from the conjugate prior distribution corresponding to the exponential family distribution of $\mathcal{X}$

Having specified a full Bayesian model, the authors then propose a hybrid Monte Carlo for estimating the model parameters. Then the matrix $\mathbf{V}$, formed by the vectors $\mathbf{v}_{i}, i=1, \ldots n$ as rows, is the analogue of the data after projection by PCA.

### 1.5.1.5 Simple Exponential Principal Component Analysis

Remark. Chapter 5 is an extension of this method, as such a detailed technical description will be given there. A brief description is given here for completeness.

Simple Exponential Principal Component Analysis (SePCA), put forward in Li and Tao (2013), uses a very similar model to that of Bishop (1999) in BPCA. The key difference is that, rather than the observed data having mean given by the product of the latent factors, it is the natural parameter which is given by this. As such, in the Gaussian case this method reproduces BPCA as a special case, while being able to model data
from any exponential family distribution. This method is also suitable for ARD, and the authors give details of how to apply ARD to automatically determine an appropriate dimension for the reduced dimension data.

### 1.5.1.6 Generalised Principal Component Analysis

Remark. Chapter 4 is an extension of this method, as such a detailed technical description will be given there. A brief description is given here for completeness.

In Landgraf and Lee (2015b) (an extension of Landgraf and Lee (2015a)), the authors propose an extension of PCA to the exponential family. Like Collins et al. (2002), this begins by recognising a connection between PCA and the estimation of the parameters of a Gaussian distribution. Generalised Principal Component Analysis (GPCA), however, uses a different form for the decomposition of the natural parameters in their extension to the exponential family, and optimises by minimising the deviance from the saturated model.

### 1.5.1.7 Sparse Probabilistic Principal Component Analysis

This method from Guan and Dy (2009) essentially extends PPCA (Tipping and Bishop 1999) in much the way that SPCA (detailed in Section 1.5.2.1) extended PCA. Specifically, they use the same model as in PPCA (Section 1.5.1.1), with a choice of additional prior distributions over $\mathbf{W}$.

The first proposed prior is a two-level hierarchical prior, equivalent to a Laplacian prior; the first level is a normal prior such that $\mathbf{W}_{i j} \mid z_{i j} \sim N\left(0, z_{i j}\right)$, the second level is an exponential distribution on each $z_{i j}$ such that $z_{i j} \sim \operatorname{Exp}(\lambda)$. When marginalising out the $z_{i j}$, this is equivalent to a Laplacian prior directly on $\mathbf{W}$. This prior is similar to the ARD prior, such as is used in BPCA, although this was not explored in the original paper.

The authors also propose an inverse-Gaussian prior and a Jeffrey's prior (that is, an uninformative prior proportional to the square root of the determinant of the Fisher information matrix). They show that these three priors are equivalent to sparsity-inducing penalty functions; the precise equivalences are shown in Table 1.2.

### 1.5.1.8 Sparse Exponential Family Principal Component Analysis

In Lu et al. (2016), the authors propose the method Sparse Exponential Family Principal Component Analysis (SEFPCA) as an ideological extension of SPCA to EF distributions. Denoting observed data by $\mathbf{X}$ composed of $n$ observations $\mathbf{x}_{i}, i=1, \ldots, n$ as the rows, they specify the form of the natural parameters for the $i^{\text {th }}$ observation as $\mathbf{W}^{\mathrm{T}} \mathbf{z}_{i}+\boldsymbol{\mu}$.

| Prior | Equivalent Penalty |
| ---: | :---: |
| Hierarchical Laplace | $\sum_{i} \sum_{j} \lambda\left\|\mathbf{W}_{i j}\right\|$ |
| Inverse Gaussian | $-\frac{1}{2} \sum_{i} \sum_{j}\left(\log \left(\mathbf{W}_{i j}^{2}+\lambda\right)+\log \left(K_{1}\left(\frac{\sqrt{\lambda} \sqrt{\mathbf{W}_{i j}^{2}+\lambda}}{\mu}\right)\right)\right)$ |
| Jeffrey's | $\sum_{i} \sum_{j} \log \left(\mathbf{z}_{i j}\right)$ |

Table 1.2: Equivalences between priors on $\mathbf{W}$ and penalty functions for Sparse Probabilistic PCA. Here, $K_{1}$ is the modified Bessel function of the second kind with order 1 and $\lambda$ is a penalty controlling the strength of the penalty function.

Here, $\mu$ specifies the "average" natural parameters which are then diverged from by the lower-dimensional latent variable $\mathbf{z}_{i}$ through $\mathbf{W}$.

Under this specification, SEFPCA seeks to minimise the function

$$
\begin{equation*}
\sum_{n} b\left(\mathbf{W}^{\mathrm{T}} \mathbf{z}_{n}+\boldsymbol{\mu}\right)-\operatorname{Trace}\left(\left(\mathbf{z W}+\mathbf{1} \boldsymbol{\mu}^{\mathrm{T}}\right) \mathbf{X}^{\mathrm{T}}\right)+P(\mathbf{W}, \boldsymbol{\mu}) \tag{1.18}
\end{equation*}
$$

where $b$ is the log-partition function as described in Section 1.4 which is specified by the choice of exponential family distribution, and $P(\mathbf{W}, \mu)$ is a penalty function to induce sparsity. The form of this penalty function is

$$
\begin{equation*}
P(\mathbf{W}, \mathbf{b}):=\lambda_{0}\left\|\mathbf{Z W}+\mathbf{1 b}^{\mathrm{T}}\right\|_{2}^{2}+\sum_{i=1}^{k} \lambda_{i}\left|\mathbf{W}_{i}\right| \tag{1.19}
\end{equation*}
$$

which, in addition to the sparsifying effect, also helps with the stability of the estimation process when the dimension of the data is higher than the number of observations.

### 1.5.1.9 Multinomial Inverse Regression

Moving away from methods derived from PCA, we will now look at a method specifically designed for dimension reduction of text data. In Taddy (2013) and Taddy (2015), Taddy introduces Multinomial Inverse Regression (MNIR), a method for supervised dimension reduction of text data, based around a multinomial topic model. Specifically, the model is

- $\mathbf{x}_{i} \sim \operatorname{MN}\left(\mathbf{q}_{i}, m_{i}\right)$ for $i=1, \ldots, n$, where $\mathbf{x}_{i}$ is the $i^{\text {th }}$ observation
- $q_{i j}=\frac{\exp \left(\eta_{i j}\right)}{\sum_{l=1}^{p} \exp \left(\eta_{i j}\right)}$ is the probability of the $j^{\text {th }}$ term appearing in the $i^{\text {th }}$ observation, $j=1, \ldots, p$
- $\eta_{i j}=\alpha_{j}+u_{i j}+\mathbf{v}_{i}^{\mathrm{T}} \boldsymbol{\phi}_{j}$

Here the $\mathbf{v}_{i}$ are $k$-dimensional "response factors" which capture the dependency of the term counts on the response $y$. That is, the $v_{i}$ are (possibly random) functions dependent on $y$. The terms $u_{i j}$ can be assembled into vectors $\mathbf{u}_{i}=\left[u_{i 1}, \ldots, u_{i p}\right]^{\mathrm{T}}$ of "subject effects". In this way, the individual term probabilities are decomposed into a "mean", a subject specific effect, and the response specific effect introduced via the inverse regression coefficients $\boldsymbol{\phi}_{j}$.

In order to estimate these parameters, Taddy specifies a full Bayesian model of priors for $\boldsymbol{\alpha}=\left[\alpha_{1}, \ldots, \alpha_{p}\right]^{\mathrm{T}}, \boldsymbol{\Phi}=\left[\boldsymbol{\phi}_{1}, \ldots, \boldsymbol{\Phi}_{p}\right]^{\mathrm{T}}$ and $\mathbf{U}=\left[\mathbf{u}_{1}, \ldots, \mathbf{u}_{n}\right]^{\mathrm{T}}$. In particular, the model specified is

- $\alpha_{j} \sim \mathrm{~N}(0,1)$, which identifies the logistic multinomial model, removing the need to specify a null category
- $\phi_{j k} \sim \operatorname{Lap}\left(0,1 / \lambda_{j k}\right)$, where each Laplace distribution is independent of the others and the $\lambda_{j k}$ are coefficient-specific precision parameters
- $\lambda_{j k} \sim \operatorname{Gamma}(s, r)$, where the $s$ and $r$ are shared hyperparameters for the shape and rate of the gamma distribution respectively
- $\exp \left(u_{i j}\right) \sim \operatorname{Gamma}(1,1)$

Importantly, Taddy shows that this model produces a sufficient dimension reduction, where

$$
\begin{equation*}
y_{i} \Perp \mathbf{x}_{i}\left|\mathbf{v}_{i} \Rightarrow y_{i} \Perp \mathbf{x}_{i}\right| \boldsymbol{\Phi}^{\mathrm{T}} \mathbf{x}_{i} \tag{1.20}
\end{equation*}
$$

They also develop a bespoke optimisation algorithm for solving this problem, stating that the typical approaches to Bayesian inference are generally too slow for application to text data which tends to be very high dimensional.

### 1.5.1.10 Latent Dirichlet Allocation

Latent Dirichlet Allocation (LDA), put forward by Blei et al. (2003), is a generative probabilistic model for text, displayed graphically as a plate diagram in Figure 1.4

To elaborate, for the $d^{\text {th }}$ "document" (that is, for each observation), we assume a vector $\theta_{d}$ of topic probabilities is drawn from a Dirichlet distribution with parameter $\boldsymbol{\alpha}$. Each of the $n_{d}$ words within that document are generated by drawing a "topic" from a multinomial distribution with topic probabilities $\theta_{d}$, then drawing a word from a multinomial distribution with probabilities $\boldsymbol{\beta}_{z}$. Here, $\boldsymbol{\beta}_{z}$ is the $z^{\text {th }}$ column of matrix $\boldsymbol{\beta}$ of word probabilities.

From a dimension reduction point of view, it is the posterior distribution of $\theta_{d}$ for $d=1, \ldots, D$ which is of interest. In particular, we can use the posterior mean for each


Figure 1.4: Plate diagram for Latent Dirichlet Allocation
$\theta_{d}$ as a lower-dimensional representation of the document. As the number of topics should be significantly lower than the number of words, this should achieve a significant reduction in the dimension of the data.

### 1.5.2 Non-Distributional Dimension Reduction

In contrast to the methods of the previous subsection, which have distributional assumptions at their heart, these methods are motivated primarily from a non-distributional perspective.

### 1.5.2.1 Sparse Principal Component Analysis

The idea behind SPCA, formulated by Zou, Hastie, and Tibshirani (2006), is to create sparse PCA loadings by enforcing an elastic net (Zou and Hastie 2005) penalty on the components of the loading matrix. To state this more precisely, the "SPCA Criterion" they gave was to find an approximation of some data $\mathbf{X} \in \mathbb{R}^{n \times p}$ of the form $\mathbf{X B A}^{\mathrm{T}}$ with $\mathrm{A}, \mathbf{B} \in \mathbb{R}^{p \times k}$ which minimises

$$
\begin{equation*}
\left\|\mathbf{X}-\mathbf{X B A}^{\mathrm{T}}\right\|_{2}^{2}+\lambda \sum_{i=1}^{k}\left\|\boldsymbol{\beta}_{i}\right\|_{2}^{2}+\sum_{j=1}^{k} \lambda_{1, j}\left\|\boldsymbol{\beta}_{j}\right\|_{1} \tag{1.21}
\end{equation*}
$$

subject to the constraint $\mathbf{A}^{\mathrm{T}} \mathbf{A}=\mathbb{I}$, and where the vectors $\boldsymbol{\beta}_{i}$ are the columns of $\mathbf{B}$. The loadings matrix is $\mathbf{B}$, and the two summation terms are the elastic net penalty, a combination of the $L^{1}$ and $L^{2}$ penalties which usually gives superior performance to either of them in isolation. This penalty has been widely applied in regression analysis with success; given the regression-like formulation of PCA in minimising squared error the authors had good reason to suspect that it would be effective in this context.

The critical difficulty in applying SPCA is in the number of tuning parameters, with one parameter for the $L^{2}$ penalty, and $k$ for the $L^{1}$ penalty. In practice, the authors suggest choosing $\lambda$ a small, positive number for data with $n>p$ where it mainly fulfils
a role in helping to reduce potential collinearity problems with the data. In order to choose the values $\left\{\lambda_{1, j}, j=1, \ldots k\right\}$, the authors suggest trying a range, determining the best choice by balancing explained variance against sparsity.

### 1.5.2.2 Joint Sparse Principal Component Analysis

The idea of Joint Sparse Principal Component Analysis (JSPCA) (Yi et al. 2017) is similar in essence to SPCA in using penalisation to induce sparsity, but rather than use the squared $L^{2}$ norm, the authors use the unsquared $L^{2}$ norm directly for both the penalty and the reconstruction error. In contrast to SPCA, they also do not use $L^{1}$ penalisation on the components. To be precise, they solve the optimisation problem

$$
\begin{equation*}
\underset{\mathrm{A}, \mathbf{B} \in \mathbb{R}^{p \times k}}{\arg \min }\left\|\mathbf{X}-\mathbf{X B A}^{\mathrm{T}}\right\|_{2}+\lambda \sum_{i=1}^{k}\left\|\boldsymbol{\beta}_{i}\right\|_{2} \tag{1.22}
\end{equation*}
$$

where again the vectors $\beta_{i}(i=1, \ldots, k)$ are the columns of $\mathbf{B}$. The first term is the reconstruction error. The authors' rationale for using the non-squared $L^{2}$ norm is to provide better robustness to outliers in the data. They use a classification problem to demonstrate this, generating a set of noise-corrupted images (with different types of noise) and applying (amongst others) classical PCA and their JSPCA. Each observation is then classified using a nearest-neighbours classifier. Using the classification accuracy as the evaluation metric, the authors show the relative performances across multiple datasets. In some of the authors' experiments, there is good evidence that JSPCA performs better than any other PCA variant on noise-corrupted data, while some of the experiments show more mixed results.

### 1.5.2.3 Robust Principal Component Analysis

Though not a sparse method, we include the work of Kwak (2008) here for its similar aim. The focus is also on a PCA variant which is robust to outliers, something which sparse methods often achieve. In contrast to JSPCA, Robust Principal Component Analysis (RPCA) achieves this from the framework of the projection variance maximisation formulation of PCA. In this formulation, the aim is to find the projection matrix $\mathbf{U} \in \mathbb{R}^{p \times k}$ which maximises

$$
\begin{equation*}
\max \left\|\mathbf{U}^{\mathrm{T}} \mathbf{X}\right\|_{1} \tag{1.23}
\end{equation*}
$$

where $\mathbf{X}$ is the matrix of observed data, and where we impose the additional constraint that $\mathbf{U}^{\mathrm{T}} \mathbf{U}=\mathbf{I}$. The authors contrast this to another possible choice of $L^{1}$ PCA-like optimisation problems

$$
\begin{equation*}
\min _{\mathbf{U} \in \mathbb{R}^{p \times k}}\left\|\mathbf{X}-\mathbf{U} \mathbf{U}^{\mathrm{T}} \mathbf{X}\right\|_{1} \tag{1.24}
\end{equation*}
$$

which, while more robust to outliers, is not invariant to rotations. The criteria they use, (1.23), is invariant to rotations and is robust to outliers. It can be seen as maximising the $\left(L^{1}\right)$ dispersion in the feature space as opposed to the input space.

It is important to note that, while maximising the $L^{2}$ dispersion in feature space and minimising the squared $L^{2}$ reconstruction error are duals to one another, and thus both result in classical PCA, the two problems (1.23) and (1.24) are not dual problems. This means that the calculated projections will not, in general, be equal.

### 1.5.2.4 Sparse Principal Component Analysis by Rotation and Truncation

Proposed by Hu et al. (2016), Sparse Principal Component Analysis by Rotation and Truncation (SPCArt) is an alternative sparse PCA method to SPCA. Unlike many of the methods discussed so far, the authors do not propose either a probabilistic model to estimate, or introduce a penalisation of the PCA optimality criterion.

Instead, their method derives from a matrix approximation problem. The authors make use of the fact that $\operatorname{Span}(\mathbf{B})=\operatorname{Span}(\mathbf{R B})$ for all full-rank rotation matrices $\mathbf{R}$. In particular, rotating the PCA loadings by a full-rank rotation produces an orthogonal basis which spans the same subspace. The authors propose to solve

$$
\begin{equation*}
\min _{\mathbf{R} \in \mathbb{R}^{k \times k}, \mathbf{R}^{\mathrm{T}} \mathbf{R}=\mathrm{I}}\left\|\mathbf{V}_{1: k} \mathbf{R}\right\|_{0} \tag{1.25}
\end{equation*}
$$

where $\mathrm{V}_{1: k}$ is the matrix of the $k$ principal components, and the $\|\cdot\|_{0}$ norm is applied to the matrix by summing the column norms. Due to the difficulty of solving this problem, the authors solve the approximation

$$
\begin{equation*}
\min _{\mathbf{R} \in \mathbb{R}^{k \times k}, \mathbf{P} \in \mathbb{R}^{n \times k}} \frac{1}{2}\left\|\mathbf{V}_{1: k}-\mathbf{P R}\right\|_{F}^{2}+\lambda \sum_{i=1}^{n}\left\|\mathbf{P}_{i}\right\|_{1} \tag{1.26}
\end{equation*}
$$

subject to $\left\|\mathbf{P}_{i}\right\|_{2}=1$ for $i=1, \ldots, n$ and $\mathbf{R}^{\mathrm{T}} \mathbf{R}=\mathbf{I}$, where $\|\cdot\|_{F}$ is the Frobenius norm. Here, the approximation to $\mathbf{V}_{1: k}$ is given by $\mathbf{P R}$, and in general it may not be orthogonal (though for a sufficiently good approximation, it should be close to orthogonal).

As part of the solution process to this approximation of the original problem, thresholding takes place. Although the solution properly calls for a type of soft thresholding, the authors propose four different types (including the soft thresholding), each of which possesses different properties. Of note is one type of thresholding, which the authors call "truncation by sparsity" which allows direct control over the sparsification of the result.

### 1.5.2.5 Non-negative Matrix Factorisation

Like SPCArt, Nonnegative Matrix Factorisation (NMF) is a method based around matrix approximation, rather than a PCA-like criteria to optimise. There are a number of distinct NMF algorithms (see Gillis (2014) for an overview), but they all focus on solving the same problem. That problem is the decomposition of a non-negative matrix $\mathbf{X}$ into the product of two non-negative matrices $\mathbf{W}$ and $\mathbf{H}$, such that $\mathbf{X} \approx \mathbf{W H}$ with $\mathbf{W}$ a $n \times k$ matrix and $\mathbf{H}$ a $k \times p$ matrix. Here we expect $k<p$. Then $\mathbf{H}$ can be interpreted as a basis matrix, with the rows of $\mathbf{W}$ giving the co-ordinates of each observation $\mathbf{x}$ forming the rows of $\mathbf{X}$ with respect to the basis $\mathbf{H}$.

In the context of text data, this factorisation has a pleasing interpretation; $\mathbf{H}$ corresponds to topics, with each column being counts (or frequencies) of terms appearing within each topic. We can then interpret $\mathbf{W}$ as decomposing each observation into its topics. Note that this interpretation is not strict, no conditions are enforced to make this interpretation exact, but it is often appropriate in practice.

In the context of dimension reduction, we can use the matrix $\mathbf{W}$ as a lower-dimensional representation of our observed data $\mathbf{X}$.

### 1.6 Conclusion

In this chapter we have introduced the concepts of text data, dimension reduction, and some of the popular related methods. In the following chapter, we will introduce a Poisson based model for text data, and from there specify a Bayesian model which will allow us to find a dimension reduction transformation for the data.

## - Chapter 2 -

## Poisson Inverse Regression

### 2.1 Introduction

As each column in a document-term matrix consists of observed counts, a natural idea is to model these variables as observations from a Poisson random variable. Conveniently, in Cook (2007) we find that such a model fits into a framework of exponential family inverse regression, which leads to a simple dimension reduction framework. In this vein, we will propose a simple Poisson model for text data, and from it show how it leads to a linear dimension reduction transformation which is statistically sufficient. The details of this will be given in Section 2.2; in Section 2.3 we will give two methods of estimating the linear transformation. Recalling the similar method of Multinomial Inverse Regression (MNIR) from Section 1.5, we will compare the two methods in Section 2.4, before concluding in Section 2.5

### 2.2 Poisson Inverse Regression

In this section, we will define the Poisson Inverse Regression (PoIR) model. We begin by discussing a very simple multi-variate extension to a Poisson model, then add dependence on the response through the natural parameter.

Recall that if $\lambda$ is the mean of a Poisson distributed random variable (i.e. the usual parameter), then its natural parameter is $\kappa:=\log (\lambda)$. We then have the form

$$
\frac{1}{x!} \exp (x \kappa-\exp (\kappa))
$$

for the probability mass function. We can extend this definition to vector-valued random variables $\mathcal{X}$ where each component is an independent Poisson distribution with natural parameter $\kappa_{i}$ :

$$
\begin{equation*}
\frac{1}{\prod_{i=1}^{p} x_{i}!} \exp \left(\mathbf{x} \cdot \boldsymbol{\kappa}-\sum_{i=1}^{p} \exp \left(\kappa_{i}\right)\right) \tag{2.1}
\end{equation*}
$$

where $\kappa_{i}$ is the natural parameter for the $i^{\text {th }}$ component of $\mathcal{X}$.

The inverse regression model we will use assumes that is it $\mathcal{X} \mid y$ which has a Poisson distribution. In order to capture this relationship, we assume that the natural parameter is a function of $y$. In particular, we assume the following form

$$
\begin{equation*}
\kappa_{i}(y)=\mu_{i}+\boldsymbol{\eta}_{i}^{\mathrm{T}} \boldsymbol{v}(y) \tag{2.2}
\end{equation*}
$$

where $y$ is the observed value of $y$. Here, $\boldsymbol{\eta}_{i}$ denotes the $i$ th column of a matrix $\mathbf{H}$ and $\boldsymbol{v}(y)$ is some vector-valued function of $y$ which, for the moment, we will place no restrictions on in terms of form. This form assumes that there is a central mean tendency to the natural parameter, captured by $\boldsymbol{\mu}$, which is shared across all values of $y$. Then the term $\boldsymbol{\eta}_{i}^{\mathrm{T}} \boldsymbol{v}(y)$ captures the divergence from that central tendency. If we briefly relate this back to the mean parameter of the Poisson distribution, then we can see that the "divergence" from the central tendency has a multiplicative effect on the expected count. In practice, this assumed form is difficult to check without making more restrictive assumptions.

We now present the following theorem, which validates the use of this method of dimension reduction.

Theorem 2.2.1. If $\mathcal{X} \mid y$ has a Poisson distribution with natural parameter of the form (2.2), then $\mathbb{S}_{\mathbf{H}}:=$ ColSpace $(\mathbf{H})$ is a dimension reduction subspace under the definition in Def. 1.1.2.

