

*Covariance Structure Regularization via Entropy
Loss Function*

Lin, Lijing and Higham, Nicholas J. and Pan, Jianxin

2013

MIMS EPrint: **2012.61**

Manchester Institute for Mathematical Sciences
School of Mathematics

The University of Manchester

Reports available from: <http://eprints.maths.manchester.ac.uk/>

And by contacting: The MIMS Secretary
School of Mathematics
The University of Manchester
Manchester, M13 9PL, UK

ISSN 1749-9097

Covariance Structure Regularization via Entropy Loss Function

Lijing Lin, Nicholas J. Higham and Jianxin Pan¹
*School of Mathematics, University of Manchester,
Manchester M13 9PL, UK*

Abstract

The need to estimate structured covariance matrices arises in a variety of applications and the problem is widely studied in statistics. A new method is proposed for regularizing the covariance structure of a given covariance matrix whose underlying structure has been blurred by random noise, particularly when the dimension of the covariance matrix is high. The regularization is made by choosing an optimal structure from an available class of covariance structures in terms of minimizing the discrepancy, defined via the entropy loss function, between the given matrix and the class. A range of potential candidate structures comprising tridiagonal Toeplitz, compound symmetry, AR(1), and banded Toeplitz is considered. It is shown that for the first three structures local or global minimizers of the discrepancy can be computed by one-dimensional optimization, while for the fourth structure Newton's method enables efficient computation of the global minimizer. Simulation studies are conducted, showing that the proposed new approach provides a reliable way to regularize covariance structures. The approach is also applied to real data analysis, demonstrating the usefulness of the proposed approach in practice.

Keywords: Covariance estimation; Covariance structure; Entropy loss function; Kullback-Leibler divergence; Regularization.

1. Introduction

The need to estimate structured covariance matrices arises in a variety of application fields including signal processing (Pascal et al., 2008), networks (Vinciotti and Hashem, 2013), and control problems (Lin and Jovanović, 2009) and the problem is widely studied in statistics; see, e.g., Pourahmadi (1999) and Pan and Mackenzie (2003). A conventional way, known as the “Burg technique”, is to find the maximum likelihood estimation for a covariance matrix that has a specific/regularized

¹Corresponding author. Tel: +44 161 2755864; Fax: +44 161 275 5819.
Email Addresses: Lijing.Lin@manchester.ac.uk (L. Lin); Nick.Higham@manchester.ac.uk (N. J. Higham); Jianxin.Pan@manchester.ac.uk (J. Pan)

structure using random samples drawn from a stochastic process (Burg et al., 1982). However, this method has some drawbacks, including that (a) it is based on the presumption that the stochastic process is multivariate normal, (b) the structure of the covariance must be prespecified, and (c) the sample covariance matrix must be available; moreover, it can be difficult to deduce the underlying covariance structure from the sample covariance matrix, because of random noise or large dimension of the matrix, for example.

To overcome these difficulties, in this paper we propose a new method for regularizing the underlying structure of a given covariance matrix. Our method is based on the entropy loss function (Dey and Srinivasan, 1985; James and Stein, 1961)

$$L(A, B) = \text{tr}(A^{-1}B) - \log(\det(A^{-1}B)) - m, \quad (1.1)$$

where A and B are $m \times m$ matrices and, to ensure that $L(A, B)$ is nonnegative, we assume that A and B are symmetric positive definite. The entropy loss function, also known as the Kullback-Leibler divergence, is a well-accepted nonsymmetric measure of the discrepancy between two probability distributions (Pan and Fang, 2002). It is a special case of the Bregman divergence (Dhillon and Tropp, 2007) and has been widely used in statistics (Pan and Fang, 2002). Recent work on using the entropy loss function in multivariate spectral estimation and multivariate process control is that of Ferrante et al. (2012) and Maboudou-Tchao and Agboto (2013), respectively. The problem of interest here is, given a covariance matrix A whose underlying structure is blurred due to random noise, particularly when the dimension m is high, to identify the underlying structure of A from a class of candidate covariance structures. To demonstrate our idea, we introduce the following notation. Let \mathcal{S} be the set of all positive definite covariance matrices with structure s . We define the discrepancy between a given positive definite covariance matrix A and the set \mathcal{S} by

$$D(A, \mathcal{S}) = \min_{B \in \mathcal{S}} L(A, B), \quad (1.2)$$

where $L(A, B)$ is the entropy loss function in (1.1). Our idea is that, among a given class of k candidate covariance structures $\{s_1, s_2, \dots, s_k\}$, the structure with which A has the smallest discrepancy can be viewed as the most likely underlying structure of A . We refer to the replacement of A by a matrix B achieving the minimum in (1.2) as the process as regularizing A . It is worth pointing out that the matrix A is not necessarily a sample covariance matrix. It can be any estimator of a covariance matrix, obtained by statistical methods such as those based on modified Cholesky decomposition methods (Pan and Mackenzie, 2003; Ye and Pan, 2006). Regularization of the given covariance matrix helps the understanding of the underlying correlation/covariance process and simplifies complicated and high-dimensional data problems.

In this paper we consider the following four candidate covariance structures that are commonly used in practice, for example, in longitudinal and spatial studies.