Proof. Substituting 2.2 into the model 2.1, we obtain

$$
\begin{aligned}
& \frac{1}{\prod_{i=1}^{k} x_{i}!} \exp \left(\mathbf{x} \cdot \boldsymbol{\mu}+\boldsymbol{v}(y)^{\mathrm{T}} \mathbf{H}^{\mathrm{T}} \mathbf{x}-\sum_{i=1}^{k} \exp \left(\mu_{i}+\boldsymbol{\eta}_{i}^{\mathrm{T}} \boldsymbol{v}(y)\right)\right) \\
= & {\left[\frac{1}{\prod_{i=1}^{k} x_{i}!} \exp \left(\mathbf{x} \cdot \boldsymbol{\mu}-\sum_{i=1}^{k} \exp \left(\mu_{i}\right)\right)\right] \times\left[\exp \left(\boldsymbol{v}(y)^{\mathrm{T}} \mathbf{H}^{\mathrm{T}} \mathbf{x}-\sum_{i=1}^{k} \exp \left(\boldsymbol{v}(y)^{\mathrm{T}} \boldsymbol{\eta}_{i}\right)\right)\right] }
\end{aligned}
$$

Notice that the left hand term is a function only of $\mathbf{x}$ and the right hand term is a function only of $\mathbf{H}^{\mathrm{T}} \mathbf{x}$ and $y$. Then through the Fisher-Neyman factorisation theorem for sufficiency we have that the distribution of $\mathcal{X} \mid\left(\mathbf{H}^{\mathrm{T}} \boldsymbol{X}, y=y\right)$ is the same as the distribution of $\mathcal{X} \mid \mathbf{H}^{\mathrm{T}} \boldsymbol{X}$ for all values of $y$. Consequently, $y \Perp \mathcal{X} \mid \mathbf{H}^{\mathrm{T}} \boldsymbol{X}$.

Now that we have established that we can, indeed, use this model to perform dimension reduction on text data, it seems prudent to name it.

Definition 2.2.1. Let $\boldsymbol{X} \mid \boldsymbol{y}=y \sim \operatorname{Po}\left(\exp \left(\mu+\boldsymbol{v}(y)^{\mathrm{T}} \mathbf{H}\right)\right)$, where we overload the notation $\operatorname{Po}(\cdot)$ for a vector parameter to mean independent Poisson distributions for each component of the random vector with mean given by the corresponding component of the vector parameter (and where we apply the exponential component-wise). This model is the "Poisson inverse regression" model.

In order to actually estimate this model, we will first establish a full Bayesian model which will be used to perform the estimation. To do this, we assign prior distributions to each of the parameters in the model. In order to encourage sparsity, we place a Laplace prior with mean 0 and on each component of $\mathbf{H}$. We use a shared rate parameter $1 / h$ for each of these priors. For flexibility, we place an exponential hyperprior on $h$ with a fixed rate parameter $1 / \rho_{h}$ to be determined by experimentation. Our experience suggests that estimation is not too sensitive to the choice of $\rho_{h}$, so we may search across a coarse grid for the best value in respect to convergence of MCMC methods.

Similarly, we place a Laplace prior on each component of $\mu$ with mean 0 and rate parameter $1 / m$. This prior is flexible enough to accommodate a wide range of values for $\mu$, while requiring increasingly more evidence from the observed data to deviate to particularly high or low values. Upon $m$ we place an exponential hyperprior with rate parameter $1 / \rho_{m}$.

For each value of $y$, we place an independent standard normal prior on $v(y)$. In the case of discrete-valued responses this is easy to accomplish; in the continuous-valued case, we suggest slicing the response into discrete groups. This prior is quite uninformative; for a given dataset one may have a better understanding of the form this function should take, in this case practitioners can use that function directly rather than attempting to estimate it.

Using sequential conditioning, we can split the joint probability function into a multiplication of several probabilities whose random variables are conditionally independent as follows. Summarised, also, are the priors we have imposed.

$$
\begin{align*}
\mathbb{P}(\mathbf{x}, y, h, \mathbf{H}, m, \mu, \boldsymbol{v})=\mathbb{P}(\mathbf{x} \mid y, h, \mathbf{H}, m, \boldsymbol{\mu}, \boldsymbol{v}) \mathbb{P}(\mathbf{H} \mid h) \mathbb{P}(h) & \\
& \times \mathbb{P}(\mu \mid m) \mathbb{P}(m) \mathbb{P}(\boldsymbol{v} \mid y) \mathbb{P}(y) \tag{2.3}
\end{align*}
$$

$$
\begin{align*}
\eta_{i j} & \sim \operatorname{Lap}\left(0, \frac{1}{h}\right)  \tag{2.4}\\
\mu_{i} & \sim \operatorname{Lap}\left(0, \frac{1}{m}\right)  \tag{2.5}\\
h & \sim \operatorname{Exp}\left(\frac{1}{\rho_{h}}\right)  \tag{2.6}\\
m & \sim \operatorname{Exp}\left(\frac{1}{\rho_{m}}\right)  \tag{2.7}\\
v_{i} \mid y & \sim \mathrm{~N}(0,1)  \tag{2.8}\\
\mathbf{x}_{i} \mid y, \mu, \mathbf{H}, \boldsymbol{v} & \sim \operatorname{Po}\left(\exp \left(\mu_{i}+\eta_{i}^{\mathrm{T}} v(y)\right)\right) \tag{2.9}
\end{align*}
$$

Here we use the convention that the pdf of a Laplace distributed variable with mean parameter $\mu$ and rate parameter $b$ is

$$
f(x)=\frac{1}{2 b} \exp \left(-\frac{|x-\mu|}{b}\right)
$$

and the convention that the pdf of an Exponential random variable with rate parameter $b$ has pdf

$$
f(x)=\frac{1}{b} \exp \left(-\frac{x}{b}\right)
$$

Having fully specified the Bayesian model, we can then proceed to use methods from Bayesian inference to estimate $\mathbf{H}$.

### 2.3 Estimation

We have used two main methods to estimate the value of $\mathbf{H}$. Firstly, we used maximum a posteriori estimation, where we maximise $\mathbb{P}(\mu, H, \nu, m, h \mid X, y)$, where $X$ is a matrix of observed data and $\mathbf{y}$ is a vector of observed responses. Using (2.3) and Bayes's theorem we can compute the a posteriori probability. For simplicity, we have used an "off the shelf" optimisation routine in R to find the maximum values. However, it would be worth investing time in deriving a bespoke routine for calculating the MAP estimate, as was done in Taddy (2013). Note that this can be extended to vector responses in a straightforward manner, in which case we have a matrix of response values $\mathbf{Y}$

We also made use of Markov Chain Monte Carlo methods; the aim of such methods is to sample from a Markov chain whose equilibrium distribution is the posterior distribution of interest. In our case, that means that we wish to construct a Markov chain whose equilibrium distribution is the PoIR posterior distribution.

To be precise, a (discrete time) Markov chain is a type of stochastic process, where the distribution of possible values at some time-point $t$ is independent of all previously
obtained values, except the value at $t-1$. This definition can be expanded to continuous time, but will not be necessary here. The formal definition is as follows:

Definition 2.3.1. A discrete-time Markov chain is given by a sequence of random variables $\boldsymbol{X}_{1}, \boldsymbol{X}_{2}, \ldots$ such that

$$
\mathbb{P}\left(X_{t}=\mathrm{x} \mid X_{1}=\mathrm{x}_{1}, \ldots, X_{t-1}=\mathrm{x}_{t-1}\right)=\mathbb{P}\left(X_{t}=\mathrm{x} \mid X_{t-1}=\mathrm{x}_{t-1}\right)
$$

The equilibrium distribution is then given by $\lim _{t \rightarrow \infty} \mathbb{P}\left(X_{t}\right)$; that is, the limiting unconditional probability density.

In Metropolis et al. (1953), a method was developed for constructing a Markov chain with the desired equilibrium distribution under some restrictions; the work of Hastings (1970) extended this to the general case. This method is known as the MetropolisHastings method for its progenitors. It requires the ability to calculate a function $f(\mathbf{x})$ which is proportional to the desired density $\mathbb{P}(\mathbf{x})$; in the Bayesian setting this function $f(\mathbf{x})$ can be taken to be the un-normalised posterior distribution. This allows us to avoid the (often difficult) task of calculating the integral required to normalise the posterior. To give some intuition, at time step $t$ with current value $\mathbf{x}_{t}$, the method works by generating a proposed new value $\mathbf{x}^{\prime}$ sampled from a "proposal distribution" $g\left(\mathbf{x}^{\prime} \mid \mathbf{x}_{t}\right)$. This proposal distribution is often taken to be a normal distribution with mean $\mathbf{x}_{t}$ for simplicity. We then calculate the "acceptance ratio"

$$
A=\min \left(\frac{f\left(\mathbf{x}^{\prime}\right)}{f\left(\mathbf{x}_{t}\right)} \frac{g\left(\mathbf{x}^{\prime} \mid \mathbf{x}_{t}\right)}{g\left(\mathbf{x}_{t} \mid \mathbf{x}^{\prime}\right)}, 1\right)=\min \left(\frac{\mathbb{P}\left(\mathbf{x}^{\prime}\right)}{\mathbb{P}\left(\mathbf{x}_{t}\right)} \frac{g\left(\mathbf{x}^{\prime} \mid \mathbf{x}_{t}\right)}{g\left(\mathbf{x}_{t} \mid \mathbf{x}^{\prime}\right)}, 1\right)
$$

and "accept" the proposed point with probability $A$. If we do not accept the proposed point, we set $\mathbf{x}_{t+1}=\mathbf{x}_{t}$. If the proposal distribution is symmetric, then the acceptance ratio simplifies to

$$
A=\min \left(\frac{f\left(\mathbf{x}^{\prime}\right)}{f\left(\mathbf{x}_{t}\right)}, 1\right)=\min \left(\frac{\mathbb{P}\left(\mathbf{x}^{\prime}\right)}{\mathbb{P}\left(\mathbf{x}_{t}\right)}\right)
$$

We can then see that, roughly, this criteria means that as $t$ increases, we are more likely to move to (or stay at) points that are more likely under the posterior likelihood.

We used a Gibbs-like version of the Metropolis-Hastings, in that we divided the sampling step up into a sampling step for each of the variables conditional on those variables previously sampled. This method derives from the work of Geman and Geman (1984); it works by sampling each individual component from the conditional distribution given the rest of the components. However, the intractability of sampling from the actual conditional distributions required us to use the Metropolis-Hastings style proposal-acceptance routine within each sampling sub-step. To expand on this, each sampling step was as follows:

1. Draw $h^{(*)}$ from a normal distribution left-truncated at 0 with mean $h^{(t)}$. Calculate the acceptance ratio

$$
A:=\frac{\mathbb{P}\left(h^{(*)} \mid \mathbf{X}, \mathbf{y}, \mathbf{H}^{(t)}, m^{(t)}, \boldsymbol{\mu}^{(t)}, \mathbf{v}^{(t)}\right)}{\mathbb{P}\left(h^{(t)} \mid \mathbf{X}, \mathbf{y}, \mathbf{H}^{(t)}, m^{(t)}, \boldsymbol{\mu}^{(t)}, \mathbf{v}^{(t)}\right)} \frac{\mathbb{P}\left(h^{(t)} \mid h^{(*)}\right)}{\mathbb{P}\left(h^{(*)} \mid h^{(t)}\right)}
$$

where the right-hand fraction is the ratio of probabilities from the truncatednormal proposal distribution. With probability $\min (1, A)$ accept $h^{(*)}$ as $h^{(t+1)}$, otherwise let $h^{(t+1)}:=h^{(t)}$
2. Draw $\mathbf{H}^{(*)}$ from an (uncorrelated) multivariate normal distribution with mean $\mathbf{H}^{(t)}$. Calculate the acceptance ratio

$$
A:=\frac{\mathbb{P}\left(\mathbf{H}^{(*)} \mid \mathbf{X}, \mathbf{y}, h^{(t+1)}, m^{(t)}, \boldsymbol{\mu}^{(t)}, \mathbf{v}^{(t)}\right)}{\mathbb{P}\left(\mathbf{H}^{(t)} \mid \mathbf{X}, \mathbf{y}, h^{(t+1)}, m^{(t)}, \boldsymbol{\mu}^{(t)}, \boldsymbol{v}^{(t)}\right)} \frac{\mathbb{P}\left(\mathbf{H}^{(t)} \mid \mathbf{H}^{(*)}\right)}{\mathbb{P}\left(\mathbf{H}^{(*)} \mid \mathbf{H}^{(t)}\right)}
$$

and accept with probability $\min (1, A)$
3. Draw $m^{(*)}$ from a normal distribution left-truncated at 0 with mean $m^{(t)}$. Calculate acceptance ratio

$$
A:=\frac{\mathbb{P}\left(m^{(*)} \mid \mathbf{X}, \mathbf{y}, h^{(t+1)}, \mathbf{H}^{(t+1)}, \boldsymbol{\mu}^{(t)}, \mathbf{v}^{(t)}\right)}{\mathbb{P}\left(m^{(t)} \mid \mathbf{X}, \mathbf{y}, h^{(t+1)}, \mathbf{H}^{(t+1)}, \boldsymbol{\mu}^{(t)}, \mathbf{v}^{(t)}\right)} \frac{\mathbb{P}\left(m^{(t)} \mid m^{(*)}\right)}{\mathbb{P}\left(m^{(*)} \mid m^{(t)}\right)}
$$

and accept $m^{(*)}$ with probability $\min (1, A)$.
4. Draw $\mu^{(*)}$ from an (uncorrelated) multivariate normal distribution with mean $\mu^{(t)}$. Calculate

$$
A:=\frac{\mathbb{P}\left(\boldsymbol{\mu}^{(*)} \mid \mathbf{X}, \mathbf{y}, h^{(t+1)}, \mathbf{H}^{(t+1)}, m^{(t+1)}, \mathbf{v}^{(t)}\right)}{\mathbb{P}\left(\boldsymbol{\mu}^{(t)} \mid \mathbf{X}, \mathbf{y}, h^{(t+1)}, \mathbf{H}^{(t+1)}, m^{(t+1)}, \boldsymbol{v}^{(t)}\right)} \frac{\mathbb{P}\left(\boldsymbol{\mu}^{(t)} \mid \boldsymbol{\mu}^{(*)}\right)}{\mathbb{P}\left(\boldsymbol{\mu}^{(*)} \mid \boldsymbol{\mu}^{(t)}\right)}
$$

accept $\boldsymbol{\mu}^{(*)}$ with probability $\min (1, A)$
5. For each value $y \in \mathscr{Y}$, draw an $\boldsymbol{v}_{y}^{(*)}$ from an uncorrelated multivariate normal distribution with mean $\boldsymbol{v}_{y}^{(t)}$. Calculate

$$
A:=\frac{\mathbb{P}\left(\boldsymbol{v}_{y}^{(*)} \mid \mathbf{X}, \mathbf{y}, h^{(t+1)}, \mathbf{H}^{(t+1)}, m^{(t+1)}, \boldsymbol{\mu}^{(t+1)}\right)}{\mathbb{P}\left(\boldsymbol{v} y^{(t)} \mid \mathbf{X}, \mathbf{y}, h^{(t+1)}, \mathbf{H}^{(t+1)}, m^{(t+1)}, \boldsymbol{\mu}^{(t+1)}\right)} \frac{\mathbb{P}\left(\boldsymbol{v}_{y}^{(t)} \mid \boldsymbol{v}_{y}^{(*)}\right)}{\mathbb{P}\left(\boldsymbol{v}_{y}^{(*)} \mid \boldsymbol{v}_{y}^{(t)}\right)}
$$

and accept with probability $\min (1, A)$.

In order to evaluate convergence, we use the Gelman-Rubin $\hat{R}$ statistic Gelman and Rubin (1992). This statistic requires us to carry out C Markov Chain Monte Carlo processes all started from different, over-dispersed initial values. By overdispersed, we mean selected from a distribution with significantly higher variance than their prior distribution. This helps to ensure exploring a larger range of the parameter space, leading to better estimation. After discarding a set number of iterations for burn-in, we calculate the following quantities from the $n$ samples each chain gives us.

$$
\begin{gathered}
W=\frac{1}{C} \sum_{j=1}^{C} s_{j}^{2} \\
B=\frac{n}{C-1} \sum_{j=1}^{C}\left(\bar{\theta}_{j}-\overline{\bar{\theta}}\right)^{2}
\end{gathered}
$$

Where $s_{j}$ is the variance in the $j$ th chain, $\bar{\theta}_{j}$ is the sample mean from chain $j$ and $\overline{\bar{\theta}}$ is the overall sample mean. For simplicity, we speak of a single parameter $\theta$, but this is extended to multiple parameters component-wise. These quantities give us estimates of the within-chain and between-chain variances, respectively. We then calculate an estimate of the variance of the so-called stationary distribution.

$$
\widehat{\operatorname{Var}(\theta)}=\left(1-\frac{1}{n}\right) W+\frac{1}{n} B
$$

Finally, we have

$$
\hat{R}=\sqrt{\frac{\widehat{\operatorname{Var}(\theta)}}{W}}
$$

The usual convention is to keep drawing samples (and possibly discarding more as burnin samples) until the value of $\hat{R}$ is approximately 1 (usually $<1.1$ ) for every parameter. Once we have that, we can pool our samples from each chain together and calculate the mean of each parameter. Asymptotically, this parameter mean is convergent to the expected value of that parameter over its true distribution.

Both the maximum a-posteriori (MAP) method and the Markov Chain Monte Carlo (MCMC) method were effective in estimating the parameters of the PoIR model. However, given that we are not as interested in the posterior distributions as we are in the posterior means, we will proceed to use the MAP as it is less computationally expensive.

### 2.4 Evaluation

The closest method to our formulated Poisson inverse regression is MNIR; in this section we will evaluate the performance of our method using MNIR as the benchmark. To


Figure 2.1: Plot of two directions recovered by POIR from the we8there data.
do this, we will use a test dataset used in Taddy (2013) and made available in the accompanying R package, namely the "we8there" dataset of restaurant reviews. The data, as made available in the "textir" package, is already preprocessed which makes it very convenient to work with. Figure 2.2 shows three directions recovered from the data using POIR, and Figure 2.3 shows the one direction MNIR is able to recover. The limitation of MNIR to one direction is due to us using only a one-dimensional response (the overall rating). MNIR is, as implemented in "textir" and as focused on in Taddy (2013), limited to recovering only as many directions as the dimension of the response variable. The way we have specified the POIR model leaves us able to estimate however many directions are appropriate.

We estimate the PoIR model using MAP, as we only need a point estimate. MNIR is estimated using the "textir" package method, detailed in Taddy (2013). In Figure 2.1 we show the results of estimating a two-dimensional PoIR model; a three-dimensional PoIR model is shown in Figure 2.2. These two figures are in the form of "pairs" plots, showing for each combination of two directions a scatterplot of the projected data (coloured according to the value of the response), the correlations between the two directions on a per-response-value level, as well as a one-dimensional density plot for each direction,


Figure 2.2: Plot of three directions recovered by POIR from the we8there data.
separated by response value. We also show the (one-dimensional) results from MNIR in Figure 2.3 in both a scatterplot of the direction against the rating, and as a density plot separated by the rating.

From Figure 2.1, we can see that two directions is not sufficient to distinguish the five response classes; in particular the three lowest ratings (1, 2, and 3 stars) are particularly difficult to distinguish from one another. It is no surprise, then, that the one-dimensional MNIR shown in Figure 2.3 also struggles significantly to distinguish the five response classes. On the other hand, Figure 2.2 shows the three-dimensional PoIR giving a more reasonable separation of the data.

As can be seen from the two figures, the extra dimensions we are able to extract using POIR give us significantly better separation of the five different response levels which are indicated using colours on both plots. One direction only, in the case of MNIR, leaves the different response levels largely inseparable; it would be difficult to identify correctly any more than some of the 1 -rated and some of the 5 -rated responses.


Figure 2.3: Scatterplot of the one MNIR direction recovered from the we8there data against the rating, and density of that direction separated by rating.

### 2.5 Conclusion

We have shown in Section 2.4 that our PoIR is an effective method for text data dimension reduction. In fact, from 2.2.1, we know that it is a sufficient dimension reduction. We have given two methods for estimation in Section 2.3, a method based on MAP, and one based on MCMC.

There are some interesting directions for future work. Firstly, the model for text data upon which PoIR is based is a very simple one which does not take into account potential correlation between terms. Secondly, the estimation algorithms given are not tuned for computational efficiency. As such, they are currently slower than the method provided in Taddy (2013) and then further improved in Taddy (2015) for MNIR. It should be possible to significantly decrease the computational time for the MAP estimation method by using a customised optimisation algorithm, rather than the "out of the box" method we used in our tests. It may, in fact, be possible to extend Taddy's work on distributed multinomial regression to the Poisson case, and adapt it for PoIR. For the MCMC method, there has been considerable research into improved algorithms for MCMC which we have not2.5. Conclusion31
taken advantage of in this work.

## — Chapter 3 -

## Healthcare Data

In this chapter, we will introduce a dataset which will be used as a "real data" example in Chapter 4 and Chapter 5. This dataset was kindly provided by Cardiff and Vale University Health Board.

This dataset was a by-product of a lexicon-based classifier built by the health board in order to classify letters sent by consultants at a large hospital to general practitioners and patients. Amongst other things, these letters either recommend further treatment, or a discharge from the out-patient service, but almost all of these letters are unlabelled. In order to better model patient pathways, the health board informatics team wanted to have labels assigned to each; given that they have over one million such letters, it was economically infeasible to label each of these by hand.

In order to deal with this problem, the health board developed a lexicon of just over 40,000 procedurally generated sentences, built from exemplar phrases and words related to either a follow-up or discharge decision. This lexicon could then be used with a nearest-neighbour classifier to predict the label for the letters. In order to validate this process, they had a number of letters manually labelled. This process also helped suggest further "seeds" for the lexicon generation process.

Some example sentences with their classification are included below.

Discharge "I DID NOT FEEL THERE WAS ANY NEED FOR HIM TO BE FOLLOWED UP"
Discharge "I DID NOT FEEL THERE WAS ANY NEED FOR HIM TO BE REVIEWED"

Discharge "I DID NOT FEEL THERE WAS ANY NEED FOR HIM TO BE SEEN AGAIN"
Follow-up "I WILL SEE HER AGAIN IN CLINIC [ ${ }^{\sim}$ CLINIC"]"
Follow-up "I WILL SEE HER AGAIN IN [ ${ }^{\text {CLINIC }}{ }^{\sim}$ ]"
Follow-up "I WILL SEE HER IN CLINIC [ ${ }^{\sim}{ }^{\text {CLINIC }}{ }^{\sim}$ ]"

Due to the procedural generation process, which produces many very similar "documents", we encountered difficulties applying any algorithms to the entire dataset. To combat this (with the added benefit of making experiments quicker to run), we worked with a subset of 300 examples from the lexicon.

The lexicon which was generated had a heavy imbalance in favour of follow-up examples. Because the methods in this work are not designed to cope with such a heavy imbalance as is present in the original dataset, this subset was randomly sampled to consist of 150 "discharge" sentences and 150 "follow-up" sentences.

In order to prepare the lexicon data for analysis with the methods developed in Chapters 4 and 5, we first stemmed the words using the R package "RTextTools". The stemmed documents were then converted into a document-term matrix using the same package.

## - Chapter 4 -

## Sparse Generalised Principal Component Analysis

In this chapter we introduce Sparse Generalised Principal Component Analysis (SGPCA), a dimension reduction technique developed from the Generalised Principal Component Analysis (GPCA) of Landgraf and Lee (2015b). We begin by a recollection of the vector-space model of text data and an explanation of how that motivates sparsity in Section 1.2.

We will introduce GPCA as an extension of Principal Component Analysis (PCA) which is suitable for data from any exponential family of distribution in Section 4.2, before defining our extension SGPCA in Section 4.3. This will involve an exploration of methods of penalisation (Section 4.3.1) and an efficient algorithm for estimation of the loadings matrix (Section 4.3.2).

In Sections 4.4 and 4.6 we will examine the efficacy of SGPCA in several ways. Firstly, we will conduct synthetic data experiments using a hidden factor model in Section 4.4.1, a two-class hidden factor model in Section 4.4.2 and a varying-noise hidden factor model to explore robustness against noise in Section 4.4.3. Section 4.6 will comprise of the application and analysis of SGPCA to a healthcare text dataset.

Finally, we will summarise the results of this chapter in Section 4.7.

### 4.1 Introduction

### 4.1.1 Text Data

Recalling from Section 1.2, in the vector-space model for text data we represent each document in a corpus as a row in a data matrix $\mathbf{X}=\left[x_{i j}\right]_{i=1, j=1}^{i=n, j=p}$. Each column of $\mathbf{X}$ represents a term (either a word or n-gram) which appears in the corpus. Then the element $x_{i j}$ is the number of times the $i^{\text {th }}$ document contains term $j$. Often, a suitable
model for the data is that each term has a Poisson distribution with some mean $\lambda_{j}$ for $j=1, \ldots, p$.

It is important for our future discussion to note two points. Firstly, that such data is not Gaussian; in fact, whilst the Gaussian distribution is a reasonable approximation to the Poisson distribution for large enough $\lambda$, it is rare in a text mining scenario for such values of $\lambda$ to occur. Secondly, across a collection of documents, we can often expect many terms to be common to most of the documents. We are often interested in either classifying or clustering documents, but these common terms are uninformative in these contexts. As such, we might expect that a dimension reduction across the corpus would not include contributions from these terms, thus motivating our belief in sparsity.

### 4.2 Generalised Principal Component Analysis

### 4.2.1 Generalised Principal Component Analysis Definition

Let us now return to the earlier definition of PCA in (1.3). In Landgraf and Lee (2015a) and Landgraf and Lee (2015b), the authors recognised that this squared reconstruction error is equivalent (up to a constant) to the deviance of a model where the observations are from a Gaussian distribution with known variance and a natural parameter matrix of the form

$$
\begin{equation*}
\widehat{\boldsymbol{\Theta}}=\mathbb{1}_{n} \boldsymbol{\mu}^{\mathrm{T}}+\left(\widetilde{\boldsymbol{\Theta}}-\mathbb{1}_{n} \boldsymbol{\mu}^{\mathrm{T}}\right) \mathbf{U U}^{\mathrm{T}} \tag{4.1}
\end{equation*}
$$

where $\widetilde{\boldsymbol{\Theta}}$ is the matrix of saturated natural parameters, $\mathbf{U} \in \mathbb{R}^{p \times k}, \mu \in \mathbb{R}^{p}$ and we constrain $\mathbf{U}^{\mathrm{T}} \mathbf{U}=\mathbb{I}$. This connection can be seen by recognising that the canonical link function for the Gaussian distribution is merely the identity function, so the saturated natural parameters are simply the observed values. Note that the matrix $\widehat{\boldsymbol{\Theta}}$ has rank $k$, compared to $\widetilde{\boldsymbol{\Theta}}$ which has rank $p$. It is precisely this low-rank approximation which can be used as a lower-dimensional analogue to the observed data. Once optimal $\boldsymbol{\mu}$ and $\mathbf{U}$ have been identified, $\widehat{\Theta}$ can be calculated.

In order to generalise this to exponential family data (binary in Landgraf and Lee (2015a), most exponential families in Landgraf and Lee (2015b)), we seek the devianceoptimal parameters U and $\boldsymbol{\mu}$. Here, the deviance refers to a measure of model quality, comparing one model (within the same family) to the saturated model (the model with natural parameters equal to the saturated natural parameters). Formally, the deviance is defined by the log of the likelihood ratio; this is equivalent to the difference of the loglikelihood. The deviance is non-negative, with minimum 0 achieved at the saturated model. We now prove a lemma regarding the form of the deviance under the GPCA model.

Lemma 4.2.1. Let $\mathrm{X} \in \mathbb{R}^{n \times p}$ be $n$ observations of a p-dimensional random vector $X=$ $\left[x_{1}, \ldots, x_{p}\right]^{\mathrm{T}}$ where $X_{j}$ has an exponential family distribution with log-partition function $b_{j}$ for $j=1, \ldots, p$. Denote by $\widetilde{\boldsymbol{\Theta}}$ the matrix of saturated natural parameters and let $\widehat{\boldsymbol{\Theta}}$ be the low-rank approximation with $i j^{\text {th }}$ component $\hat{\theta}_{i j}$. Then the deviance up to an additive constant and multiplication by a positive factor is given by

$$
\begin{equation*}
D(\mathbf{X} \mid \widehat{\boldsymbol{\Theta}}) \stackrel{+\times}{\propto} \sum_{i=1}^{n} \sum_{j=1}^{p}\left[b_{j}\left(\hat{\theta}_{i j}\right)-x_{i j} \hat{\theta}_{i j}\right] \tag{4.2}
\end{equation*}
$$

Proof. Properly, we have

$$
\begin{align*}
D(\mathbf{X} \mid \widehat{\boldsymbol{\Theta}}) & =-2(\log \mathbb{P}(\mathbf{X} \mid \widehat{\boldsymbol{\Theta}})-\log \mathbb{P}(\mathbf{X} \mid \widetilde{\boldsymbol{\Theta}}))  \tag{4.3}\\
& +\times  \tag{4.4}\\
& =\log \prod_{i=1}^{n} \prod_{j=1}^{p}\left[h_{j}\left(x_{i j}\right) \exp \left(x_{i j} \hat{\theta}_{i j}-b_{j}\left(\hat{\theta}_{i j}\right)\right)\right]  \tag{4.5}\\
& =-\sum_{i=1}^{n} \sum_{j=1}^{p}\left[\log h_{j}\left(x_{i j}\right)+x_{i j} \hat{\theta}_{i j}-b_{j}\left(\hat{\theta}_{i j}\right)\right]  \tag{4.6}\\
& \propto \sum_{i=1}^{n} \sum_{j=1}^{p}\left[b_{j}\left(\hat{\theta}_{i j}\right)-x_{i j} \hat{\theta}_{i j}\right] \tag{4.7}
\end{align*}
$$

Remark. From this point on, when we refer to $D(\mathbf{X} \mid \widehat{\boldsymbol{\Theta}})$ we will actually refer to the right hand side of (4.2) unless otherwise stated. From the point of view of optimisation, the two are entirely equivalent.

Definition 4.2.1. For some observed data $\mathbf{X} \in \mathbb{R}^{n \times p}$, the rank $k$ GPCA approximation is the pair $\mathbf{U}$ and $\mu$ satisfying

$$
\mathbf{U}, \boldsymbol{\mu}=\underset{\mathbf{U} \in \mathbb{R}^{p \times k}, \boldsymbol{\mu} \in \mathbb{R}^{p}}{\arg \min } D\left(\mathbf{X} \mid \mathbb{1}_{n} \boldsymbol{\mu}^{\mathrm{T}}+\left(\widetilde{\boldsymbol{\Theta}}-\mathbb{1}_{n} \boldsymbol{\mu}^{\mathrm{T}}\right) \mathbf{U} \mathbf{U}^{\mathrm{T}}\right) \quad \text { subject to } \mathbf{U}^{\mathrm{T}} \mathbf{U}=\mathbb{I}
$$

### 4.3 Sparse Generalised Principal Component Analysis

### 4.3.1 Penalisation

As the goal of SGPCA is to provide a method comparable to GPCA but with sparse loadings, the first step to defining it is to discuss penalisation in the context of penalised optimisation. We will be placing a penalty function $P(\mathbf{U})$ on the components of $\mathbf{U}$, and rather than optimising the deviance directly, we shall optimise the deviance plus the
penalty function, with the aim to jointly minimise both, finding a good compromise between sparsity and deviance-optimality. We will consider two well-used penalty functions in order to construct our penalty function, the Smoothly Clipped Absolute Deviation (SCAD) penalty and the $L^{1}$ penalty. Both of these have well-understood properties, and are often used in the sparsity context.

Remark. We also briefly investigated the $L^{2}$ penalty due to its frequent use in sparsifying methods (especially in conjuction with the $L^{1}$ penalty in the so-called elastic net (Zou and Hastie 2005)). Unfortunately, due to the semi-orthogonality constraint on $\mathbf{U}$, the $L^{2}$ penalty on matrices (often known as the Frobenius norm) would be always equal to $k$ on $\mathbf{U}$.

### 4.3.1.1 Smoothly Clipped Absolute Deviation Penalty

The SCAD penalty, due to Fan and Li (2001), is usually defined by its derivative

$$
\begin{equation*}
P_{\mathrm{S}}^{\prime}(\theta ; \lambda, a):=\lambda I(\theta \leq \lambda)+\frac{(a \lambda-\theta)_{+}}{a-1} I(\theta>\lambda) \tag{4.8}
\end{equation*}
$$

where $I(A)$ is the indicator function on the proposition $A$ and $(x)_{+}=\max \{x, 0\}$. An application of integration shows that the penalty function itself is

$$
P_{\mathrm{S}}(\theta ; \lambda, a)= \begin{cases}\lambda \theta & 0<\theta \leq \lambda  \tag{4.9}\\ -\frac{\theta^{2}-2 a \lambda \theta+\lambda^{2}}{2(a-1)} & \lambda<\theta \leq a \lambda \\ \frac{\left(a+1 \lambda^{2}\right.}{2} & a \lambda<\theta\end{cases}
$$

A graph of $P_{S}(|\theta|)$ is shown in Figure 4.1. Examination either of (4.9) or Figure 4.1 will reveal that within $\lambda$ of the origin the penalty decreases towards 0 linearly, more than $a \lambda$ from the origin the penalty is a (positive) constant, and between the two regions the penalty is a quadratic joining the two regions. With this in mind, one can understand the motivation behind the penalty - to drive coefficients close to 0 down to 0 , to penalise coefficients larger than a certain threshold only a constant amount, under the recognition that large enough coefficients are probably needed by the model, and to transition smoothly between the two behaviours. To illustrate the way in which the penalty changes with $\lambda$, in Figure 4.2 we show the penalty for three different values of $\lambda$.

We will be applying the SCAD penalty to $\mathbf{U}$ as follows:

$$
\begin{equation*}
P_{\mathrm{S}}(\mathbf{U} ; \lambda, a)=\sum_{i=1}^{n} \sum_{j=1}^{p} P_{\mathrm{S}}\left(\left|U_{i j}\right| ; \lambda, a\right) \tag{4.10}
\end{equation*}
$$



Figure 4.1: The (symmetrised) SCAD penalty, with $\lambda=0.1$ and $a=3.7$. In the region $[-\lambda, \lambda]$, the function is shown dashed; in the region $(-\infty,-a \lambda) \cup(a \lambda, \infty)$ it is shown dotted; the quadratic joining sections are shown solid.


Figure 4.2: The (symmetrised) SCAD penalty shown for $\lambda=0.1$ (solid line), $\lambda=0.15$ (dashed line), and $\lambda=0.2$ (dotted line).

Note that the SCAD penalty takes two parameters, $\lambda$ and $a$. The latter is usually taken to be 3.7 due to an argument in the original paper, but the former must be chosen on a case-by-case basis. We will usually take it to be a small value on the order of $10^{-1}$ or smaller.

### 4.3.1.2 $\quad L^{1}$ Penalty

The $L^{1}$ penalty on $\mathbf{U}$ is simply $\|\mathbf{U}\|_{1}=\sum_{i=1}^{n} \sum_{j=1}^{p}\left|U_{i j}\right|$. Its simplicity in comparison with the SCAD penalty is obvious, but it does come with the disadvantage of reduced flexibility and power. We will investigate whether this is truly a disadvantage in our application, or whether its performance is sufficient.

### 4.3.1.3 Total Penalty Function

We define the total penalty function

$$
\begin{equation*}
P\left(\mathbf{U} ; a, \lambda, \lambda_{\mathrm{S}}, \lambda_{L}\right):=\lambda_{\mathrm{S}} P_{\mathrm{S}}(\mathbf{U} ; a, \lambda)+\lambda_{L}\|\mathbf{U}\|_{1} \tag{4.11}
\end{equation*}
$$

Note that this penalty function has two additional parameters, $\lambda_{S}$ and $\lambda_{L}$; these are used both to control the relative importance of each penalty function to each other, and the relative importance to the deviance. A linear combination of the two penalties was chosen in order to make their comparison easier; in practice we would not advise using both penalties, although we will later compare the performance of this "doubly penalised" model against the singly penalised versions for thoroughness. SGPCA is fairly sensitive to appropriate specification of these parameters to manage the desire for sparsity against the need for a deviance-optimal fit. We will discuss some heuristic strategies for choosing appropriate values in Sec 4.7, along with the need to develop a general data-driven method for their selection.

### 4.3.1.4 Definition of SGPCA

We can now define the SGPCA approximation.
Definition 4.3.1. Let $\mathbf{X} \in \mathbb{R}^{n \times p}$ be a matrix of $n$ observations of $p$ variables and let the $j^{\text {th }}$ variable have an exponential family distribution with $\log$-partition function $b_{j}$. Let $\widetilde{\Theta} \in \mathbb{R}^{n \times p}$ denote the matrix of saturated natural parameters. Then, letting $D$ and $P$ have the definitions given in (4.2) (4.11). Then the SGPCA approximation for $\mathbf{X}$ is the pair $\mathbf{U} \in \mathbb{R}^{p \times k}$ and $\boldsymbol{\mu} \in \mathbb{R}^{p}$ satisfying

$$
\mathbf{U}, \boldsymbol{\mu}=\underset{\mathbf{U} \in \mathbb{R}^{p \times k}, \mu \in \mathbb{R}^{p}}{\arg \min }\left[D(\mathbf{X} \mid \widehat{\boldsymbol{\Theta}})+P\left(\mathbf{U} ; a, \lambda, \lambda_{\mathrm{S}}, \lambda_{L}\right)\right]
$$

### 4.3.2 Estimation

Now that we have defined the SGPCA approximation, it remains to discuss how to estimate the optimum $\mathbf{U}$ and $\mu$. Unfortunately, this is not an easy task due to the singularities of the penalty function at 0 components and the semi-orthogonality condition on $\mathbf{U}$.

This is further complicated by the non-convexity of the objective function. In general, this means that finding the global minimum is not guaranteed. In practice, we have not found this to be particularly problematic, as local minima have been sufficiently good for our purposes.

In order to optimise the objective function, we would like to be able to use gradientbased methods, but the singularities of the penalty function preclude doing so directly. To avoid this problem, we will use the method of Hunter and Li (2005) to derive an appropriate, differentiable and non-singular majoriser of the penalties. To summarise, they showed that, for penalties of the form $P(\theta)=\lambda p\left(\left|\theta_{i}\right|\right)$ where $p$ is non-decreasing, convex and have $\lim _{\theta \downarrow 0} p^{\prime}(\theta)<\infty$ the function

$$
\Phi_{\theta_{0}}(\theta):=\lambda p(|\theta|)+\frac{\left(\theta^{2}-\theta_{0}^{2}\right) \lambda p^{\prime}\left(\left|\theta_{0}\right|+\right)}{2\left|\theta_{0}\right|}
$$

is a majoriser. That is, $\Phi_{\theta_{0}} \geq \lambda p(|\theta|)$ for all $\theta$ with equality when $\theta=\theta_{0}$. Note that here $p^{\prime}\left(\left|\theta_{0}\right|+\right)$ means the right hand limit of the function at $\left|\theta_{0}\right|$.

As the SCAD component function is convex, and our total SCAD penalty is merely a weighted sum (with non-negative weights) of the component functions evaluated at the absolute value of the components, we can apply this result to majorise the SCAD term of our penalty function. Under the notation of Hunter and Li, the $L^{1}$ penalty is also of this form, as the corresponding $p$ function is merely the identity function which is convex. Thus, we can majorise our total penalty function $P\left(\mathbf{U} ; a, \lambda, \lambda_{S}, \lambda_{L}\right)$ by an appropriately calculated combination of $\Phi_{\theta_{0}}$ functions.

Lemma 4.3.1. The function $\Phi^{(t-1)}$ defined below is a majoriser of the total penalty function.

$$
\begin{align*}
\Phi^{(t-1)}:= & \sum_{i=1}^{n} \sum_{j=1}^{p}\left\{\lambda_{L}\left[\left|u_{i j}^{(t-1)}\right|+\frac{u_{i j}^{2}-\left(u_{i j}^{(t-1)}\right)^{2}}{2\left|u_{i j}^{(t-1)}\right|}\right]+\right. \\
& \left.\lambda_{\mathrm{S}}\left[P_{\mathrm{S}}\left(\left|u_{i j}^{(t-1)}\right| ; \lambda, a\right)+\frac{\left(u_{i j}^{2}-\left(u_{i j}^{(t-1)}\right)^{2}\right) P_{\mathrm{S}}^{\prime}\left(\left|u_{i j}^{(t-1)}\right|+; \lambda, a\right)}{2\left|u_{i j}^{(t-1)}\right|}\right]\right\} \tag{4.12}
\end{align*}
$$

Proof. An application of the result of Hunter and Li (2005).

However, our majoriser is not defined everywhere; it contains a singularity at all $U_{0}$ with at least one component equal to 0 . Fortunately, we can remove this singularity with a perturbed version of the majoriser:

$$
\begin{align*}
\Phi^{(t-1)}:= & \sum_{i=1}^{n} \sum_{j=1}^{p}\left\{\lambda_{L}\left[\left|u_{i j}^{(t-1)}\right|+\frac{u_{i j}^{2}-\left(u_{i j}^{(t-1)}\right)^{2}}{2\left(\left|u_{i j}^{(t-1)}\right|+\varepsilon\right)}\right]+\right. \\
& \left.\lambda_{\mathrm{S}}\left[P_{\mathrm{S}}\left(\left|u_{i j}^{(t-1)}\right| ; \lambda, a\right)+\frac{\left(u_{i j}^{2}-\left(u_{i j}^{(t-1)}\right)^{2}\right) P_{\mathrm{S}}^{\prime}\left(\left|u_{i j}^{(t-1)}\right|+; \lambda, a\right)}{2\left(\left|u_{i j}^{(t-1)}\right|+\varepsilon\right)}\right]\right\} \tag{4.13}
\end{align*}
$$

This new perturbed majoriser is not, in fact, a majoriser for our penalty function except in the limit $\varepsilon \downarrow 0$, but we will refer to it as such under this understanding. Hunter and Li showed that such perturbed majorisers do majorise a perturbed version of the problem, differing by a term with a multiplicative factor of $\varepsilon$. As such, we can hope that with small enough $\varepsilon$, solving the perturbed problem will provide a good approximation to solving the original problem.

Now we have an expression for a (perturbed) majoriser which is defined everywhere and is differentiable, we can proceed to use gradient methods to find the optimum $\mathbf{U}$ and $\mu$. We will treat the two quantities separately, in a batched coordinate-descent style. As optimising over $\mu$ is fairly simple, we will concentrate on how to perform optimisation over $\mathbf{U}$. We need to preserve the left semi-orthogonality condition; we could use a Langragian method, but we prefer to use a more direct method from Wen and Yin (2013). Given a feasible point $\mathbf{U}$ and the gradient at that point $\mathbf{G}$, they first define the matrix $\mathbf{A}:=\mathbf{G U}^{\mathrm{T}}-\mathbf{U G}^{\mathrm{T}}$ (in notation consistent with our use) which is skewsymmetric, that is, $\mathbf{A}^{T}=-\mathbf{A}$. Next, they define $\mathbf{Y}(\tau)$ by the Crank-Nicholson-like scheme $\mathbf{Y}(\tau)=\mathbf{U}-\frac{\tau}{2} \mathbf{A}(\mathbf{U}+\mathbf{Y}(\tau))$. The solution to this can be given explicitly as

$$
\begin{equation*}
\mathbf{Y}(\tau)=\left(\mathbb{I}+\frac{\tau}{2} \mathbf{A}\right)^{-1}\left(\mathbb{I}-\frac{\tau}{2} \mathbf{A}\right) \mathbf{U} \tag{4.14}
\end{equation*}
$$

Wen and Yin show that $\mathbf{Y}(\tau)$ possesses several very important features: for all $\tau \geq 0$ we have that $(\mathbf{Y}(\tau))^{\mathrm{T}} \mathbf{Y}(\tau)=\mathbf{U}^{\mathrm{T}} \mathbf{U}$, it is smooth in $\tau, \mathbf{Y}(0)=\mathbf{U}$, and finally $\mathbf{Y}(\tau)$ defines a descent path for $\tau \geq 0$. It is important to note that, while this function does involve two matrix inversions, an efficient algorithm for calculating them is given in Wen and Yin (2013) which we recommend.

Now all that is needed is the appropriate gradients for the objective function which we will denote by $S$.

Lemma 4.3.2. The gradient of the objective function

1. with respect to $U_{k l}$ is

$$
\begin{align*}
& \frac{\partial S}{\partial \mathbf{U}_{k l}}=\sum_{i=1}^{n} \sum_{j=1}^{p}\left\{\left(b_{j}^{\prime}\left(\mu_{j}+\left[\mathbf{U U}^{\mathrm{T}}\left(\tilde{\boldsymbol{\theta}}_{i}-\mu\right)\right]_{j}\right)-x_{i j}\right)\right. \\
& \left.\times\left(\delta_{k j} \mathbf{U}_{[l]}^{\mathrm{T}}\left(\tilde{\theta}_{i}-\mu\right)+\mathbf{U}_{j l}\left(\tilde{\theta}_{i k}-\mu_{k}\right)\right)\right\} \\
& \quad+\frac{\left(\lambda_{L}+\lambda_{S} P_{S}^{\prime}\left(\left|\mathbf{U}_{k l}^{(t-1)}\right|+; \lambda, a\right)\right) \mathbf{U}_{k l}}{\varepsilon+\left|\mathbf{U}_{k l}^{(t-1)}\right|} \tag{4.15}
\end{align*}
$$

2. with respect to $\mu_{k}$ is

$$
\begin{equation*}
\frac{\partial S}{\partial \mu_{k}}=\sum_{i=1}^{n} \sum_{j=1}^{p}\left(b_{j}^{\prime}\left(\mu_{j}+\left[\mathbf{U U}^{\mathrm{T}}\left(\tilde{\boldsymbol{\theta}}_{i}-\mu\right)\right]_{j}\right)-x_{i j}\right)\left(\delta_{j k}-\left[\mathbf{U U}^{\mathrm{T}}\right]_{j k}\right) \tag{4.16}
\end{equation*}
$$

Proof. First, we calculate the derivative of the majorised penalty with respect to the kl element of $\mathbf{U}$, i.e. $\frac{\partial}{\partial \mathbf{U}_{r s}} M_{P}^{\varepsilon}\left(\mathbf{U} \mid \mathbf{U}^{(t-1)}\right)$

$$
\begin{aligned}
\frac{\partial M_{P}^{\varepsilon}}{\partial \mathbf{U}_{k l}}= & \frac{\partial}{\partial \mathbf{U}_{k l}} \sum_{i=1}^{n} \sum_{j=1}^{p}\left\{\lambda_{L}\left|\mathbf{U}_{i j}^{(t-1)}\right|+\lambda_{\mathrm{S}} P_{\mathrm{S}}\left(\left|\mathbf{U}_{i j}^{(t-1)}\right|\right)\right. \\
& \left.+\frac{\left[\mathbf{U}_{i j}^{2}-\left(\mathbf{U}_{i j}^{(t-1)}\right)^{2}\right]\left[\lambda_{L}+\lambda_{\mathrm{S}} P_{\mathrm{S}}^{\prime}\left(\left|\mathbf{U}_{i j}^{(t-1)}\right|\right)\right]}{2\left[\varepsilon+\left|\mathbf{U}_{i j}^{(t-1)}\right|\right]}\right\} \\
= & \sum_{i=1}^{n} \sum_{j=1}^{p}\left\{\frac{\partial}{\partial \mathbf{U}_{k l}}\left[\mathbf{U}_{i j}^{2}-\left(\mathbf{U}_{i j}^{(t-1)}\right)^{2}\right] \frac{\lambda_{L}+\lambda_{\mathrm{S}} P_{\mathrm{S}}^{\prime}\left(\left|\mathbf{U}_{i j}^{(t-1)}\right|\right)}{2\left[\varepsilon+\left|\mathbf{U}_{i j}^{(t-1)}\right|\right]}\right\} \\
= & \frac{2 \mathbf{U}_{k l}\left(\lambda_{L}+\lambda_{\mathrm{S}} P_{\mathrm{S}}^{\prime}\left(\left|\mathbf{U}_{k l}^{(t-1)}\right|\right)\right)}{2\left[\varepsilon+\left|\mathbf{U}_{k l}^{(t-1)}\right|\right]}=\frac{\left.\mathbf{U}_{k l}\left(\lambda_{L}+\lambda_{\mathrm{S}} P_{\mathrm{S}}^{\prime}| | \mathbf{U}_{k l}^{(t-1)} \mid\right)\right)}{\left[\varepsilon+\left|\mathbf{U}_{k l}^{(t-1)}\right|\right]}
\end{aligned}
$$

where we suppress the arguments $P_{\mathrm{S}}$ for notational clarity.
Next, we calculate the gradient of the deviance function with respect to $\mathbf{U}_{k l}$

$$
\begin{aligned}
\frac{\partial D}{\partial \mathbf{U}_{k l}} & =\frac{\partial}{\partial \mathbf{U}_{k l}}\left\{\sum_{i=1}^{n} \sum_{j=1}^{p}\left\{b_{j}\left(\mu_{j}+\left[\mathbf{U U}^{\mathrm{T}}\left\{\tilde{\boldsymbol{\theta}}_{i}-\boldsymbol{\mu}\right\}\right]_{j}\right)-x_{i j}\left\{\mu_{j}+\left[\mathbf{U U}^{\mathrm{T}}\left\{\tilde{\boldsymbol{\theta}}_{i}-\boldsymbol{\mu}\right\}\right]_{j}\right\}\right\}\right. \\
& =\sum_{i=1}^{n} \sum_{j=1}^{p}\left\{\left\{b_{j}^{\prime}\left(\mu_{j}+\left[\mathbf{U U}^{\mathrm{T}}\left\{\tilde{\boldsymbol{\theta}}_{i}-\boldsymbol{\mu}\right\}\right]_{j}\right)-x_{i j}\right\} \frac{\partial}{\partial \mathbf{U}_{k l}}\left[\mu_{j}+\left[\mathbf{U U}^{\mathrm{T}}\left\{\tilde{\boldsymbol{\theta}}_{i}-\boldsymbol{\mu}\right\}\right]_{j}\right]\right\} \\
& =\sum_{i=1}^{n} \sum_{j=1}^{p}\{\left\{b_{j}^{\prime}\left(\mu_{j}+\left[\mathbf{U U}^{\mathrm{T}}\left\{\tilde{\boldsymbol{\theta}}_{i}-\boldsymbol{\mu}\right\}\right]_{j}\right)-x_{i j}\right\}\{\underbrace{\delta_{k j} \mathbf{U}_{[l]}^{\mathrm{T}}\left[\tilde{\boldsymbol{\theta}}_{i}-\boldsymbol{\mu}\right]}_{(*)}+\underbrace{\mathbf{U}_{j l}\left[\tilde{\theta}_{i k}-\mu_{k}\right]}_{(\hat{*})}\}\}
\end{aligned}
$$

Where ( $*$ ) comes from the terms quadratic in $\mathbf{U}_{k l}$ and $(\uparrow)$ is from the terms linear in $\mathbf{U}_{k l}$. Combining the two expressions, we get (4.15).

Next, we calculate the derivative of the objective function with respect to $\mu_{k}$. First, note that the majorised penalty does not depend on $\mu$, so we need only calculate the gradient of the deviance function.

$$
\begin{aligned}
& \frac{\partial}{\partial \mu_{k}}\left\{\sum_{i=1}^{n} \sum_{j=1}^{p}\left\{b_{j}\left(\mu_{j}+\left[\mathbf{U U}^{\mathrm{T}}\left\{\tilde{\boldsymbol{\theta}}_{i}-\boldsymbol{\mu}\right\}\right]_{j}\right)-x_{i j}\left\{\mu_{j}+\left[\mathbf{U U}^{\mathrm{T}}\left\{\tilde{\boldsymbol{\theta}}_{i}-\boldsymbol{\mu}\right\}\right]_{j}\right\}\right\}\right. \\
= & \sum_{i=1}^{n} \sum_{j=1}^{p}\left\{\left\{b_{j}^{\prime}\left(\mu_{j}+\left[\mathbf{U U}^{\mathrm{T}}\left\{\tilde{\boldsymbol{\theta}}_{i}-\boldsymbol{\mu}\right\}\right]_{j}\right)-x_{i j}\right\} \frac{\partial}{\partial \mu_{k}}\left[\mu_{j}+\left[\mathbf{U U}^{\mathrm{T}}\left\{\tilde{\boldsymbol{\theta}}_{i}-\boldsymbol{\mu}\right\}\right]_{j}\right]\right\} \\
= & \sum_{i=1}^{n} \sum_{j=1}^{p}\left\{\left\{b_{j}^{\prime}\left(\mu_{j}+\left[\mathbf{U} \mathbf{U}^{\mathrm{T}}\left\{\tilde{\boldsymbol{\theta}}_{i}-\boldsymbol{\mu}\right\}\right]_{j}\right)-x_{i j}\right\}\left\{\delta_{j k}-\frac{\partial}{\partial \mu_{k}}\left[\mathbf{U U}^{\mathrm{T}} \boldsymbol{\mu}\right]_{j}\right\}\right\} \\
= & \sum_{i=1}^{n} \sum_{j=1}^{p}\left\{\left\{b_{j}^{\prime}\left(\mu_{j}+\left[\mathbf{U U ^ { \mathrm { T } }}\left\{\tilde{\boldsymbol{\theta}}_{i}-\boldsymbol{\mu}\right\}\right]_{j}\right)-x_{i j}\right\}\left\{\delta_{j k}-\left[\mathbf{U U}^{\mathrm{T}}\right]_{j k}\right\}\right\}
\end{aligned}
$$

Remark. Although the gradients are given in Lemma 4.3.2, in practice we approximate the gradients by finite differences for reasons of computational efficiency.

In Algorithm 1 we summarise the procedure for calculating the estimated $\mathbf{U}$ and $\mu$ for SGPCA. In our implementation, we find the minimising value of $\mu$ using the ортім function in R (R Foundation for Statistical Computing, Vienna 2011). We find the appropriate value of $\tau$ by the line search algorithm of Nocedal and Wright (1999); we do not give explanation of this method here as the exact implementation of this aspect of the algorithm is fairly arbitrary, any line search algorithm would be sufficient. The algorithm of Nocedal and Wright is chosen upon the recommendation in Wen and Yin (2013).

A naive implementation of this algorithm tends to be computationally expensive, but with suitable care taken with evaluating the necessary functions it can be evaluated in a reasonable time-frame. In particular, we recommend using the efficient algorithm for the matrix inverses in $\mathbf{Y}(\tau)$. We have also found success using an automatic differentiation scheme to calculate both the objective function and its gradient simultaneously, though we do not yet have a concrete recommendation for this.

Remark. We remark that this algorithm for estimation does not come with any guarantees of convergence to the global minimum. Fortunately, our experience in practice
is that starting at the recommended initial values leads to convergence to a minimum which is sufficiently good.

### 4.4 Synthetic Data Examples

We will now investigate the performance of SGPCA in three situations using synthetic data; firstly, we will explore its performance on some synthetic data generated by a hidden factor model and evaluate the ability of SGPCA to identify directions corresponding to those hidden factors, secondly we will perform a similar analysis but using data drawn from two different hidden factor models and test the ability of SGPCA to find directions which can differentiate the two classes of data, finally we will test the robustness of the SGPCA loadings to noise by repeating the analysis from the first synthetic model but with varying levels of noise.

Throughout this analysis, we will use the following distribution for noise, which we will denote by $\mathcal{E}$ and observations from it by $\epsilon$.

$$
\begin{equation*}
\mathcal{E}(\eta):=\mathcal{P} \times(-1)^{\mathcal{B}}, \quad \mathcal{P} \sim \operatorname{Po}(\eta), \quad \mathcal{B} \sim \operatorname{Bernoulli}(1 / 2) \tag{4.17}
\end{equation*}
$$

That is, it is a Poisson random variable of mean $\eta$ multiplied with equal probability by 1 or -1 .

We will differentiate between three types of SGPCA in this and the following section: $L^{1}$ penalised SGPCA, SCAD penalised SGPCA and the combined SGPCA. The first will denote the case where $\lambda_{\mathrm{S}}=0$, the second will denote the case where $\lambda_{L}=0$ and the third will denote the case where both $\lambda_{S}$ and $\lambda_{L}$ are non-zero. This will allow us to analyse the performance of each penalty together and in isolation.

### 4.4.1 "Classless" Data Exploration

The hidden factor model we will use to test the performance of SGPCA will be constructed around three hidden factors, $\mathcal{V}_{1}, \mathcal{V}_{2}$ and $\mathcal{V}_{3}$ which will be a linear combination of the previous two. In particular,

$$
\mathcal{V}_{1} \sim \operatorname{Po}(25) \quad \mathcal{V}_{2} \sim \operatorname{Po}(30) \quad \mathcal{V}_{3}=1 \mathcal{V}_{1}+3 \mathcal{V}_{2}
$$

Each observations $\mathbf{x}_{i}$ will be constructed by drawing $v_{1 i}, v_{2 i}$ and $v_{3 i}$ from $\mathcal{V}_{1}, \nu_{2}$ and $\nu_{3}$ respectively, then denoting the $j^{\text {th }}$ component of $\mathbf{x}_{i}$ by $x_{i j}$, we have

$$
x_{i j}=v_{1 i}+\epsilon_{i j}, j \in\{1,2,3,4\} \quad x_{i j}=v_{2 i}+\epsilon_{i j}, j \in\{5,6,7,8\} \quad x_{i j}=v_{3 i}+\epsilon_{i j}, j \in\{9,10\}
$$

where each $\epsilon_{i j} \stackrel{\text { i.i.d }}{\sim} \mathcal{E}(2)$. We drew 100 observations, and performed $L^{1}$ SGPCA with $\lambda_{L}=10^{7}$, SCAD SGPCA with $\lambda=0.1$ and $\lambda_{\mathrm{S}}=10^{6}$, and the combined penalty SGPCA with $\lambda_{L}=10^{6}, \lambda_{S}=10^{6}$ and $\lambda=0.05$. We also performed PCA, Sparse Principal Component Analysis (SPCA), Robust Principal Component Analysis (RPCA), Nonnegative Matrix Factorisation (NMF) and Latent Dirichlet Allocation (LDA). The first loadings/directions from all algorithms are shown in Table 4.1, and the second in Table 4.2.

Analysis of Table 4.1 shows that for each of the three SGPCA variants, the first loading picks a direction corresponding primarily with the second hidden factor, and secondarily with the third hidden factor. Picking out both of these hidden factors together is to be expected, as these two factors are very highly correlated. On the other hand, GPCA's first loading does not strongly identify a direction associated with any hidden factor. PCA, SPCA, SPPCA and RPCA all have quite similar performance, all finding a direction most associated with the third hidden factor, with some contribution from the second. NMF finds a direction primarily associated with the third hidden factor, with small contribution from the second hidden factor and slightly larger contribution from the first hidden factor. LDA primarily identifies the third hidden factor.

Looking at the second loadings/directions in Table 4.2, we see that all three SGPCA variants strongly identify the first hidden factor. GPCA returns a similar loading, but with smaller coefficients for the components of the first hidden factor, and fairly large coefficients (of opposite sign) for the second. It also includes contributions from the third hidden factor. PCA returns fairly similar loadings to GPCA, with slightly larger contributions from the first hidden factor and slightly smaller from the second and third. SPCA, SPPCA and RPCA perform very simiarly to PCA. NMF, on the other hand, produces a loading similar in characteristics to its first loading, strongly identifying the third hidden factor with some contribution from the second. LDA also does not identify anything more than its first direction, finding primarily the third hidden factor.

Of all the algorithms considered, the performance of SGPCA is the best. Its first loading identifies the second hidden factor and the strongly correlated third hidden factor, and the second identifies the first hidden factor. Although GPCA, PCA, SPCA and RPCA provide similar identification in the second loading, they all include larger contributions from the second hidden factor, somewhat reducing interpretability.

### 4.4.2 Classed Synthetic Data Exploration

Although SGPCA is not a supervised dimension reduction method, we believe it should be capable of identifying directions suitable for differentiating between classes. To investigate this, we perform a similar study to that in Section 4.4.1, but drawing 100

Data: X, tol, $\lambda_{S}, \lambda, a, \lambda_{L}$
Result: U, $\mu$
calculate saturated natural parameters matrix $\widetilde{\boldsymbol{\Theta}}$;
set $\mu^{(0)}$ equal to the column means of $\widetilde{\boldsymbol{\Theta}}$;
set $\mathbf{U}^{(0)}$ equal to the first $k$ eigenvalues of $\widetilde{\boldsymbol{\Theta}}-\mathbb{1} \boldsymbol{\mu}^{\mathrm{T}}$;
set $t=1$;
while $\left\|\mathbf{U}^{(t)}-\mathbf{U}^{(t-1)}\right\|_{2}>$ tol or $\left\|\boldsymbol{\mu}^{(t)}-\boldsymbol{\mu}^{(t-1)}\right\|_{2}>$ tol or $t=1$ do
set $\mu^{(t)}$ to the minimiser of the objective function w.r.t. $\mu$;
set $\mathbf{G}$ to the gradient of the objective function w.r.t. $\mathbf{U}$;
set $\mathbf{A}=\mathbf{G}\left(\mathbf{U}^{(t-1)}\right)^{\mathrm{T}}-\mathbf{U}^{(t-1)} \mathbf{G}^{\mathrm{T}}$;
define function $\mathbf{Y}(\tau)=\left(\mathbb{I}+\frac{\tau}{2} \mathbf{A}\right)^{-1}\left(\mathbb{I}-\frac{\tau}{2} \mathbf{A}\right) \mathbf{U}^{(t-1)}$;
set $\tau^{*}$ to the value of $\tau$ which minimises the objective function along the path of $\mathbf{Y}(\tau)$, using $\mu^{(t)}$ as the value of $\mu$;
set $\mathbf{U}^{(t)}=\mathbf{Y}\left(\tau^{*}\right)$;
set $t=t+1$;
end
Algorithm 1: SGPCA Estimation Procedure

| $L^{1}$ | SCAD | Both | GPCA | PCA | SPCA | RPCA | NMF | LDA | SPPCA |
| ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: |
| 0.085 | 0.074 | 0.107 | -0.283 | 0.045 | 0.000 | -0.022 | 0.074 | 0.203 | -0.195 |
| 0.094 | 0.095 | 0.094 | -0.306 | 0.051 | 0.002 | -0.024 | 0.074 | 0.178 | -0.198 |
| 0.055 | 0.058 | 0.049 | -0.290 | 0.040 | 0.000 | -0.013 | 0.069 | 0.118 | -0.191 |
| 0.045 | 0.052 | 0.032 | -0.284 | 0.034 | 0.026 | -0.019 | 0.068 | 0.243 | -0.189 |
| 0.418 | 0.426 | 0.396 | -0.336 | 0.193 | 0.132 | -0.205 | 0.186 | 0.085 | -0.275 |
| 0.440 | 0.434 | 0.461 | -0.334 | 0.207 | 0.133 | -0.212 | 0.194 | 0.180 | -0.282 |
| 0.420 | 0.418 | 0.433 | -0.334 | 0.201 | 0.123 | -0.198 | 0.190 | 0.229 | -0.279 |
| 0.445 | 0.447 | 0.441 | -0.358 | 0.209 | 0.160 | -0.219 | 0.189 | 0.141 | -0.283 |
| 0.342 | 0.341 | 0.330 | -0.311 | 0.641 | 0.668 | -0.639 | 0.647 | 0.805 | -0.517 |
| 0.345 | 0.345 | 0.340 | -0.316 | 0.646 | 0.691 | -0.645 | 0.645 | 0.307 | -0.520 |

Table 4.1: First loading/direction for synthetic data, using $L^{1}$ penalised SGPCA, SCAD penalised SGPCA, the combined penalty SGPCA, GPCA, PCA, SPCA, RPCA, NMF, LDA and Sparse Probabilistic Principal Component Analysis (SPPCA)

| $L^{1}$ | SCAD | Both | GPCA | PCA | SPCA | RPCA | NMF | LDA | SPPCA |
| ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: |
| 0.470 | 0.452 | 0.493 | -0.373 | -0.458 | -0.485 | -0.443 | 0.217 | 0.059 | 0.408 |
| 0.481 | 0.475 | 0.490 | -0.380 | -0.466 | -0.483 | -0.457 | 0.218 | 0.082 | 0.413 |
| 0.511 | 0.517 | 0.505 | -0.428 | -0.496 | -0.526 | -0.519 | 0.227 | 0.137 | 0.428 |
| 0.505 | 0.522 | 0.484 | -0.429 | -0.493 | -0.479 | -0.486 | 0.228 | 0.028 | 0.427 |
| -0.082 | -0.084 | -0.068 | 0.270 | 0.133 | 0.065 | 0.148 | 0.126 | 0.209 | 0.236 |
| -0.106 | -0.108 | -0.094 | 0.299 | 0.175 | 0.111 | 0.164 | 0.123 | 0.133 | 0.259 |
| -0.082 | -0.070 | -0.098 | 0.265 | 0.130 | 0.047 | 0.096 | 0.135 | 0.095 | 0.235 |
| -0.077 | -0.079 | -0.062 | 0.277 | 0.131 | 0.076 | 0.152 | 0.126 | 0.163 | 0.235 |
| 0.018 | 0.015 | -0.006 | 0.143 | -0.017 | 0.000 | -0.070 | 0.608 | 0.412 | 0.175 |
| 0.022 | 0.019 | 0.008 | 0.141 | -0.036 | 0.000 | -0.056 | 0.606 | 0.839 | 0.185 |

Table 4.2: Second loading/direction for synthetic data, using $L^{1}$ penalised SGPCA, SCAD penalised SGPCA, the combined penalty SGPCA, GPCA, PCA, SPCA, RPCA, NMF and LDA
observations each from two different hidden factor models. The first 100 observations have hidden factors

$$
\mathcal{V}_{1} \sim \operatorname{Po}(25) \quad \mathcal{V}_{2} \sim \operatorname{Po}(25) \quad \mathcal{V}_{3}=1 \mathcal{V}_{1}+3 \mathcal{V}_{2}
$$

and the second 100 observations have hidden factors

$$
\nu_{1} \sim \operatorname{Po}(25) \quad V_{2} \sim \operatorname{Po}(35) \quad V_{3}=2 \mathcal{V}_{1}+1 \mathcal{V}_{2}
$$

We construct all of the observations by the same method as before. We then perform the same algorithms as before, with the addition this time of Multinomial Inverse Regression (MNIR), using as a response the class identifier 0 for the first 100 observations and 1 for the second 100. The loadings/directions for all algorithms are given in Table 4.3.

The primary differentiating factors between the two classes of data are the second hidden factor and the third hidden factor. As such, the strong identification of the second hidden factor by the loading of the $L^{1}$ penalised SGPCA is excellent. Interestingly, the SCAD and combined penalty SGPCA loadings are somewhat confused, suggesting that the $L^{1}$ penalty may be preferable. The first GPCA loading assigns roughly similar coefficients to the first and second hidden factors, and slightly lower coefficients to the third. Once again, PCA, SPCA, SPPCA and RPCA give similar loadings, concentrating mostly on the third hidden factor with a smaller contribution of the first. NMF also gives highest weighting to the third hidden factor, with a smaller contribution from the first and an even smaller by the second. LDA assigns mostly small coefficients, except to one of the components containing the third hidden factor. MNIR performs quite similarly to GPCA, giving very similar coefficients to both the first and second hidden factors, and smaller coefficients to the third hidden factor.

| $L^{1}$ | SCAD | Both | GPCA | PCA | SPCA | RPCA | NMF | LDA | MNIR | SPPCA |
| ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: |
| -0.227 | -0.196 | -0.184 | -0.193 | -0.123 | -0.116 | -0.126 | 0.162 | 0.198 | 0.000 | 0.240 |
| 0.042 | 0.023 | 0.023 | -0.225 | -0.110 | -0.109 | -0.113 | 0.165 | 0.263 | 0.000 | 0.233 |
| -0.119 | -0.101 | -0.089 | -0.226 | -0.127 | -0.125 | -0.131 | 0.164 | 0.098 | 0.000 | 0.243 |
| -0.025 | -0.030 | -0.025 | -0.239 | -0.120 | -0.119 | -0.129 | 0.163 | 0.183 | 0.000 | 0.237 |
| -0.078 | -0.069 | -0.062 | -0.431 | -0.130 | -0.128 | -0.134 | 0.196 | 0.245 | 0.447 | 0.243 |
| -0.036 | -0.058 | -0.068 | -0.437 | -0.130 | -0.131 | -0.125 | 0.197 | 0.208 | 0.469 | 0.243 |
| -0.066 | -0.042 | -0.027 | -0.441 | -0.129 | -0.129 | -0.127 | 0.197 | 0.174 | 0.448 | 0.243 |
| -0.062 | -0.053 | -0.045 | -0.428 | -0.126 | -0.123 | -0.127 | 0.197 | 0.097 | 0.446 | 0.240 |
| -0.674 | -0.682 | -0.684 | -0.156 | -0.656 | -0.657 | -0.650 | 0.605 | 0.307 | -0.301 | 0.515 |
| -0.679 | -0.687 | -0.691 | -0.159 | -0.668 | -0.670 | -0.670 | 0.610 | 0.782 | -0.300 | 0.524 |

Table 4.3: Loadings/directions for classed synthetic data, using $L^{1}$ penalised SGPCA, SCAD penalised SGPCA, the combined penalty SGPCA, GPCA, PCA, SPCA, RPCA, NMF, LDA and MNIR

### 4.4.3 Robustness Against Noise

In order to investigate how robust the loadings given by SGPCA are to noise, we will perform a similar synthetic data analysis to that in Section 4.4.1, but varying the parameter of noise. That is,

$$
\nu_{1} \sim \operatorname{Po}(25) \quad \nu_{2} \sim \operatorname{Po}(30) \quad \nu_{3}=1 \nu_{1}+3 \nu_{2}
$$

We then construct each observation $\mathbf{x}_{i}$ by drawing $v_{1 i}$ from $\mathcal{V}_{1}, v_{2 i}$ from $\mathcal{V}_{2}$ and constructing $v_{3 i}=3 v_{1 i}+v_{2 i}$, then we have

$$
x_{i j}=v_{1 i}+\epsilon_{i j}, j \in\{1,2,3,4\} \quad x_{i j}=v_{2 i}+\epsilon_{i j}, j \in\{5,6,7,8\} \quad x_{i j}=v_{3 i}+\epsilon_{i j}, j \in\{9,10\}
$$

where the $\epsilon_{i j}$ are independent observations from $\mathcal{E}(\eta)$. Then $\eta$ is our noise parameter, we will find SGPCA loadings for each $\eta \in\{1,2,3,4\}$. We again drew 100 observations and performed $L^{1}$ SGPCA with $\lambda_{L}=10^{7}$, SCAD SGPCA with $\lambda=0.1$ and $\lambda_{\mathrm{S}}=10^{6}$, and the combined penalty SGPCA with $\lambda_{L}=10^{6}, \lambda_{\mathrm{S}}=10^{6}$ and $\lambda=0.05$. The loadings for the $L^{1}$ penalised SGPCA are displayed in Table 4.4a, for the SCAD penalised SGPCA in Table 4.4b, and for the combined penalty SGPCA in Table 4.4c.

An examination of Table 4.4a suggests that, for all four magnitudes of noise, the loadings of $L^{1}$ penalised SGPCA are very similar (up to sign changes), with the first loading consistently identifying a direction primarily capturing the second and third hidden factors, and the second loading capturing the first hidden factor. Table 4.4c suggests that the combined penalty SGPCA has the same characteristics. However, Table 4.4b suggests that the SCAD penalised SGPCA is not quite so robust; both loadings seem to

|  | $\eta=1$ | $\eta=2$ | $\eta=3$ | $\eta=4$ |  | $\eta=1$ | $\eta=2$ | $\eta=3$ | $\eta=4$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | 0.1267 | 0.0846 | 0.0355 | -0.0182 |  | -0.4553 | 0.4699 | -0.5496 | -0.4738 |
|  | 0.1385 | 0.0941 | 0.0265 | -0.0007 |  | -0.4751 | 0.4809 | -0.4895 | -0.5268 |
|  | 0.1303 | 0.0548 | 0.0250 | 0.0473 |  | -0.4807 | 0.5112 | -0.4823 | -0.4865 |
|  | 0.1582 | 0.0454 | 0.0602 | -0.0237 |  | -0.4941 | 0.5052 | -0.4524 | -0.4912 |
|  | 0.4128 | 0.4176 | 0.4151 | -0.4448 |  | 0.1482 | -0.0824 | 0.0629 | 0.0641 |
|  | 0.4150 | 0.4401 | 0.4358 | -0.3682 |  | 0.1402 | -0.1060 | 0.0512 | 0.0434 |
|  | 0.4012 | 0.4203 | 0.5003 | -0.4802 |  | 0.1534 | -0.0816 | 0.0637 | -0.0300 |
|  | 0.4293 | 0.4452 | 0.3950 | -0.4873 |  | 0.1601 | -0.0771 | 0.0684 | 0.0329 |
|  | 0.3385 | 0.3417 | 0.3362 | -0.3153 |  | 0.0200 | 0.0181 | -0.0454 | -0.0777 |
|  | 0.3468 | 0.3453 | 0.3352 | -0.3098 |  | 0.0227 | 0.0221 | -0.0606 | -0.0773 |

(a) Two loadings of $L^{1}$ penalised SGPCA under varying levels of noise on synthetic data

|  | $\eta=1$ | $\eta=$ | $\eta=3$ | $\eta=$ |  | $\eta=1$ | $\eta=2$ | $\eta=3$ | $\eta=4$ |
| :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: | :---: |
|  | 0.1266 | 0.0744 | 0.0361 | -0.0347 | $\begin{aligned} & \text { N } \\ & \text { on } \\ & \text { 苐 } \\ & 0 \\ & \hline \end{aligned}$ | -0.4556 | 0.4524 | -0.5505 | -0.0414 |
|  | 0.1388 | 0.0953 | 0.0262 | -0.0237 |  | -0.4753 | 0.4750 | -0.4890 | -0.0324 |
|  | 0.1304 | 0.0580 | 0.0249 | 0.0019 |  | -0.4806 | 0.5168 | -0.4826 | 0.0375 |
|  | 0.1579 | 0.0517 | 0.0599 | -0.0195 |  | -0.4939 | 0.5216 | -0.4516 | -0.0254 |
|  | 0.4126 | 0.4255 | 0.4152 | -0.0392 |  | 0.1481 | -0.0843 | 0.0627 | -0.1338 |
|  | 0.4154 | 0.4339 | 0.4355 | -0.0510 |  | 0.1400 | -0.1084 | 0.0514 | -0.1471 |
|  | 0.4015 | 0.4177 | 0.5005 | -0.0430 |  | 0.1534 | -0.0700 | 0.0637 | -0.1214 |
|  | 0.4287 | 0.4470 | 0.3949 | -0.0336 |  | 0.1601 | -0.0792 | 0.0684 | -0.1211 |
|  | 0.3386 | 0.3414 | 0.3362 | 0.8715 |  | 0.0202 | 0.0146 | -0.0454 | -0.4866 |
|  | 0.3468 | 0.3452 | 0.3352 | -0.4808 |  | 0.0229 | 0.0195 | -0.0606 | -0.8303 |

(b) Two loadings of SCAD penalised SGPCA under varying levels of noise on synthetic data

| $\eta=1$ | $\eta=2$ | $\eta=3$ | $\eta=4$ |  | $\eta=1$ | $\eta=2$ | $\eta=3$ |
| ---: | ---: | ---: | ---: | ---: | ---: | ---: | ---: |$\quad \eta=4$

(c) Two loadings of the combined penalty SGPCA under varying levels of noise on synthetic data

Table 4.4: Investigations of the performance of all three SGPCA variants across varying levels of noise.
vary more as $\eta$ increases, with the loadings for $\eta=4$ being very different from those for $\eta=1$, no longer giving near-equal weights amongst components corresponding to the same hidden factors. This is, perhaps, not surprising, given the use of the $L^{1}$ penalty in RPCA and Joint Sparse Principal Component Analysis (JSPCA) for its robustness to outliers (which become more probable with increasing magnitude of noise).

### 4.5 Dependence on Tolerance

As one might expect, the time taken to estimate SGPCA is heavily dependent on the numerical tolerance; decreasing the tolerance by an order of magnitude has, in our experience, roughly the effect of increasing the time taken by an order of magnitude. Thus it is important to assess how critical a small tolerance is to accurate estimation. In order to determine this, we ran a further simulation study using the same data generation setup as Section 4.4.1. We fitted the three SGPCA variants with the same parameters, varying only the tolerance. We chose a range of tolerance values between $10^{-4}$ and $10^{-9}$, taking the latter as sufficiently converged to use as a reference standard.

In Figure 4.3a we graph the Euclidean distances between the loadings from the reference fit and each of the lower tolerance fits. In Figure 4.3b we graph the deviances of each fit, normalised by the deviance of the reference fit. Both of these figures demonstrate the same behaviour: a period of nearly identical fits followed by rapid convergence to our reference fit. However, it is important to notice the scales on the two figures: the largest deviation amongst all fits is less than one half of one percent from the reference fit, and the greatest Euclidean distance is of order $10^{-2}$. Consequently, we suggest that a practical strategy for choosing an appropriate tolerance value would be to choose the smallest value which allows for feasible optimisation times. If this study is representative of the general situation, a good starting point would be $10^{-7}$.

### 4.6 Healthcare Data

As a "real-world" example, we apply GPCA, SGPCA and classical PCA to the healthcare data, described in Chapter 3. We chose to estimate three directions with each algorithm; experimenting with fewer directions failed to adequately separate the data with any of the methods. The results of this are shown in Figure 5.4.

Certainly it seems that both GPCA and SGPCA provide better separation of the data by pairs of components. The performances of GPCA and SGPCA, while producing rather different pictures in appearance, are comparable in quality. We should be reassured that,
(a) Euclidean distances between loadings from fits with different tolerance values for each of the SGPCA variants. The first loading is shown dashed, and the second is shown solid.

(b) Deviances from each fit across a range of tolerance values, divided by the deviance from the fit with the smallest tolerance. All three SGPCA variants are shown.


Figure 4.3: Behaviour of all SGPCA variants as tolerance is varied.
despite inducing sparse loadings, SGPCA's performance is on-par with the current state of the art on a real-world task.

### 4.7 Discussion

In this chapter, we have extended GPCA by introducing a sparsifying penalty, and shown that we achieve as-good or superior performance in extracting appropriate principal component loadings which prioritise the selection of informative variables in each principal component, assigning much smaller coefficients to non- or less-informative variables.

Comparison of the efficacy of the three penalties we used ( $L^{1}$, SCAD and the combined penalty) suggest that, despite the more complex nature of the SCAD penalty, the $L^{1}$ penalty has the best performance. As such, our preference is for the $L^{1}$ penalty which is significantly simpler than the SCAD penalty, and whose evaluation computationally


Figure 4.4: Plots of pairs of the first three principal components for the healthcare dataset obtained from SGPCA, GPCA and PCA. Red " + " symbols denote the "discharge" class; black dots represent the "follow-up" class.
is cheaper, not requiring any conditional statements and requiring fewer multiplications and no additions. It has the further advantage of only requiring one parameter $\left(\lambda_{L}\right)$, unlike the SCAD penalty which requires $\lambda_{\mathrm{S}}, a$ and $\lambda$. As finding optimal choices of these parameters is most easily accomplished by repeated application of the SGPCA algorithm until one finds a suitable compromise between interpretability and predictive/discriminative power, having only one parameter to find is a significant advantage. We note that this is the method by which we chose the values of these parameters in the applications shown.

We suggest several important topics to extend this work. Firstly, a data-driven method for choosing these tuning parameters which does not require manual examination of the generated coefficients, thus facilitating an automated choice. Secondly, the construction of an appropriate order determination test to choice the required value of $k$. Such a test would likely be applicable also to GPCA. Thirdly, we suggest that this method could likely be extended to a supervised method for dimension reduction, similar to partial least squares (PLS). Finally, we suggest a more challenging problem: to develop a kernel-based extension of GPCA and SGPCA for the extraction of non-linear features.

## — Chapter 5 -

## Sparse Simple Exponential Family Principal Component Analysis

In this chapter, we will introduce an extension to Simple Exponential Principal Component Analysis (SePCA), the work of Li and Tao (2013). As with Generalised Principal Component Analysis (GPCA) in Chapter 4, we will propose a method which incorporates sparsity in the (analogy to) the principal component loadings in order to improve interpretability and performance. Unlike GPCA, SePCA starts with a fully specified Bayesian model, though the two methods are ultimately similar in concept, applying to any exponential family distribution. Here, we will focus on the Poisson case for its application to text data.

### 5.1 Simple Exponential Principal Component Analysis

In Li and Tao (2013), the authors specify a Bayesian model for some data $\mathbf{X}$ following an exponential family (EF) distribution. A plate diagram for the model is given in Figure 5.1. Specifically, they prescribe the model:

- $\boldsymbol{x}_{i} \mid \mathcal{W}, \boldsymbol{y}_{i} \sim \operatorname{ExpFam}\left(\mathcal{W} \boldsymbol{y}_{i}\right)$ for $i=1, \ldots, n$
- $\boldsymbol{y}_{i} \sim \mathrm{~N}(\mathbf{0}, \mathbf{I})$ for $i=1, \ldots, n$
- $\mathcal{W}_{j} \mid \alpha_{j} \sim \mathrm{~N}\left(\mathbf{0}, \alpha_{j}^{-1} \mathbf{I}\right)$ for $j=1, \ldots, k$
where each $\mathcal{X}_{i} \in \mathbb{R}^{p}, \boldsymbol{y}_{i} \in \mathbb{R}^{k}, \mathcal{W} \in \mathbb{R}^{p \times k}$. We denote by $\mathcal{W}_{j}$ the $j^{\text {th }}$ row of $\mathcal{W}$. Finally, $\alpha:=\left(\alpha_{1}, \ldots, \alpha_{k}\right)$ is a vector of precision hyperparameters.

Here we identify the observed value $\mathbf{W}$ of $\mathcal{W}$ with the classical loadings matrix and the observed value $\mathbf{Y}$ formed of the $n$ random vectors $\boldsymbol{y}_{i}(i=1, \ldots, n)$ with the classical scores matrix. The integer $k$ here (as previously) is the number of principal components. It is these three parameters which we are particularly interested in. From a dimension
reduction standpoint, Y is the most important, it is this which we suggest can be used as a lower-dimensional analogue to the observed data.


Figure 5.1: Plate diagram for SePCA

We now draw attention to one of SePCA's greatest advantages: using Automatic Relevance Determination (ARD), introduced in Mackay (1995), SePCA is able to determine the appropriate number of principal components $k$. Note that, according to the model we have given for SePCA, if (for some j) $\alpha_{j} \rightarrow \infty$, then $\mathbf{W}_{j} \xrightarrow{p} \mathbf{0}$. Thus, if we have a sufficiently large $\alpha_{j}$, then we can reasonably determine that the corresponding principal component is not useful to the model. We will refer to the critical value for this as $M$, and remove the column of $\mathbf{W}$ and row of $\mathbf{Y}$ corresponding to any $\alpha_{j}$ which exceeds this threshold. We note that it is precisely because of the precision hyperparameters for $\mathcal{W}$ that we can apply ARD.

### 5.2 Sparse Simple Exponential Principal Component Analysis

In the same spirit as Sparse Generalised Principal Component Analysis (SGPCA) (Chapter 4), we will impose a sparsifying penalty on the loadings in the interest of increasing performance and interpretability. In the case of SePCA, the prior distribution for each $\mathcal{W}_{j} \mid \alpha_{j}$ roughly corresponds to an $L^{2}$ penalty on its entries, so we will not consider such a penalty. Instead, we approach the problem using the method of Frommlet and Nuel (2016) to place an (approximate) $L^{0}$ penalty on $\mathbf{W}$ through an iterative method. We could consider the same penalties as in SGPCA (namely, the SCAD and $L^{1}$ penalties), and we suspect that they would be reasonable choices. Our motivation for choosing the $L^{0}$ penalty here is primarily out of curiosity; it is known to be a strongly sparsifying penalty, and the new development of Frommlet and Nuel's method for approximating it allowed us to use it. Note that

$$
\begin{equation*}
\|\mathbf{W}\|_{0}=\sum_{i=1}^{p} \sum_{j=1}^{k} I\left(\mathbf{W}_{i j} \neq 0\right) \tag{5.1}
\end{equation*}
$$

where $I(\cdot)$ is the indicator function. Frommlet and Nuel use the following approximation (written in the specific case of our problem):

$$
\begin{equation*}
\|\mathbf{W}\|_{0} \approx \sum_{i=1}^{p} \sum_{j=1}^{k} \frac{\mathbf{W}_{i j}^{2}}{\left(\mathbf{W}_{i j}^{(t-1)}\right)^{2}+\delta} \tag{5.2}
\end{equation*}
$$

where $\mathbf{W}^{(t-1)}$ is the value of $\mathbf{W}$ from the previous interation and $\delta>0$ is a very small value. The advantage to this method is the differentiability of (5.2), allowing us to use gradient-based optimisation methods in the estimation procedure. The $L^{0}$ penalty, on the other hand, is not differentiable at 0 .

Remark. In Frommlet and Nuel (2016), the authors' method requires that each $L^{0}$ penalty be approximate by multiple interations of the approximation (5.2). However, in our method we will use only a single interation. This is due to the iterative nature of the estimation algorithm we will develop in Section 5.3.

We will precisely define Sparse Simple Exponential Principal Component Analysis (SSePCA) in the following section by its objective function.

Remark. Note that, like SePCA, SSePCA is also amenable to using ARD to determine the number of principal components needed.

### 5.3 Estimation

In this section, we will describe jointly the process of estimating SePCA and SSePCA. We fit both by maximum a-posteriori (MAP) estimation, with the caveat that SSePCA is technically penalised MAP. Both methods require the specification of the exponential family in order to write down explicitly the objective functions and gradients, so we give them for the specific case of the Poisson distribution and will note which parts must be adapted for other EF distributions.

We begin by deriving the objective function for SePCA, from which we can calculate the objective function for SSePCA. We begin with the conditional probability $\mathbb{P}(\mathbf{X} \mid \mathbf{W}, \mathbf{Y})$.

As this is in the exponential family, we know the general form:

$$
\begin{aligned}
\mathbb{P}(\mathbf{X} \mid \mathbf{W}, \mathbf{Y}) & \propto \exp \sum_{i=1}^{n}\left[\mathbf{x}_{i}^{\mathrm{T}} \mathbf{W} \mathbf{y}_{i}+g\left(\mathbf{W} \mathbf{y}_{i}\right)\right] \\
& \propto \exp \sum_{i=1}^{n}\left[\mathbf{x}_{i}^{\mathrm{T}} \mathbf{W} \mathbf{y}_{i}-\sum_{j=1}^{k} \exp \left(\left[\mathbf{W y}_{i}\right]_{j}\right)\right] \\
& \propto \exp \left[\operatorname{Trace}\left(\mathbf{X}^{\mathrm{T}} \mathbf{W Y}\right)-\sum_{i=1}^{n} \sum_{j=1}^{p} \exp \left([\mathbf{W Y}]_{i j}\right)\right]
\end{aligned}
$$

where $\propto$ indicates proportionality up to a constant multiple (with respect to X , which is fixed during an optimisation over $\mathbf{W}$ and $\mathbf{Y}$ ).

Similarly, the joint $\log$-posterior of $\mathbf{W}, \mathbf{Y} \mid \mathbf{X}, \boldsymbol{\alpha}$ is, up to addition of a constant:

$$
\log \mathbb{P}(\mathbf{X}, \mathbf{Y}, \mathbf{W}, \boldsymbol{\alpha})=\log \mathbb{P}(\mathbf{X} \mid \mathbf{W}, \mathbf{Y})+\log \mathbb{P}(\mathbf{Y})+\log \mathbb{P}(\mathbf{W} \mid \boldsymbol{\alpha})+\text { const }
$$

Then the objective function for SePCA is

$$
\begin{equation*}
P=\operatorname{Trace}\left(\mathbf{X}^{\mathrm{T}} \mathbf{W} \mathbf{Y}\right)-\sum_{i=1}^{n} \sum_{j=1}^{p} \exp (\mathbf{W} \mathbf{Y})-\frac{1}{2} \operatorname{Trace}\left(\mathbf{Y}^{\mathrm{T}} \mathbf{Y}\right)-\frac{1}{2} \operatorname{Trace}\left(\mathbf{W}^{\mathrm{T}} \mathbf{W} \operatorname{Diag}(\boldsymbol{\alpha})\right) \tag{5.3}
\end{equation*}
$$

where we can ignore the constant.
It is now very simple to define the objective function for SSePCA at the $t^{\text {th }}$ iteration as

$$
\begin{equation*}
P^{\mathrm{s}}=P-\lambda \sum_{i=1}^{p} \sum_{j=1}^{k} \frac{\mathbf{W}_{i j}^{2}}{\left(\mathbf{W}_{i j}^{(t-1)}\right)^{2}+\delta} \tag{5.4}
\end{equation*}
$$

where the second term comes from (5.2).
We see now that SSePCA is the weighted sum of the SePCA objective function and the approximate $L^{0}$ penalty. The value of $\lambda$ controls to what extent we prioritise sparsity over explanatory power for the observed data; determining the appropriate value is a heavily data- and context-dependent task. Our practical advice is to try a logarithmic sequence of $\lambda$ values and select the value which best balances sparsity against predictive/discriminative power in your application.

In order to optimise these two objective functions, we make use of the differentiability with respect to both $\mathbf{W}$ and $\mathbf{Y}$. The required derivatives are

$$
\begin{gather*}
\frac{\partial P}{\partial \mathbf{W}}=\mathbf{X} \mathbf{Y}^{\mathrm{T}}-\exp (\mathbf{W} \mathbf{Y}) \mathbf{Y}^{\mathrm{T}}-\mathbf{W} \operatorname{Diag}(\boldsymbol{\alpha})  \tag{5.5}\\
\frac{\partial P^{\mathrm{s}}}{\partial \mathbf{W}}=\frac{\partial P}{\partial \mathbf{W}}-2 \lambda \frac{\mathbf{W}}{\left(\mathbf{W}^{(t-1)}\right)^{2}+\delta}  \tag{5.6}\\
\frac{\partial P}{\partial \mathbf{W}}=\frac{\partial P^{\mathrm{s}}}{\partial \mathbf{W}}=\mathbf{W}^{\mathrm{T}} \mathbf{X}-\mathbf{W}^{\mathrm{T}} \exp (\mathbf{W Y})-\mathbf{Y} \tag{5.7}
\end{gather*}
$$

We can now give an algorithm for the estimation of SePCA and SSePCA in Algorithm 2. This algorithm is novel, being unlike the algorithm originally given for SePCA in Li and Tao 2013. The original algorithm used an expectation-maximisation scheme, where we directly optimise the posterior likelihood using gradient information for efficiency. We note that, like the algorithm for SGPCA, this algorithm does not have a guarantee of finding the global optimum. One strategy may be to initialise the algorithm at randomly distributed points and choose the best performing solution. In practice, we have not found this to be necessary; the (most likely local) optimum found from the proposed initial values has been sufficiently performant.

Remark. It is worth noting that we have generally found it useful to temporarily increase the value of $M$ for the first ten iterations of the optimisation as this time is when we are most likely to spuriously remove components due to a poor starting point.

### 5.4 Synthetic Data Studies

All investigations in this section will use the same basic model, with some small adaptations. We will use the two hidden factors

$$
\begin{equation*}
V_{1} \sim \text { Poisson }(20) \quad V_{2} \sim \text { Poisson }(30) \tag{5.8}
\end{equation*}
$$

We will also use an error distribution $E$, constructed by drawing an observation from a Poisson(2) distribution and multiplying by 1 or -1 with equal probability.

The first analysis will use two datasets, with "true" dimensions 1 and 2 respectively, which we will refer to as X1D and X2D. Each consists of 100 observations of a random vector of length 10, but the construction of that vector differs. For the component selection procedure we set $M=100$ except the first 10 iterations where $M=500$ (as was mentioned in Section 5.3 we do this to avoid removing components too early). Also for the Sparse Simple Poisson Principal Component Analysis (SSPPCA) algorithm $\delta=10^{-8}$.

To construct X1D, let $v_{1 i}, i=1, \ldots, 100$ be independently observed values of $V_{1}$ and let $\varepsilon_{i j}, i=1, \ldots, 100, j=1, \ldots, 10$ be independently observed values of $E$. Then
the $i^{\text {th }}$ observation in X1D has its first two components equal to $v_{1 i}$ plus error, and the remaining eight components are equal to $2 * v_{1 i}$ plus error. Formally each observation has the form

$$
\left[\left(v_{1 i}+\varepsilon_{i 1}\right),\left(v_{1 i}+\varepsilon_{i 2}\right),\left(2 v_{1 i}+\varepsilon_{i 3}\right), \ldots,\left(2 v_{1 i}+\varepsilon_{i 10}\right)\right]^{\mathrm{T}}
$$

To give a bit more insight here, one should expect that a good dimension reduction in this case will identify that we need exactly one component, which has larger coefficients to variables 3 to $p$ and it has smaller coefficients for variables 1 and 2 .

Similarly, the $i^{\text {th }}$ observation in X2D has its first two components equal to an observed value $v_{1 i}$ of $V_{1}$ plus independent errors, its second two components equal to an observed value $v_{2 i}$ of $V_{2}$ plus independent errors, and its final six components equation to $v_{1 i}+3 v_{2 i}$ plus independent errors:

$$
\left[\left(v_{1 i}+\varepsilon_{i 1}\right),\left(v_{1 i}+\varepsilon_{i 2}\right),\left(v_{2 i}+\varepsilon_{i 3}\right),\left(v_{2 i}+\varepsilon_{i 4}\right),\left(v_{1 i}+3 v_{2 i}+\varepsilon_{i 5}\right) \ldots,\left(v_{1 i}+3 v_{2 i}+\varepsilon_{i 10}\right)\right]^{\mathrm{T}}
$$

To both of these datasets we applied each of Sparse Probabilistic Principal Component Analysis (SPPCA), SSPPCA, PCA, Sparse Principal Component Analysis (SPCA), GPCA and SGPCA. For the latter three we needed to specify the dimension; for SPPCA and SSPPCA the automatic relevance determination criterion successfully identified the true dimension. The loadings for the one-dimensional data are given in Table 5.1; SPPCA, SSPPCA, PCA and SPCA all give very similar results qualitatively, giving equal weighting to components three through ten (corresponding to the $2 v_{1}$ term) and slightly smaller values to the first two components corresponding to the $v_{1}$ term. GPCA gives approximately equal weighting to all the terms. SGPCA, on the other hand, gives considerably more sporadic loadings. This is perhaps due to the lack of sparsity of the underlying data.

In Table 5.2 we give the two loadings for the two-dimension data. Here, the first SPPCA loading gives roughly equal weight to the first two and last six components, corresponding to the $v_{1}$ and $v_{1}+3 v_{2}$ terms respectively, and a slightly lower loading to the second two components (corresponding to the $v_{2}$ terms). The second SPPCA loading gives most weight to the last six components, with small weights for the second pair of components and the lowest weights to the first pair of components. The performance of SSPPCA is more easily interpretable; the first loading gives highest weighting to the last six components, with smaller weight for the first four; the second loading strongly identifies the first two components with near-zero weighting given to all other terms. PCA's first loading primarily identifies the $v_{1}+3 v_{2}$ term, with its second primarily identifying the $v_{1}$ term; SPCA does similarly with sparser loadings. GPCA's first loading gives
approximately equal weighting to all terms (except for the very first component), with its second primarily emphasising the $v_{1}$ components. Finally, SGPCA's first loading identifies a combination of the $v_{1}$ and $v_{1}+3 v_{2}$ terms, while its second fairly strong identifies the $v_{1}$ components. Of all the loadings, the most successful at identifying the hidden factors are the second loadings of SSPPCA, PCA, SPCA, GPCA and SGPCA, with SSPPCA, SPCA and SGPCA slightly better as the other components are driven closer to 0 .

### 5.4.1 Order Determination

In order to investigate the accuracy of the order determination provided by ARD, we conducted similar experiments to those in Section 5.4, varying several parameters. We looked at $p \in\{10,20\}, p \in\{25,50,100,200\}, k \in\{1,2,3\}$. For each combination of parameters, we constructed data by the following method and used both SPPCA and SSPPCA to estimate $k$, repeating this 50 times in order to understand the average behaviour. When $k=1$, the $i^{\text {th }}$ observation $(i=1, \ldots, n)$ was given by

$$
\left[\left(v_{1 i}+\varepsilon_{i 1}\right),\left(v_{1 i}+\varepsilon_{i 2}\right),\left(2 v_{1 i}+\varepsilon_{i 3}\right), \ldots,\left(2 v_{1 i}+\varepsilon_{i p}\right)\right]^{\mathrm{T}}
$$

When $k=2$, it was given by

$$
\left[\left(v_{1 i}+\varepsilon_{i 1}\right),\left(v_{1 i}+\varepsilon_{i 2}\right),\left(v_{2 i}+\varepsilon_{i 3}\right),\left(v_{2 i}+\varepsilon_{i 4}\right),\left(v_{1 i}+3 v_{2 i}+\varepsilon_{i 5}\right) \ldots,\left(v_{1 i}+3 v_{2 i}+\varepsilon_{i p}\right)\right]^{\mathrm{T}}
$$

Finally, when $k=3$ it was given by

$$
\begin{aligned}
& {\left[\left(v_{1 i}+\varepsilon_{i 1}\right),\left(v_{1 i}+\varepsilon_{i 2}\right),\left(v_{2 i}+\varepsilon_{i 3}\right),\left(v_{2 i}+\varepsilon_{i 4}\right),\left(v_{3 i}+\varepsilon_{i 5}\right),\left(v_{3 i}+\varepsilon_{i 6}\right),\right.} \\
& \\
& \left.\left(3 v_{1 i}+2 v_{2 i}+2 v_{3 i}+\varepsilon_{i 7}\right) \ldots,\left(3 v_{1 i}+2 v_{2 i}+2 v_{3 i}+\varepsilon_{i k}\right)\right]^{\mathrm{T}}
\end{aligned}
$$

For each of these, $v_{1 i}, v_{2 i}$ and $v_{3 i}$ denote observations from $V_{1}, V_{2}$ and $V_{3}$ defined in (5.8) respectively.

Tables 5.3 a and 5.3 b give the percentage of times each algorithm correctly identified $k$ for a given choice of $n$ and $k$ with $p=20$. Generally, it appears that SPPCA performs better for small $n$, but its performance degrades as $n$ increases. However, our proposed SSPPCA's performance improves as $n$ increases and in fact performs significantly better by $n=200$.

### 5.4.2 Synthetic Data with Classes

Although SPPCA and SSPPCA are not supervised methods, it is instructive to see whether, given data arising from two or more classes, they are able to find principal components which are able to distinguish between these classes. This gives some indication of their

Data: X, $\lambda, M$, tol
Result: W, Y, $k$
initialise $k=p-1$;
initialise $\mathbf{W}^{(0)}$ and $\mathbf{Y}^{(0)}$ as the loadings and score vectors of standard Principal Component Analysis (PCA);
initialise $\boldsymbol{\alpha}^{(0)}=(1, \ldots, 1)$;
set $t=1$;
while $P$ changed more than tol or components removed on previous interation do set $\mathbf{W}^{(t)}, \mathbf{Y}^{(t)}=\arg \max _{\mathbf{W}, \mathbf{Y}} P$; for $j=1$ to $k$ do
set $\alpha_{j}=p /\left\|\mathbf{W}_{j}^{(t)}\right\|_{2}^{2} ;$
end
if $\alpha_{j}>M$ for any $j$ then /* this block is the application of ARD */
remove $j^{\text {th }}$ component from $\alpha$;
remove $j^{\text {th }}$ column from $\mathbf{W}$;
remove $j^{\text {th }}$ row from $\mathbf{Y}$;
reorder components of $\alpha$ from smallest to largest;
reorder columns of $\mathbf{W}$ and rows of $\mathbf{Y}$ in the same order;
decrement $k$;
end
increment $t$;
end
Algorithm 2: SePCA and SSePCA Estimation Procedure.

| SPPCA | SSPPCA | PCA | SPCA | GPCA | SGPCA |
| :---: | :---: | :---: | :---: | :---: | :---: |
| -0.26 | -0.26 | 0.17 | 0.00 | 0.34 | 0.06 |
| -0.27 | -0.27 | 0.17 | 0.00 | 0.33 | 0.42 |
| -0.33 | -0.33 | 0.35 | 0.47 | 0.32 | -0.00 |
| -0.33 | -0.33 | 0.33 | 0.27 | 0.29 | -0.50 |
| -0.33 | -0.33 | 0.34 | 0.35 | 0.30 | -0.43 |
| -0.33 | -0.33 | 0.36 | 0.31 | 0.33 | -0.06 |
| -0.33 | -0.33 | 0.34 | 0.37 | 0.31 | -0.44 |
| -0.33 | -0.33 | 0.33 | 0.32 | 0.31 | 0.19 |
| -0.33 | -0.33 | 0.36 | 0.39 | 0.33 | -0.01 |
| -0.33 | -0.33 | 0.35 | 0.32 | 0.31 | -0.40 |

Table 5.1: Loadings for X1D

| SPPCA | SSPPCA | PCA | SPCA | GPCA | SGPCA |
| :---: | :---: | :---: | :---: | :---: | :---: |
| -0.33 | -0.23 | 0.01 | 0.00 | -0.18 | 0.37 |
| -0.31 | -0.23 | 0.03 | 0.00 | -0.31 | 0.04 |
| -0.23 | -0.26 | 0.12 | 0.00 | -0.35 | 0.10 |
| -0.23 | -0.26 | 0.12 | 0.00 | -0.35 | 0.06 |
| -0.34 | -0.36 | 0.41 | 0.40 | -0.33 | 0.40 |
| -0.34 | -0.36 | 0.40 | 0.36 | -0.32 | 0.40 |
| -0.34 | -0.36 | 0.40 | 0.42 | -0.32 | 0.38 |
| -0.34 | -0.36 | 0.41 | 0.40 | -0.33 | 0.37 |
| -0.34 | -0.36 | 0.40 | 0.45 | -0.33 | 0.31 |
| -0.34 | -0.36 | 0.39 | 0.42 | -0.31 | 0.38 |

(a) First loading

| SPPCA | SSPPCA | PCA | SPCA | GPCA | SGPCA |
| :---: | :---: | :---: | :---: | :---: | :---: |
| 0.12 | -0.76 | 0.73 | -0.82 | 0.71 | 0.67 |
| 0.14 | -0.64 | 0.64 | -0.57 | 0.61 | 0.68 |
| 0.29 | 0.09 | -0.18 | 0.00 | -0.20 | -0.10 |
| 0.28 | 0.08 | -0.16 | 0.00 | -0.19 | -0.05 |
| 0.37 | 0.00 | 0.01 | 0.00 | -0.09 | -0.11 |
| 0.36 | 0.00 | 0.07 | 0.00 | -0.08 | -0.09 |
| 0.37 | -0.00 | -0.00 | 0.00 | -0.09 | -0.12 |
| 0.37 | 0.00 | -0.02 | 0.00 | -0.10 | -0.12 |
| 0.37 | 0.00 | -0.00 | 0.00 | -0.09 | -0.08 |
| 0.37 | 0.00 | -0.02 | 0.00 | -0.09 | -0.15 |

(b) Second Loading

Table 5.2: Two loadings from X2D

|  | $d$ |  |  |
| :---: | :---: | :---: | :---: |
| $N$ | 1 | 2 | 3 |
| 25 | 94 | 24 | 18 |
| 50 | 82 | 62 | 26 |
| 100 | 62 | 24 | 26 |
| 200 | 24 | 16 | 14 |

(a) SPPCA

|  | $d$ |  |  |
| :---: | :---: | :---: | :---: |
| $N$ | 1 | 2 | 3 |
| 25 | 2 | 8 | 4 |
| 50 | 42 | 10 | 8 |
| 100 | 82 | 60 | 18 |
| 200 | 78 | 70 | 50 |

(b) SSPPCA

Table 5.3: Percentage of correct identification of $d$ for SPPCA and SSPPCA
suitability for use as a step before applying a clustering or classification algorithm (depending on whether labels are available or not). To this end, we construct two sets of classed data; the first having observations from two classes with equal sample sizes from both, the second having three classes with imbalanced sample sizes.

We will use again the hidden factors from (5.8) and both datasets have dimension $p=10$ and total sample size $n=100$. We will denote the two-class data by X2C and the three-class data by X3C. The first class for both datasets will have its first two components equal to observations $v_{2}$ of $V_{2}$ with independent error $E$ and the remaining eight components equal to $3 v_{2}$ with independent error. The second class for both will have first two components equal to $2 v_{3}$ with independent error and the remaining eight components equal to $v_{3}$, where the $v_{3}$ are observations of $V_{3}$. The third class will have all components equal to observations from $V_{1}$ with independent error. The two-class data X2C has 50 observations from the first class and 50 from the second. The three-class data X3C is divided between 25 observations of the first class, 25 observations of the second class, and 50 observations of the third class.

The loadings from applying SPPCA, SSPPCA, GPCA, SGPCA, PCA and SPCA to X2C are given in Figure 5.2. For GPCA, SGPCA, PCA and SPCA we must specify a dimension: as both SPPCA and SSPPCA choose $k=2$ we use that value. All six algorithms achieve good separation of the two classes. Visually, it appears that SPPCA and SSPPCA (in Figures 5.2 a and 5.2 b respectively) give the best clustering of the two classes. We use the method of silhouettes put forward by Rousseeuw (1987) to confirm this, using the Euclidean distance metric and clusters found using $k$-medioid clustering. The silhouette of the $i^{\text {th }}$ observation is given by

$$
\frac{b(i)-a(i)}{\max \{a(i), b(i)\}}
$$

where $a(i)$ is the average dissimilarity of the $i^{\text {th }}$ observation to the other members of its cluster and $b(i)$ is the lowest average dissimilarity of the $i^{\text {th }}$ observation to any other cluster. We can thus interpret the silhouette as a measure of how well a data point is assigned to its cluster; the average silhouette over a dataset gives a measure for how well clustered the data is. Average silhouette values range between -1 and 1 ; the closer to 1 the better the clustering. In Table 5.4 we give average silhouettes for X2C for each of the six algorithms. Our visual intuition that SPPCA and SSPPCA give the best clustering is confirmed. For X3C we can see that SSPPCA is actually the only algorithm that is able to achieve separation of the 3 classes using only the first principal component. All other methods need the second direction to achieve this separation. When we compare the average silhouette measure in Table 5.4 we see that SSPPCA is actually behind the other
methods which have similar values but this may as well be due to the fact that we do not apply the penalty for multiple iterations as was explained in Section 5.3.


Figure 5.2: Two directions from each algorithm for X2C, with one class shown with red squares, the other with black triangles.

|  | SPPCA | SSPPCA | PCA | SPCA | GPCA | SGPCA |
| :--- | :---: | :---: | :---: | :---: | :---: | :---: |
| X2C | 0.94 | 0.95 | 0.75 | 0.75 | 0.75 | 0.67 |
| X3C | 0.86 | 0.86 | 0.78 | 0.79 | 0.78 | 0.78 |

Table 5.4: Average (Euclidean) silhouettes

### 5.5 Healthcare Data

We will now examine the efficacy of SPPCA and SSPPCA in reducing the dimension of the healthcare dataset as detailed in Chapter 3.

In Figure 5.4 we show the results of applying SPPCA, SSPPCA, GPCA and SGPCA to this dataset. Discharge data points are shown with crosses and follow up points are


Figure 5.3: Two directions from each algorithm for X3C
shown with circles. For both SPPCA and SSPPCA we used $M=100$, for the latter we chose $\lambda=0.06$. From Figure 5.4 a we can see that SPPCA estimated $k$ as 1 ; on the other hand, from Figure 5.4b we see that SSPPCA chose $k=2$. Based on this, we chose $k=2$ for GPCA and SGPCA, which require a fixed value. The difficulty inherent in achieving good class separation as a result of the strong class imbalance is evident. Of the four, GPCA (Figure 5.4 c ) is the worst, with the discharge points not particularly tightly clustered and difficult to separate from the follow up points. SGPCA (Figure 5.4d, on the other hand, tightly clusters the discharge points. SPPCA (Figure 5.4a, also manages to tightly cluster the discharge points. SSPPCA (Figure 5.4b) does not cluster the discharge points particularly closely, but does achieve better visual separation than GPCA.

In order to better quantify the clustering, we give the average (Euclidean) silhouettes in Table 5.5. Based on this performance metric, SPPCA and SGPCA are the best performers. Average silhouette is designed to measure and inform about the performance of clustering algorithms and not the accuracy of feature extraction. The fact that SPPCA seems to be worst may either be due to this or to the fact that we are not doing
multiple iterations of the penalty.


Figure 5.4: The resulting principal components from applying SPPCA, SSPPCA, GPCA and SGPCA to the healthcare data. The "discharge" class is plotted as red squares, the "follow-up" class is shown as black triangles.

### 5.6 Discussion

In this paper we have developed a Poisson based PCA algorithm which we called SPPCA and which was based on the SePCA (Li and Tao 2013). We use a different algorithm

| SPPCA | SSPPCA | GPCA | SGPCA | PCA |
| :---: | :---: | :---: | :---: | :---: |
| 0.38 | 0.34 | 0.23 | 0.17 | 0.21 |

Table 5.5: Average silhouettes first the healthcare data.
for inference on $\mathbf{W}$ and $\mathbf{Y}$ than SePCA. We have illustrated that by example in the specific case where the distribution is Poisson, to give SPPCA. We have also introduced an approximate $L^{0}$ sparsity penalty in this context to allow for Sparse SPPCA. In a more general framework this can be seen as a unified way of achieving sparse or non-sparse feature extraction from a Poisson-based PCA algorithm. At the same time this algorithm should be straightforward to extend to other distributions in the exponential family by modifying appropriately the formulas.

The sparse algorithm performs particularly well, both in latent dimension discovery and in class separation for multi class Poisson data. Computation times are acceptable for small samples ( $n \leq 500$ ), but become a slightly more burdensome for larger samples. It is worth noting that there exist multiple solutions or local maxima. This is also dealt with simply, by evaluating multiple optima using fully specified probability model. In practice, we have found that this has not been necessary, the maxima obtained starting from the Gaussian PCA have performed perfectly well.

There is scope for extension of this work. First of all it is interesting to introduce different more complex sparsity penalties, such as the $L^{1}$ or SCAD penalties and compare their performance. Another possible extension is the development of nonlinear feature extraction methods as well as sparse nonlinear feature extraction method in the generalised PCA setting for non-Gaussian data.

## - Chapter 6 -

## Quasi-Likelihood Principal Component Analysis

### 6.1 Introduction

The two Principal Component Analysis (PCA)-like methods we have concentrated on so far share a number of similarities - both aim to optimise some criterion. In the case of (S)GPCA this is the (penalised) deviance, and in the case of (S)SePCA it is the (penalised) posterior likelihood. As an optimisation problem of fairly simple functions, we can express these methods as finding the zeros of the derivatives. In fact, there is a rich theory for estimating statistical parameters from such a premise, the theory of estimating equations. However, this theory has, to date, focused on solving problems with scalar or vector parameters, whilst dimension reduction methods like we are interested in require higher dimensional analogues. In this chapter, we will develop the theory of estimating equations for tensor parameters. We shall do so with the additional restriction of requiring that our estimating equations are themselves tensors in order to have a pleasing invariance with respect to changes of parametrisations.

The purpose of this theory is to provide a framework for the development and understanding of PCA-like methods. As such, having developed a theory for tensor-parameter estimating equations, we will then apply that theory to a range of dimension reduction methods in the literature in order to better understand the differences between each. To illustrate this technique, we shall apply it to three important methods: Generalised Principal Component Analysis (Landgraf and Lee 2015b), the method of Collins et al. (2002), and Simple Exponential Family Principal Component Analysis (Li and Tao 2013).

It is worth noting that, by expressing these methods as estimating equations, one can very easily write code to find the dimension reduction parameters using a standard root-finding algorithm. In many programming languages, it is even possible to input the
appropriate equations symbolically. While this will rarely be the most computationally efficient method, it illustrates a very compelling feature of the estimating equations viewpoint - newly devised methods can be tried experimentally with very little time investment.

### 6.2 Tensor Estimating Equations

Before defining tensor estimating equations, we will recap the vector case, in order to illustrate the concepts and provide motivation to the tensor extension. We will then introduce tensors and some key properties, before showing the extension of vector estimating equations to the tensor case.

### 6.2.1 Vector Parameter Estimating Equations

Suppose we are given a sample of data $\left\{\mathbf{x}_{1}, \ldots, \mathbf{x}_{n}\right\}$ from a random variable $\mathbf{X}$ which takes values in $\mathscr{X} \subseteq \mathbb{R}^{p}$ with a family of probability distributions $\mathscr{P}:=\left\{\mathbb{P}_{\theta}: \theta \in \Theta\right\}$ and set of possible parameters $\Theta$ being an open subset of $\mathbb{R}^{q}$. As a method for estimating the "true" parameter value $\theta$, we are interested in determining an optimal estimating function $\mathbf{G}: \mathscr{X}^{n} \times \Theta \rightarrow \mathbb{R}^{q}$ whose roots give good estimates of the "true" $\theta$.

Let $\mathscr{G}$ be the class of estimation functions. We require that each G has zero mean, is square-integrable, and for which $\mathbb{E}(\dot{\mathbf{G}}):=\left(\mathbb{E}\left(\frac{\partial \mathbf{G}_{i}}{\partial \theta_{j}}\right)\right)$ and $\mathbb{E}\left(\mathbf{G G}^{\mathrm{T}}\right)$ are non-singular. To compare estimating functions, we will need a notion of standardisation, given in Definition 6.2.1. This will allow us to find an optimal estimating function for parameter estimation.

Definition 6.2.1 (Standardised Estimating Function). Given an estimating function $\mathbf{G} \in$ $\mathscr{G}$, the corresponding standardised estimating function $\mathbb{G}$ is

$$
-\mathbb{E}\left(\dot{\mathbf{G}}^{\mathrm{T}}\right)\left(\mathbb{E}\left(\mathbf{G G}^{\mathrm{T}}\right)\right)^{-1} \mathbf{G}
$$

We now seek a notion of finite-sample optimality, which we will denote $O_{F}$-optimality. This will use an information criterion $\mathscr{E}(\cdot)$, whose definition is given in Definition 6.2.2, which is an analogue to the Fisher Information.

Definition 6.2.2 (Information Criterion). The information criterion $\mathscr{E}(\mathbf{G})$ is given by

$$
\mathscr{E}(\mathbf{G}):=\mathbb{E}\left(\mathbf{G G}_{\sim}^{\mathrm{T}}\right)=(\mathbb{E}(\dot{\mathbf{G}}))^{\mathrm{T}}\left(\mathbb{E}\left(\mathbf{G G}^{\mathrm{T}}\right)\right)^{-1}(\mathbb{E}(\dot{\mathbf{G}}))
$$

We are now able to give the definition of $O_{F}$-optimality. It will be given in terms of a subclass $\mathscr{H}$ of estimating functions, as this will usually be the case in practice.

Definition 6.2.3 ( $O_{F}$-Optimality). Let $\mathscr{H} \subseteq \mathscr{G}$. Then we call $\mathbf{G}_{*} \in \mathscr{H}$ the $O_{F}$-optimal estimating function if $\mathscr{E}\left(\mathbf{G}_{*}\right)-\mathscr{E}(\mathbf{G})$ is positive semi-definite for all $\mathbf{G} \in \mathscr{H}, \theta \in \Theta$ and $\mathbb{P}_{\boldsymbol{\theta}} \in \mathscr{P}$.

This concept is known as Loewner optimality in experimental design, its name originating from the Loewner partial ordering on positive semi-definite matrices. Should the score function $\mathbf{U}$ exist, there is an alternative and equivalent definition:

Definition 6.2.4. The estimating function $\mathbf{G}_{*} \in \mathscr{H}$ is $O_{F}$-optimal in $\mathscr{H}$ if for all $\mathbf{G} \in \mathscr{H}$, $\theta \in \Theta$ and $\mathbb{P}_{\theta} \in \mathscr{P}$

$$
\mathbb{E}\left((\underset{\sim}{\mathbf{U}}-\underset{\sim}{\mathbf{G}})(\underset{\sim}{\mathbf{U}}-\underset{\sim}{\mathbf{G}})^{\mathrm{T}}\right)-\mathbb{E}\left(\left(\underset{\sim}{\mathbf{U}}-{\underset{\sim}{\mathbf{G}}}_{*}\right)\left(\underset{\sim}{\mathbf{U}}-{\underset{\sim}{G}}_{*}\right)^{\mathrm{T}}\right)
$$

is positive semi-definite.
Definition 6.2 .4 can be seen as minimising the dispersion distance of $\mathbf{G}_{*}$ from the score function. For this reason, such optimal estimating functions are also known as "quasi-score" functions, as they behave like score functions in many respects and they are the optimal approximation to this score function (if it exists) within a given class of functions.

Finally, we give one more equivalent criterion for $O_{F}$-optimality which will be useful in practice.

Theorem 6.2.1. The estimating function $\mathbf{G}_{*} \in \mathscr{H} \subseteq \mathscr{G}$ is $O_{F}$-optimal if

$$
\begin{equation*}
\mathbb{E}\left({\underset{\sim}{G}}_{*} \mathbf{G}^{\mathrm{T}}\right)=\mathbb{E}\left({\underset{\sim}{G} \mathbf{G}_{*}}^{\mathrm{T}}\right)=\mathbb{E}\left(\mathbf{G}_{\sim}^{\mathbf{G}}{ }^{\mathrm{T}}\right) \tag{6.1}
\end{equation*}
$$

for all $\mathbf{G} \in \mathscr{H}$. Equivalently, if $(\mathbb{E}(\dot{\mathbf{G}}))^{-1} \mathbb{E}(\mathbf{G}) \mathbb{E}\left(\mathbf{G}_{*}\right)$ is a constant matrix. In the converse direction, if $\mathbf{G}_{*} \in \mathscr{H}$ is an $O_{F}$-optimal estimating function and $\mathscr{H}$ is convex then (6.1) holds.

Theorem 6.2.2. If the score function $\mathbf{u}$ exists and $\mathbf{u} \in \mathscr{H}$, then $\mathbf{u}$ is the optimal estimating function.

### 6.2.2 Tensor Preliminaries

Perhaps the most familiar way to view tensors is as multi-dimensional arrays. Much as vectors are "larger" scalars, and matrices are "larger" vectors, tensors form a natural way to think about such objects in arbitrary dimensions. To be more precise, any tensor can be represented as a multi-dimensional array with respect to a basis, just as any vector can be represented by a one-dimensional array with respect to a basis in the appropriate vector space.

The final distinctive element of tensors is the concept of contravariance and covariance. These describe the way that a given vector transforms under a change of basis. Briefly, these are defined below.

Definition 6.2.5. A vector $\mathbf{v}$ with components $v^{1}, \ldots, v^{n}$ with respect to a basis $\mathbf{e}_{1}, \ldots, \mathbf{e}_{n}$ is contravariant if the new components $\hat{v}^{1}, \ldots, \hat{v}^{n}$ with respect to a new basis $\hat{\mathbf{e}}_{1}, \ldots, \hat{\mathbf{e}}_{n}$ are given by

$$
\hat{v}^{i}=\sum_{j=1}^{n}\left(\mathbf{R}^{-1}\right)_{j}^{i} \nu^{j}
$$

where $\mathbf{R}$ is the mapping such that $\hat{\mathbf{e}}_{i}=\sum_{j=1}^{n} \mathbf{e}_{j} \mathbf{R}_{i}^{j}$.
Definition 6.2.6. A vector $\mathbf{w}$ with components $w_{1}, \ldots, w_{n}$ with respect to a basis $\mathbf{e}_{1}, \ldots, \mathbf{e}_{n}$ is covariant if the new components $\hat{w}_{1}, \ldots, \hat{w}_{n}$ with respect to the new basis $\hat{\mathbf{e}}_{1}, \ldots, \hat{\mathbf{e}}_{n}$ are given by

$$
\hat{w}_{i}=\sum_{j=1}^{n} w_{j} \mathbf{R}_{i}^{j}
$$

with $\mathbf{R}$ again the mapping such that $\hat{\mathbf{e}}_{i}=\sum_{j=1}^{n} \mathbf{e}_{j} \mathbf{R}_{i}^{j}$.
The above definitions also introduce a piece of common notation when working with tensors, where subscript indices are used for covariant components and superscript indices are used for contravariant components. We will make use of this notation going forward. A second piece of notation which is very common is that of "Einstein notation" or the "Einstein summation convention". Many tensor definitions and operations involve summations, in order to reduce the notation needed, it will be taken as understood that a repeated algebraic index such as $i$ or $j$ is to be summed over. For example, we can write the following:

$$
\hat{w}_{i}=\sum_{j=1}^{n} w_{j} \mathbf{R}_{i}^{j}=w_{j} \mathbf{R}_{i}^{j}
$$

We will use this convention for the remainder of this chapter.
We are now in a position to give a formal definition for a tensor.
Definition 6.2.7. A tensor of type $(p, q)$ has a multidimensional array associated with respect to the basis $\mathbf{f}$, denoted by $T_{j_{1}, \ldots, j_{q}}^{i_{1}, \ldots, i_{p}}[\mathbf{f}]$ which, under the basis transformation $\mathbf{f} \mapsto$ $\mathbf{f} \cdot \mathbf{R}=\left(\mathbf{e}_{i} \mathbf{R}_{1}^{i}, \ldots, \mathbf{e}_{i} \mathbf{R}_{n}^{i}\right)$ transforms as

$$
T_{j_{1}^{\prime}, \ldots, j_{q}^{\prime}}^{i_{1}^{\prime}, \ldots, i_{p}^{\prime}}[\mathbf{f} \cdot \mathbf{R}]=\left(R^{-1}\right)_{i_{1}}^{i_{1}^{\prime}} \cdot\left(R^{-1}\right)_{i_{p}}^{i_{p}^{\prime}} T_{j_{1}, \ldots, j_{q}}^{i_{1}, \ldots, i_{p}} R_{j_{1}^{\prime}}^{j_{1}} \cdot R_{j_{q}^{\prime}}^{j_{q}}
$$

In this definition, type $(p, q)$ means that the tensor has $p$ contravariant indices and $q$ covariant indices. We use the prime indices (e.g. $i_{1}^{\prime}$ ) to denote indices under the new basis. From here, we will usually drop the explicit reference to the basis in the notation, so we write (for example) $T_{j_{1}, \ldots, j_{q}}^{i_{1}, \ldots, i_{p}}$, and will use an asterisk $(*)$ to denote the tensor under a change of basis.

### 6.2.3 Vector Estimating Equations as Tensors

In our extension to tensor estimating equations, we will begin by considering the simple case of a vector estimating equation, and give the following theorem showing when such objects are tensors.

Theorem 6.2.3. If, for a given estimating equation $\mathbf{G}$, there exists a scalar potential function $g$ such that $\mathbf{G}=\nabla_{\theta} g$, then $\mathbf{G}$ obeys the tensor transformation laws over all differentiable transformations of the parameters.

Proof. Denote the original parameters by $\theta=\theta^{1}, \ldots, \theta^{p}$, and the new parameters by $\boldsymbol{\phi}=\phi^{1}, \ldots, \phi^{p}$. Let $\mathbf{a}: \Phi \rightarrow \Theta$ map each parameter $\boldsymbol{\phi}$ to the corresponding $\theta$. Then our potential function under change of variables is $g^{*}$, defined as $g^{*}:=g \circ \mathbf{a}$. The estimating equation generated by $g^{*}$ is $\mathbf{G}^{*}:=\nabla_{\phi}$. Then we have

$$
\begin{equation*}
\mathbf{G}_{r}^{*}=\mathbf{G}_{i} a_{r}^{i} \tag{6.2}
\end{equation*}
$$

where $a_{j}^{i}=\partial a^{i} / \partial \phi^{j}$. This is precisely the tensor transformation rule for a covector.
We must now check that standardisation of estimating equations preserves the tensor transformation rule. First, let us write the standardisation in tensor notation. We will denote by $\sigma_{i j}$ the covariance matrix $\mathbb{E}\left(\mathbf{G G} \mathbf{G}^{\mathrm{T}}\right)$, and by $\sigma^{i j}$ its inverse. The expectation of the first derivative, $\mathbb{E}(\dot{\mathbf{G}})$ will be denoted by $\gamma_{i j}$. Then we have:

$$
\begin{equation*}
\underline{\sim}_{r}=-\gamma_{i r} \sigma^{i j} \mathbf{G}_{j} \tag{6.3}
\end{equation*}
$$

and

$$
\begin{equation*}
\mathbf{G}_{r}=-\gamma^{i j} \sigma_{r j} \mathbf{G}_{i} \tag{6.4}
\end{equation*}
$$

Then, with $\tau_{i j}$ being the covariance matrix of $\mathbb{E}\left(\mathbf{G}^{*} \mathbf{G}^{* T}\right)$ and $\tau^{i j}$ its inverse and $\varsigma_{i j}$ the expectation of the first derivative, we also have

$$
\begin{equation*}
\underline{\mathbf{G}}_{r}^{*}=-\zeta_{i r} \tau^{i j} \mathbf{G}_{j}^{*} \tag{6.5}
\end{equation*}
$$

For later simplification, we first give a lemma regarding $\tau$ and $\varsigma$ :

Lemma 6.2.1. $\tau$, its inverse and $\varsigma$ are given by the following expressions

$$
\begin{align*}
\tau_{i j} & =\gamma_{i k} \gamma_{j l} \sigma^{k l}  \tag{6.6}\\
\tau^{i j} & =\gamma^{i k} \gamma^{j l} \sigma_{k l}  \tag{6.7}\\
\varsigma_{i j} & =\gamma_{i k} \gamma_{j l} \sigma^{k l} \tag{6.8}
\end{align*}
$$

We now give a theorem for the transformation law of standardised estimating equations.

Theorem 6.2.4. Standardised estimating equations follow a tensor transformation law.
Proof. Combining (6.2) and (6.5), we obtain:

$$
\begin{equation*}
\underline{\mathbf{G}}_{r}^{*}=-\varsigma_{i r} \tau^{i j}\left(\mathbf{G}_{k} a_{j}^{k}\right)=-a_{j}^{k} \tau^{i j} \varsigma_{i r} \mathbf{G}_{k}=c_{r}^{i} \mathbf{G}_{i} \tag{6.9}
\end{equation*}
$$

where $c_{r}^{i}:=-a_{j}^{i} \tau^{k j} \varsigma_{k r}$. Then, using (6.4), we obtain:

$$
\begin{equation*}
\underline{\mathbf{G}}_{r}^{*}=a_{j}^{i} \tau^{k j} \varsigma_{k r} \gamma^{l m} \sigma_{i m} \mathbf{G}_{l} \tag{6.10}
\end{equation*}
$$

Finally, using (6.8) and (6.7), we obtain

$$
\begin{equation*}
\underline{\mathbf{G}}_{r}^{*}=a_{j}^{i} \gamma^{l k} \gamma^{m j} \sigma_{l m} \gamma_{k n} \gamma_{r o} \sigma^{n o} \gamma^{p q} \sigma_{i q} \underline{\mathbf{G}}_{p}=a_{r}^{i} \sigma_{i k} \gamma^{k j} \underline{\mathbf{G}}_{j} \tag{6.11}
\end{equation*}
$$

Finally, we give the information criterion and optimality criterion in tensor notation:
Definition 6.2.8. Denoting, as in Theorem 6.2.4, the the covariance matrix $\mathbb{E}\left(\mathbf{G G}^{T}\right)$ by $\sigma_{i j}$, and by $\sigma^{i j}$ its inverse and $\mathbb{E}(\dot{\mathbf{G}})$ by $\gamma_{i j}$ then the information criterion $\mathscr{E}(\mathbf{G})$, denoted by $\mathscr{E}_{i j}$ is:

$$
\mathscr{E}_{i j}:=\gamma_{i k} \gamma_{j l} \sigma^{k l}
$$

Theorem 6.2.5. The information criterion is invariant under change of coordinates.
Proof. Using the previous notation, we have:

$$
\begin{equation*}
\mathscr{E}\left(\mathbf{G}^{*}\right)_{i j}=\varsigma_{k i} \tau^{k l} \varsigma_{l j} \tag{6.12}
\end{equation*}
$$

Substituting using (6.7) and (6.8), we obtain:

$$
\begin{align*}
\mathscr{E}\left(\mathbf{G}^{*}\right)_{i j}= & \gamma_{k m} \gamma_{i n} \sigma^{m n} \gamma^{k o} \gamma^{l p} \sigma_{o p} \gamma_{l q} \gamma_{j r} \sigma^{q r} \\
= & \gamma_{k m} \gamma^{k o} \gamma_{i n} \gamma^{k o} \gamma_{l q} \gamma^{l p} \sigma_{o p} \sigma^{q r} \\
& \mathscr{E}\left(\mathbf{G}^{*}\right)_{i j}=\gamma_{i k} \gamma_{j l} \sigma^{k l} \tag{6.13}
\end{align*}
$$

### 6.2.4 Tensor Estimating Equations

We are now ready to formulate the theory of tensor estimating equations. We will make the following necessary definitions, analogues of those given for vector estimating equations in the previous section.

Definition 6.2.9. A tensor estimating equation of order $\left(p_{1}, \ldots, p_{D}\right)$ is a multi-index array of functions $\mathbf{G}_{i_{1}, \ldots, i_{p}}: \mathscr{X} \times \Theta \rightarrow \mathbb{R}$, where the indices $i_{1}, \ldots, i_{D}$ run from 1 through to the corresponding element of $\left(p_{1}, \ldots, p_{D}\right)$. When there is no confusion, we will write $\mathscr{I}$ for $i_{1}, \ldots, i_{D}$, and similarly with other indices. Using this notation, we will write an estimating equation of order $\left(p_{1}, \ldots, p_{D}\right)$ as $\mathbf{G}_{\mathscr{G}}$. We require the following properties:

- There exists a potential $g$ such that $\mathbf{G}_{\mathscr{I}}=\frac{\partial g}{\partial \theta_{\mathscr{I}}}$
- $\mathbb{E}\left(\mathbf{G}_{\mathscr{I}}\right)=0$
- $\mathbb{E}\left(\dot{\mathbf{G}}_{\mathscr{I} \mathscr{G}}\right):=\mathbb{E}\left(\frac{\partial G_{\mathscr{g}}}{\partial \theta_{\mathscr{g}}}\right)$ is non-singular, i.e. invertible.

Remark. We remark that the first condition, the existence of a potential function $g$ is not standard in the consideration of general estimating equations. However, general multi-index arrays satisfying the conditions of estimating equations will fail the tensor transformation laws. The restriction to the class of multi-index arrays generated as the gradient of a scalar function guarantees the desired tensor properties, as we will prove. For its importance, we will refer to this condition as the assumption of potential (AOP).

Definition 6.2.10. The variance-covariance tensor for a tensor of order $\left(p_{1}, \ldots, p_{D}\right)$ is

$$
\begin{equation*}
\sigma_{\mathscr{I} \mathscr{F}}=\mathbb{E}\left(\mathbf{G}_{\mathscr{I}} \mathbf{G}_{\mathscr{F}}\right) \tag{6.14}
\end{equation*}
$$

Definition 6.2.11. Given an estimating equation $G$, the standardised estimating equation $\underset{\sim}{G}$ is defined by:

$$
\begin{equation*}
\mathbf{G}_{\mathscr{R}}:=-\gamma_{\mathscr{I} \mathscr{R}} \sigma^{\mathscr{I}} \mathbf{G}_{\mathscr{J}} \tag{6.15}
\end{equation*}
$$

where $\gamma_{\mathscr{I}}$ is the expected first derivative of $\mathbf{G}_{\mathscr{I}}$ with respect to $\theta_{\mathscr{g}}$.
Definition 6.2.12. The information criterion $\mathscr{E}_{\mathscr{I} \mathscr{G}}$ associated with an estimating equation $\mathbf{G}_{\mathscr{I}}$ is

$$
\begin{equation*}
\mathscr{E}_{\mathscr{I} \mathscr{I}}=\gamma_{\mathscr{I} \mathscr{K}} \gamma_{\mathscr{L} \mathscr{L}} \sigma^{\mathscr{K} \mathscr{L}} \tag{6.16}
\end{equation*}
$$

We are now ready to prove the tensor versions of the three theorems of Section 6.2.3. The proofs are almost identical, save for the additional complication that where we before largely dealt with pairs of indices, we are now dealing with pairs of index
sets. However, the choice of notation largely protects us from this additional complexity. Firstly, we will prove that our referring to these objects as tensor estimating equations is, in fact, justified.

Theorem 6.2.6. Tensor estimating equations are tensors.
Proof. Denote the original parameters by $\theta_{\mathscr{\mathscr { C }}}$, and the new parameters by $\phi_{\mathscr{G}}$. Let $a: \Phi \rightarrow \Theta$ map each parameter $\phi_{\mathscr{I}}$ to the corresponding $\theta_{\mathscr{G}}$. Here $a$ is a multidimensional array of the same dimensions as $\theta$. Then our potential function under change of variables is $g^{*}$, defined as $g^{*}:=g \circ a$. The estimating equation generated by $g^{*}$ is $\mathbf{G}^{*}:=\nabla_{\phi} g^{*}$. Then we have

$$
\begin{equation*}
\mathbf{G}_{\mathscr{R}}^{*}=\mathbf{G}_{\mathscr{P}} a_{\mathscr{R}}^{\mathscr{G}} \tag{6.17}
\end{equation*}
$$

where $a_{\mathscr{g}}^{\mathscr{g}}:=\partial a^{\mathscr{G}} / \partial \phi^{\mathscr{L}}=\left(\nabla_{\phi} a\right)_{\mathscr{g}}^{\mathscr{g}}$. Then we have that estimating equations transform according to the law for a covector; that is, they are tensors.

As in Section 6.2.3, we will now prove a short lemma giving some useful identities which will be necessary to prove that standardised estimating equations also followed a tensor transformation law.

Lemma 6.2.2. $\tau$, its inverse and $\varsigma$ are given by the following expressions

$$
\begin{align*}
& \tau_{\mathscr{G} \mathscr{G}}=\gamma_{\mathscr{H}} \gamma_{\mathscr{L} \mathscr{L}} \sigma^{\mathscr{L}}  \tag{6.18}\\
& \tau^{\mathscr{G}}=\gamma^{\mathscr{H}} \gamma^{\mathscr{L}} \sigma_{\mathscr{K} \mathscr{L}}  \tag{6.19}\\
& \varsigma_{\mathscr{G}}=\gamma_{\mathscr{H}} \gamma_{\mathscr{L} \mathscr{L}} \sigma^{\mathscr{L}} \tag{6.20}
\end{align*}
$$

Theorem 6.2.7. Standardised tensor estimating equations are tensors.
Proof. Denoting, as before, the variance tensor of the estimating equation under a change of coordinates $\left(\mathbf{G}_{* \mathscr{G}}\right)$ by $\varsigma_{\mathscr{g}}$ and the expected first derivative by $\tau_{\mathscr{\mathscr { L }}}$, we have:

$$
\begin{equation*}
{\underline{\mathbf{G}^{*}}}_{\mathscr{R}}=-\varsigma_{\mathscr{G} \mathscr{R}} \tau^{\mathscr{q}} \mathbf{G}_{\mathscr{G}}^{*} \tag{6.21}
\end{equation*}
$$

We will also need the following expression for the inverse of standardisation.

$$
\begin{equation*}
\mathbf{G}_{\mathscr{R}}=-\gamma^{\mathscr{I}} \sigma_{\mathscr{R}} \mathbf{G}_{\mathscr{I}} \tag{6.22}
\end{equation*}
$$

Combining (6.17) and (6.21), we obtain:

$$
\begin{equation*}
\mathbf{G}_{\mathscr{R}}^{*}=-\varsigma_{\mathscr{A}} \tau^{\mathscr{I}}\left(\mathbf{G}_{\mathscr{K}} a_{\mathscr{H}}^{\mathscr{K}}\right)=-a_{\mathscr{\mathscr { K }}}^{\mathscr{K}} \tau^{\mathscr{I}} \varsigma_{\mathscr{A} \mathscr{R}} \mathbf{G}_{\mathscr{K}}=c_{\mathscr{R}}^{\mathscr{G}} \mathbf{G}_{\mathscr{I}} \tag{6.23}
\end{equation*}
$$

where $c_{\mathscr{R}}^{\mathscr{I}}:=-a_{\mathscr{G}}^{\mathscr{G}} \tau^{\mathscr{K}} \varsigma_{\mathscr{K} \mathscr{R}}$. Then, using (6.22), we obtain:

$$
\begin{equation*}
\underline{\mathbf{G}}_{\mathscr{R}}^{*}=a_{\mathscr{I}}^{\mathscr{I}} \tau^{\mathscr{K} \mathscr{\mathscr { L }}} \varsigma_{\mathscr{K} \mathscr{R}} \gamma^{\mathscr{L} \mathscr{M}_{\mathscr{M}}} \sigma_{\mathscr{L}} \tag{6.24}
\end{equation*}
$$

Finally, using (6.20) and (6.19), we obtain

$$
\begin{equation*}
\underline{\mathbf{G}}_{\mathscr{R}}^{*}=a_{\mathscr{J}}^{\mathscr{I}} \gamma^{\mathscr{L} \mathscr{K}} \gamma^{\mathscr{M} \mathscr{\mathscr { L }}} \sigma_{\mathscr{L} \mathscr{M}} \gamma_{\mathscr{K} \mathscr{N}} \gamma_{\mathscr{R} \mathscr{O}} \sigma^{\mathscr{N}} \gamma^{\mathscr{P} \mathscr{Q}} \sigma_{\mathscr{I} \mathscr{Q}} \mathbf{G}_{\mathscr{P}}=a_{\mathscr{R}}^{\mathscr{I}} \sigma_{\mathscr{I} \mathscr{K}} \gamma^{\mathscr{K} \mathscr{J}} \underline{\mathbf{G}}_{\mathscr{I}} \tag{6.25}
\end{equation*}
$$

which is a tensor transformation law.

Finally, we will show that, as in the vector case, the information criterion remains an invariant under coordinate transformations. This result is fundamental to the use of estimating equations, reassuring us that our choice of coordinate system for parameter space does not influence the optimality of our estimating equations. As such, we are free to choose convenient systems of coordinates, whether they simplify the formulation of our models or of our estimating equations.

Theorem 6.2.8. The information criterion is an invariant.
Proof. As before, we have:

$$
\begin{equation*}
\mathscr{E}\left(\mathbf{G}^{*}\right)_{\mathscr{I} \mathscr{\mathscr { L }}}=\varsigma_{\mathscr{K} \mathscr{I}} \tau^{\mathscr{K} \mathscr{L}} \varsigma_{\mathscr{L} \mathscr{\mathscr { L }}} \tag{6.26}
\end{equation*}
$$

Using the expressions given in (6.19) and (6.20), we obtain:

$$
\begin{gather*}
\mathscr{E}\left(\mathbf{G}^{*}\right)_{\mathscr{I} \mathscr{I}}=\gamma_{\mathscr{K} \mathscr{M}} \gamma_{\mathscr{I} \mathscr{N}} \sigma^{\mathscr{M} \mathscr{N}} \gamma^{\mathscr{K} O} \gamma^{\mathscr{L} \mathscr{P}} \sigma_{\mathscr{O} \mathscr{P}} \gamma_{\mathscr{L} \mathscr{Q}} \gamma_{\mathscr{O} \mathscr{R}} \sigma^{\mathscr{Q} \mathscr{R}} \\
=\gamma_{\mathscr{K} \mathscr{M}} \gamma^{\mathscr{K} \mathscr{O}} \gamma_{\mathscr{I} \mathscr{N}} \gamma^{\mathscr{K} \mathscr{O}} \gamma_{\mathscr{L} \mathscr{Q}} \gamma^{\mathscr{L} \mathscr{P}} \sigma_{\mathscr{O} \mathscr{P}} \sigma^{\mathscr{Q} \mathscr{R}} \\
\mathscr{E}\left(\mathbf{G}^{*}\right)_{\mathscr{I} \mathscr{L}}=\gamma_{\mathscr{I} \mathscr{K}} \gamma_{\mathscr{I} \mathscr{L}} \sigma^{\mathscr{K} \mathscr{L}} \tag{6.27}
\end{gather*}
$$

which is precisely the expression for $\mathscr{E}(\mathbf{G})_{\mathscr{G} \mathscr{\mathscr { L }}}$, so the information criterion remains invariant under change of coordinates.

### 6.2.5 Asymptotic Consistency

In order to establish consistency of tensor estimating equations, we will begin by using the uniform law of large numbers to establish that there is a continuous function $\mathbb{E}(g(X, \theta))$ to which the sample mean of the estimating function $g(x, \theta)$ converges almost surely. First, we state the uniform law of large numbers for convenience.

Theorem 6.2.9 (Uniform Law of Large Numbers). Given

1. $\Theta$ is compact
2. The potential function $g(x, \theta)$ is continuous for each $\theta \in \Theta$ for almost all $x$
3. $g(x, \theta)$ is a measurable function of $x$ for each $\theta$
4. There is a dominating function $d(x)$ satisfying
a) $\mathbb{E}(d(x))<\infty$
b) $|g(x, \theta)| \leq d(x) \forall \theta \in \Theta$
then $\mathbb{E}(g(X, \theta))$ is continuous in $\theta$ and

$$
\sup _{\theta \in \Theta}\left|\frac{1}{T} \sum_{i=1}^{T} g\left(X_{i}, \theta\right)-\mathbb{E}(g(X, \theta))\right| \xrightarrow[\rightarrow]{\text { a.s. }} 0
$$

We now give a theorem using the uniform law of large numbers which establishes, under some conditions, that tensor estimating equations are consistent. At a high level, we establish that they are well-behaved extremum estimators, which Wald (1949) proved are consistent.

Theorem 6.2.10. Under the following assumptions, the tensor estimating equation estimate is consistent. We shall use $\theta_{0}$ to denote the true parameter value.
A. $1 \mathbb{E}_{\theta}\left(\left|g\left(x, \theta_{0}\right)\right|\right)<\infty$
A. $2 \mathbb{E}_{\theta}(g(x, \theta))<\infty$ for all $\theta \in \Theta$
A. 3 There exists a score function $U$ for the underlying distribution of the data, generated by log-likelihood function $\ell$
A. 4 The argument maxima of $\ell$ and $g$ coincide (i.e. both are maximised at $\theta_{0}$ )
A. 5 The parameter space $\Theta$ is a compact subset of $\mathbb{R}^{p_{1} \times \ldots \times p_{D}}$
A. $6 g(x, \theta)$ is continuous in $\theta$ over the parameter space for almost all values of $x$
A. $7 g(x, \theta)$ is a measurable function of $x$ for all $\theta \in \Theta$

Proof. We now need to establish a dominating function $d(x)$, as required for the Uniform Law of Large Numbers. Our candidate will be the function in A.1, i.e. we define

$$
d(x):=\mathbb{E}_{\theta}\left(g\left(x, \theta_{0}\right)\right)
$$

We seek to establish that

$$
\mathbb{E}_{\theta}(g(X, \theta))<\mathbb{E}_{\theta}\left(g\left(X, \theta_{0}\right)\right) \quad \forall \theta \neq \theta_{0}
$$

Now, if $\mathbb{E}_{\theta}(g(X, \theta))=-\infty$ then this holds trivially, so we will consider only the case when $\mathbb{E}_{\theta}(g(X, \theta))>-\infty$. Then by A. 2 we have that

$$
\mathbb{E}_{\theta}(|g(X, \theta)|)<\infty
$$

Now consider the random variable $m:=g(X, \theta)-g\left(X, \theta_{0}\right)$. We can write $g=\ell+f$, thus we have

$$
m(X)=\ell(X, \theta)-\ell\left(X, \theta_{0}\right)-\left(f\left(X, \theta_{0}\right)-f(X, \theta)\right)
$$

We now use the gradient theorem to write

$$
f\left(X, \theta_{0}\right)-f(X, \theta)=\int_{\gamma\left[\theta, \theta_{0}\right]} \nabla_{\theta} f_{\mathscr{I}}(X, r) d r_{\mathscr{I}}
$$

where $\gamma\left[\theta, \theta_{0}\right]$ is any path through the parameter space $\Theta$ from $\theta$ to $\theta_{0}$. Now write $f=g-\ell$, so we have

$$
f\left(X, \theta_{0}\right)-f(X, \theta)=\int_{\gamma\left[\theta, \theta_{0}\right]} \nabla_{\theta}(g-\ell)_{\mathscr{I}}(X, r) d r_{\mathscr{I}}=\int_{\gamma\left[\theta, \theta_{0}\right]}(G-U)_{\mathscr{I}}(X, r) d r_{\mathscr{I}}
$$

Then we consider the expectation

$$
\begin{aligned}
\mathbb{E}\left(f\left(X, \theta_{0}\right)-f(X, \theta)\right) & =\mathbb{E}\left(\int_{\gamma\left[\theta, \theta_{0}\right]}(G-U)_{\mathscr{I}}(X, r) d r_{\mathscr{I}}\right) \\
& =\int_{\gamma\left[\theta, \theta_{0}\right]} \mathbb{E}((G-U)(X, r))_{\mathscr{I}} d r_{\mathscr{I}} \\
& =\int_{\gamma\left[\theta, \theta_{0}\right]} 0_{\mathscr{I}} d r_{\mathscr{I}} \\
& =0
\end{aligned}
$$

Thus we have that

$$
\mathbb{E}(m(X))=\mathbb{E}\left(\ell(X, \theta)-\ell\left(X, \theta_{0}\right)\right)
$$

From Wald (1949), then, we have that $\mathbb{E}(m(X))<0$, i.e. that $d(X)$ is the dominating function we require. Thus, by the uniform law of large numbers, we have the result

$$
\sup _{\theta \in \Theta}\left|\frac{1}{T} \sum_{i=1}^{T} g\left(X_{i}, \theta\right)-\mathbb{E}(g(X, \theta))\right| \xrightarrow{\text { a.s. }} 0
$$

Then, as an extremum estimator, the estimating function consistently estimates its maximum $\theta_{0}$.

### 6.3 Categorising Generalisations of PCA

Having developed the theory of tensor estimating equations, we will now apply it to the methods of Landgraf and Lee (2015b), Collins et al. (2002) and Li and Tao (2013).

### 6.3.1 Generalised Principal Component Analysis

We now return to the Generalised Principal Component Analysis of Landgraf and Lee (2015b), as detailed in Section 4.2.

As we have a full expression for the likelihood of $\mathbf{U}$ given data $\mathbf{X}$ which is differentiable, we can construct the score function. For convenience, we write it as two separate estimating functions.

$$
\begin{align*}
& \mathbf{G}_{\mathbf{U}}=\left(\left(\mathbf{x}-b^{\prime}(\hat{\boldsymbol{\Theta}})\right)^{\mathrm{T}}\left(\tilde{\boldsymbol{\Theta}}-\mathbf{1} \boldsymbol{\mu}^{\mathrm{T}}\right)+\left(\tilde{\boldsymbol{\Theta}}-\mathbf{1} \boldsymbol{\mu}^{\mathrm{T}}\right)^{\mathrm{T}}\left(\mathbf{X}-b^{\prime}(\hat{\boldsymbol{\Theta}})\right)\right) \mathbf{U}  \tag{6.28}\\
& \mathbf{G}_{\mu}=\left(\mathbf{I}-\mathbf{U U}^{\mathrm{T}}\right)\left(\mathbf{X}-b^{\prime}(\hat{\boldsymbol{\Theta}})\right)^{\mathrm{T}} \mathbf{1} \tag{6.29}
\end{align*}
$$

where $\tilde{\boldsymbol{\Theta}}$ is the matrix of saturated natural parameters, $\hat{\boldsymbol{\Theta}}$ is given by $\mathbf{1} \boldsymbol{\mu}^{\mathrm{T}}+\mathbf{U} \mathbf{U}^{\mathrm{T}}\left(\tilde{\boldsymbol{\Theta}}-\mathbf{1} \boldsymbol{\mu}^{\mathrm{T}}\right)$ and $b^{\prime}$ is the inverse of the canonical link function.

As a score function, by Theorem 6.2.2 we know that this is an optimal estimating function if we choose $\mathscr{H}$ to include it. However, Generalised Principal Component Analysis (GPCA) requires that $\mathbf{U}$ satisfy the left semi-orthonormality constraint; as such we must add a non-linear constraint to be enforced. We denote the constraint by $\phi(\mathbf{U}):=\mathbf{U}^{\mathrm{T}} \mathbf{U}$ and require that our solution to (6.28) and (6.29) satisfies $\phi(\mathbf{U})=\mathbf{I}$. For theoretical simplicity, we can now write this as a third estimating function

$$
\begin{equation*}
\mathrm{G}_{\text {orth. }}=\mathbf{U}^{\mathrm{T}} \mathbf{U}-\mathbf{I} \tag{6.30}
\end{equation*}
$$

We can also think of this constraint as a change of space $\mathscr{H}$ on which we seek the optimal estimating function. In this case, the estimating functions (6.28) and (6.29) will again be optimal.

### 6.3.2 Collins et al.

This method, though typically formulated in terms of Bregmann distances, is ultimately based on an exponential family distribution, so we can be confident that it will satisfy the requirements of a tensor estimating equation. We now give the estimating functions.

$$
\begin{align*}
& \mathbf{G}_{\mathrm{A}}=\left(G^{\prime}(\mathbf{A V})-\mathbf{X}\right) \mathbf{V}^{\mathrm{T}}  \tag{6.31}\\
& \mathbf{G}_{\mathrm{V}}=\mathbf{A}^{\mathrm{T}}\left(G^{\prime}(\mathbf{A V})-\mathbf{X}\right) \tag{6.32}
\end{align*}
$$

From the estimating equations, we can see that, heuristically, we are attempting to make the difference between the observed data and the lower-dimensional approximation of it equal to zero, with some appropriate waiting (the $V^{T}$ and $\mathbf{A}^{T}$ terms respectively) in each of $\mathbf{G}_{\mathrm{A}}$ and $\mathrm{G}_{\mathrm{V}}$.

### 6.3.3 Simple Poisson Principal Component Analysis

Due to the complexity introduced by the unspecified functions involved in the general Simple Exponential PCA of Li and Tao (2013), we will give as an example the estimating functions of the Poisson case:

$$
\begin{align*}
\mathbf{G}_{\mathbf{W}} & =\mathbf{X Y}^{\mathrm{T}}-\exp (\mathbf{W Y}) \mathbf{Y}^{\mathrm{T}}-\mathbf{W} \operatorname{Diag}(\boldsymbol{\alpha})  \tag{6.33}\\
\mathbf{G}_{\mathbf{Y}} & =\mathbf{W}^{\mathrm{T}} \mathbf{X}-\mathbf{W}^{\mathrm{T}} \exp (\mathbf{W Y})-\mathbf{Y}  \tag{6.34}\\
\mathbf{G}_{\alpha j} & =p-\left\|\mathbf{w}_{j}\right\|^{2} \boldsymbol{\alpha}_{j} \tag{6.35}
\end{align*}
$$

Here, the exponential function is applied component-wise, and we use the convention of Li and Tao (2013) and Smallman, Underwood, et al. (2019) for the arrangement of the data matrix $\mathbf{X}$ as $n$ columns of observations. Note that (6.35) is derived from the approximation of $\boldsymbol{\alpha}$ proposed by Li and Tao, rather than the conditional posterior likelihood.

### 6.4 Comparisons

In estimating equation form, it is now easy to compare how these three methods different. Let us start by comparing GPCA and the method of Collins et al. (2002). If we assume $\mu=0$ in (6.28), then we have (with a change of sign)

$$
\mathbf{G}_{\mathbf{U}}=\left(b^{\prime}\left(\mathbf{U} \mathbf{U}^{\mathrm{T}} \tilde{\boldsymbol{\Theta}}\right)-\mathbf{X}\right)^{\mathrm{T}} \tilde{\boldsymbol{\Theta}} \mathbf{U}+\tilde{\boldsymbol{\Theta}}^{\mathrm{T}}\left(b^{\prime}\left(\mathbf{U} \mathbf{U}^{\mathrm{T}} \tilde{\boldsymbol{\Theta}}\right)-\mathbf{x}\right) \mathbf{U}
$$

Comparing this to (6.31) and (6.32), we then see that this estimating equation also considers some (differently weighted) difference between the lower-dimensional representation and the observations. The difference, then, is in the weighting and the use of $\mu$ as a centring parameter (which requires the inclusion of another estimating equation).

To compare Simple Exponential Principal Component Analysis (SePCA) with these methods, let us make some changes of notation. We rewrite $\exp (\mathbf{W Y})$ as $\mathbf{X}$, recalling that exp maps the natural parameters $\mathbf{W Y}$ to the expected value of the distribution, i.e.
our best guess for $\mathbf{X}$ without further information. Then we have

$$
\begin{aligned}
\mathbf{G}_{\mathbf{W}} & =\mathbf{X Y}^{\mathrm{T}}-\tilde{\mathbf{X}} \mathbf{Y}^{\mathrm{T}}-\mathbf{W} \operatorname{Diag}(\boldsymbol{\alpha}) & & =(\mathbf{X}-\tilde{\mathbf{X}}) \mathbf{Y}^{\mathrm{T}}-\mathbf{W} \operatorname{Diag}(\boldsymbol{\alpha}) \\
\mathbf{G}_{\mathbf{Y}} & =\mathbf{W}^{\mathrm{T}} \mathbf{X}-\mathbf{W}^{\mathrm{T}} \tilde{\mathbf{X}}-\mathbf{Y} & & =\mathbf{W}^{\mathrm{T}}(\mathbf{X}-\tilde{\mathbf{X}})-\mathbf{Y} \\
\mathbf{G}_{\alpha j} & =p-\left\|\mathbf{w}_{j}\right\|^{2} \boldsymbol{\alpha}_{j} & &
\end{aligned}
$$

In this form, we can see how SePCA relates to GPCA and Collins et al. (2002); once again we are looking at weighted differences between the estimated and observed data. In fact, if it were not for the $-\mathbf{W} \operatorname{Diag}(\boldsymbol{\alpha})$ ) and -Y terms in the first two equations, and the existence of the third equation, this would be precisely the method of Collins et al. (2002).

### 6.5 Asymptotic Consistency

None of the methods considered in this chapter have previously had any asymptotic results published in the literature, which is a significant disadvantage to the justification of their use.

Recall Theorem 6.2.10. Given that assumptions A. 1 through A. 7 hold, we can establish the asymptotic consistency of each of the three methods. It is easy to see that all but A. 4 hold for each of GPCA, SePCA and Collins et al. (2002). In general, A. 4 will likely have to be assumed, as it cannot be easily verified. However, in these cases, all three methods can be readily seen to be equivalent to maximum likelihood, so A. 4 will also hold. Precise details can be found in the respective papers, but briefly: GPCA is estimated by minimum deviance which is equivalent to maximum likelihood, SePCA is estimated by maximum a-posteriori (MAP) which is equivalent to maximum likelihood for the posterior likelihood, and Collins et al. (2002) is derived from maximum likelihood via Bregman divergences.

For illustration, we will briefly investigate each of these assumptions for the Poisson case of GPCA. Recall that here, the potential function is the deviance of a Poisson model.
A. 1 Write the deviance as

$$
\log \left(\frac{1}{x!}\right)+x \theta_{s}-\exp \theta_{s}-\log (1 x!)-x \theta_{0}-\exp \theta_{0}
$$

where $\theta_{s}$ is the saturated natural parameters. This simplifies to

$$
x\left(\theta_{s}-\theta_{0}\right)+\exp \theta_{s}-\exp \theta_{0}
$$

which is almost-everywhere finite for finite values of $\theta_{s}$ and $\theta_{0}$. Thus its absolute value is also almost everywhere finite, and it follows that so is the expectation
thereof. Of course, this is for the one-dimensional case, but the true vector case follows along the same lines with more complicated notation
A. 2 The first term of the deviance is the log-likelihood under the saturated model, which is finite and has finite expectation, so this assumption is equivalent to verifying that

$$
\mathbb{E}(-\log (\mathbb{P}(x, \theta)))
$$

is finite. In the Poisson case, this is

$$
-\mathbb{E}\left(\log \left(\frac{1}{x!}\right)+x \theta-\exp \theta\right)
$$

thus, as the expectation of an almost-everywhere finite quantity, it is also finite. Again, this is the one-dimensional case, but the true case follows simply from this
A. 3 Trivially true by assumption
A. 4 As we discussed, because the minimum deviance estimator is identical to the maximum likelihood estimator, we can be assured that this holds
A. 5 As the parameter space is simply $\mathbb{R}^{p_{1} \times \ldots \times p_{D}}$, this is trivially true
A. 6 The deviance function is everywhere continuous in $\theta$ for all values of $x$ (which are restricted to vectors of non-negative integers)
A. 7 Satisfied

### 6.6 Conclusions

In this chapter, we have seen how the framework of (tensor) estimating equations enables specifications of dimension reduction methods in a way which allows both for universal computational algorithms (via standard root-finding algorithms) and direct comparison between methods.

In particular, we have stated three important methods of exponential family dimension reduction in estimating equation form, and used that form to see how fundamentally similar these methods are. At their core, all of them are concerned by the difference between the observed and expected data, differing primarily in how they are weighted. This does, of course, provide a framework for exploring new dimension reduction techniques by modifying the estimating equations.

We have also used the asymptotic results for tensor estimating equations to show consistency of each of the three studied methods.

Any important future direction for this work is the ability to incorporate sparsifying penalties. This will allow almost all of the exponential family dimension reduction methods to be analysed under the same framework. Currently, the obstacle for this is the use of penalties which are not everywhere differentiable, such as the $L^{1}$ and SCAD penalties. This could, perhaps, be done by using differentiable approximations to these penalties, but it is not yet clear how to treat these in the asymptotic framework.

## — Chapter 7 —

## Conclusions

In this work, we have given three novel methods for exponential family dimension reduction. Firstly, we derived Poisson Inverse Regression (PoIR) in Chapter 2, which is specialised to the Poisson distribution for its use with text data. We gave two methods of estimation, using either maximum a-posteriori (MAP) or Markov Chain Monte Carlo (MCMC), and demonstrated its efficacy using real text data against its closest competitor, Multinomial Inverse Regression (MNIR).

We then extended the work of Landgraf and Lee (2015b) (known as Generalised Principal Component Analysis (GPCA)) in Chapter 4, incorporating a sparsifying penalty which improves performance and intepretability of the resulting loadings. Using both simulation studies and a real text data example, we showed that our proposed method performs on par or better than the current state of the art.

In Chapter 5, we extended Simple Exponential Principal Component Analysis (SePCA) (Li and Tao 2013) in a similar fashion, applying a sparsifying penalty to the loadings matrix. Again, we demonstrated the advantages of our new method with simulation and real data studies. In particular, we showed that our method performed better at estimating the underlying structural dimension of the data than the original method.

Taking a new approach, in Chapter 6 we presented a framework for tensor-valued estimating equations. We showed several important results, including that tensor estimating equations are asymptotically consistent. The development of this framework then allowed us to re-examine three important methods for exponential family dimension reduction in a new light in. We used the framework to show how essentially similar these methods are, and precisely how they differ. We also applied the asymptotic theory from tensor estimating equations to show that each of these methods is consistent.

## Bibliography

Bishop, Christopher M (1999). "Bayesian pca". In: Advances in neural information processing systems, pp. 382-388.
Blei, David M, Andrew Y Ng, and Michael I Jordan (2003). "Latent Dirichlet Allocation". In: Journal of Machine Learning Research 3, pp. 993-1022. arXiv: 1111.6189v1.
Collins, Michael, S Dasgupta, and Robert E Schapire (2002). "A Generalization of Principal Components Analysis to the Exponential Family". In: Advances in Neural Information Processing Systems 14, pp. 617-624.
Cook, R. Dennis (2007). "Fisher Lecture: Dimension Reduction in Regression". In: Statistical Science 22.1, pp. 1-26. arXiv: 0708.3774.
Eaton, Morris L (1986). "A characterization of spherical distributions". In: Journal of Multivariate Analysis 20.2, pp. 272-276.
Fan, Jianqing and Runze Li (2001). "Variable Selection via Nonconcave Penalized Likelihood and its Oracle Properties". In: Journal of the American Statistical Association 96.456, pp. 1348-1360.

Frommlet, Florian and Grégory Nuel (2016). "An adaptive ridge procedure for L0 regularization". In: PLoS ONE 11.2, pp. 1-23.
Gelman, A and D B Rubin (1992). "A Single Series from the Gibbs Sampler Provides a False Sense of Security". In: Bayesian Statistics 4.July, pp. 625-631.
Geman, Stuart and Donald Geman (1984). "Stochastic relaxation, Gibbs distributions, and the Bayesian restoration of images". In: IEEE Transactions on pattern analysis and machine intelligence 6, pp. 721-741.
Gillis, Nicolas (2014). "The Why and How of Nonnegative Matrix Factorization". In: pp. 1-25. arXiv: 1401.5226.
Guan, Y and Jg Dy (2009). "Sparse probabilistic principal component analysis". In: Proceedings of the 12th International Conference on Artificial Intelligence and Statistics, pp. 185-192.
Hastings, W Keith (1970). "Monte Carlo sampling methods using Markov chains and their applications". In:

Hotelling, Harold (1933). "Analysis of a complex of statistical variables into principal components." In: Journal of Educational Psychology 24.6, pp. 417-441.
Hu, Zhenfang et al. (2016). "Sparse Principal Component Analysis via Rotation and Truncation". In: IEEE Transactions on Neural Networks and Learning Systems 27.4, pp. 875-890. arXiv: arXiv:1403.1430v2.
Hunter, David R and Runze Li (2005). "Variable selection using MM algorithms". In: Annals of Statistics 33.4, pp. 1617-1642. arXiv: 0508278 [math].
Kwak, Nojun (2008). "Principal component analysis based on L1-norm maximization". In: IEEE Transactions on Pattern Analysis and Machine Intelligence 30.9, pp. 16721680.

Landgraf, Andrew J and Yoonkyung Lee (Oct. 2015a). "Dimensionality Reduction for Binary Data through the Projection of Natural Parameters". In: arXiv:1510.06112 [stat.ML]. arXiv: 1510.06112.

- (2015b). "Generalized Principal Component Analysis: Projection of Saturated Model Parameters". In: Ohio State University Statistics Department Technical Report 892.892.
Li, Jun and Dacheng Tao (2013). "Simple exponential family PCA". In: IEEE Transactions on Neural Networks and Learning Systems 24.3, pp. 485-497.
Li, Ker-Chau (1991). "Sliced inverse regression for dimension reduction". In: Journal of the American Statistical Association 86.414, pp. 316-327.

Lu, Meng, Jianhua Z. Huang, and Xiaoning Qian (Dec. 2016). "Sparse exponential family Principal Component Analysis". In: Pattern Recognition 60, pp. 681-691.
Mackay, David J C (1995). "Probable networks and plausible predictionsa review of practical Bayesian methods for supervised neural networks". In: Network: Computation in Neural Systems 6.3, pp. 469-505.
Metropolis, Nicholas et al. (1953). "Equation of state calculations by fast computing machines". In: The journal of chemical physics 21.6, pp. 1087-1092.
Mohamed, Shakir, Katherine Heller, and Zoubin Ghahramani (2009). "Bayesian exponential family PCA". In: Advances in Neural Information Processing Systems, pp. 10891096.

Nocedal, J and S J Wright (1999). Numerical Optimization. Vol. 43. 2, pp. 164-75. arXiv: NIHMS150003.

Pearson, Karl (1901). "On lines and planes of closest fit to systems of points in space". In: The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science 2.1, pp. 559-572.

R Foundation for Statistical Computing, Vienna, Austria. ISBN 3-900051-07-0 (2011). "R Development Core Team". In: R: A Language and Environment for Statistical Computing 55, pp. 275-286.

Rousseeuw, Peter J. (Nov. 1987). "Silhouettes: A graphical aid to the interpretation and validation of cluster analysis". In: Journal of Computational and Applied Mathematics 20, pp. 53-65. arXiv: z0024.

Smallman, Luke, Andreas Artemiou, and Jennifer Morgan (2018). "Sparse generalised principal component analysis". In: Pattern Recognition 83, pp. 443-455.
Smallman, Luke, William Underwood, and Andreas Artemiou (2019). "Simple Poisson PCA: an algorithm for (sparse) feature extraction with simultaneous dimension determination". In: Computational Statistics, pp. 1-19.
Taddy, Matt (2013). "Multinomial Inverse Regression for Text Analysis". In: Journal of the American Statistical Association 108.503, pp. 755-770. arXiv: 1012.2098.

- (Sept. 2015). "Distributed multinomial regression". In: The Annals of Applied Statistics 9.3, pp. 1394-1414.
Tipping, Michael E. and Christopher M. Bishop (1999). "Probabilistic Principal Component Analysis". In: Journal of the Royal Statistical Society. Series B (Statistical Methodology) 61.3, pp. 611-622.
Wald, Abraham (1949). "Note on the consistency of the maximum likelihood estimate". In: The Annals of Mathematical Statistics 20.4, pp. 595-601.
Wen, Zaiwen and Wotao Yin (Dec. 2013). "A feasible method for optimization with orthogonality constraints". In: Mathematical Programming 142.1-2, pp. 397-434.
Yi, Shuangyan et al. (Jan. 2017). "Joint sparse principal component analysis". In: Pattern Recognition 61, pp. 524-536.
Zou, Hui and Trevor Hastie (2005). "Regularization and Variable Selection via the Elastic Net". In: Journal of the Royal Statistical Society. Series B (Methodological) 67.2, pp. 301-320.
Zou, Hui, Trevor Hastie, and Robert Tibshirani (June 2006). "Sparse Principal Component Analysis". In: Journal of Computational and Graphical Statistics 15.2, pp. 265286. eprint: 1205.0121 v 2.


[^0]:    "Begin at the beginning," the King said gravely, "and go on till you come to the end: then stop."

    - Lewis Carroll, Alice in Wonderland