- (1) The order-1 moving average structure, MA(1), has a tridiagonal and Toeplitz covariance matrix

$$B = \sigma^2 \begin{bmatrix} 1 & c & 0 & \cdots & 0 \\ c & 1 & c & \ddots & \vdots \\ 0 & \ddots & \ddots & \ddots & 0 \\ \vdots & \ddots & \ddots & 1 & c \\ 0 & \cdots & 0 & c & 1 \end{bmatrix}, \quad (1.3)$$

where $\sigma^2 > 0$ and $-1/(2 \cos(\pi/(m+1))) < c < 1/(2 \cos(\pi/(m+1)))$.

- (2) The covariance of compound symmetry (CS) structure assumes that the correlation coefficients of any two observations are the same, i.e.,

$$B = \sigma^2 \begin{bmatrix} 1 & c & c & \cdots & c \\ c & 1 & c & \ddots & \vdots \\ c & \ddots & \ddots & \ddots & c \\ \vdots & \ddots & \ddots & 1 & c \\ c & \cdots & c & c & 1 \end{bmatrix}, \quad (1.4)$$

where $\sigma^2 > 0$ and $-1/(m-1) < c < 1$.

- (3) The covariance of autoregression of order 1, AR(1), has the property that the correlation between any pair of observations decays exponentially towards zero as the distance between two observations increases. It is of the form

$$B = \sigma^2 \begin{bmatrix} 1 & c & c^2 & \cdots & c^{m-1} \\ c & 1 & c & \cdots & c^{m-2} \\ c^2 & c & 1 & \cdots & c^{m-3} \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ c^{m-1} & c^{m-2} & \cdots & c & 1 \end{bmatrix}, \quad (1.5)$$

where $\sigma^2 > 0$ and $-1 < c < 1$.

- (4) More generally, banded Toeplitz covariance matrices have constant subdiagonal entries, i.e., constants at lag 1 (i.e., the covariance of the repeated measures one unit of time apart), lag 2, ..., and lag p :

$$B = \sigma^2 \begin{bmatrix} 1 & c_1 & \cdots & c_p & 0 & \cdots & 0 \\ c_1 & 1 & c_1 & \ddots & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & \ddots & \ddots & 0 \\ c_p & \ddots & \ddots & \ddots & \ddots & \ddots & c_p \\ 0 & \ddots & \ddots & \ddots & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & \ddots & 1 & c_1 \\ 0 & \cdots & 0 & c_p & \cdots & c_1 & 1 \end{bmatrix}, \quad (1.6)$$

where $\sigma^2 > 0$, c_1, c_2, \dots, c_p are nonzero and all other off-diagonal elements are zero.

The main task now is to calculate the discrepancy $D(A, \mathcal{S})$ for each of the candidate covariance structures listed in (1.3)–(1.6) above, where the covariance matrix A is given. Equivalently, it is to find for each covariance structure a positive definite matrix B that minimizes the discrepancy $L(A, B)$ over the set of matrices with that structure. Accordingly, structure s in the candidate class $\{s_1, \dots, s_k\}$ that has the smallest discrepancy is the most likely covariance structure, among the candidate classes, for the matrix A .

The rest of this paper is organized as follows. In section 2, we formulate our problem of interest into an optimization problem and explore some of its general properties. We show that for \mathcal{S} the set of MA(1), CS or Toeplitz covariance matrices, the optimization problem we need to solve is a convex problem and has a unique global minimum. We then show in section 3 that the problem of finding B with structure MA(1), CS or AR(1) that minimizes $L(A, B)$ reduces to computing the zeros of a nonlinear function in one variable. In section 4 we explain how Newton’s method can be used to solve the problem for Toeplitz covariance structure. In section 5 we carry out simulation studies, illustrating how our techniques of computing the structured covariance matrix that minimizes the entropy loss function can be used in regularizing the underlying covariance structure. We also apply the proposed approach to two real data experiments. Some further remarks and discussion are given in section 6.

2. Problem of interest and its properties

We start by formulating the problem of interest and exploring some of its properties. Define $f : \mathbb{R}_+^{m \times m} \rightarrow \mathbb{R}$ where $\mathbb{R}_+^{m \times m}$ is the set of all $m \times m$ symmetric positive definite matrices and $f(B) := L(A, B)$. Let $\Omega \subset \mathbb{R}_+^{m \times m}$ be a set of structured positive definite matrices. Our problem now is

$$\min \quad f(B) \tag{2.7a}$$

$$\text{subject to } B \in \Omega. \tag{2.7b}$$

We denote by $\nabla_B f = (\partial f / \partial b_{ij})$ the gradient of f , where b_{ij} is the (i, j) entry of B . Ignoring the symmetry of A and B and using results from Magnus and Neudecker (1999) we have

$$\begin{aligned} \nabla_B \text{tr}(A^{-1}B) &= A^{-T} = A^{-1}, \\ \nabla_B \log \det(B) &= B^{-T} = B^{-1}, \end{aligned}$$

and then

$$\nabla_B f = A^{-1} - B^{-1}. \tag{2.8}$$

Write $b = \text{vec}(B) \in \mathbb{R}^{m^2}$ where vec denotes the vector obtained by stacking the columns of its matrix argument on top of each other from first to last. Taking f as

a function from \mathbb{R}^{m^2} to \mathbb{R} , the Hessian of f is then given by

$$\nabla_b^2 f := \left(\frac{\partial^2 f}{\partial b_i \partial b_j} \right) = B^{-T} \otimes B^{-1} = B^{-1} \otimes B^{-1}, \quad (2.9)$$

(Magnus and Neudecker, 1999). Since B is positive definite, $B^{-1} \otimes B^{-1}$ is positive definite, and so $f(B)$ is a strictly convex function of B . The strict convexity of $f(B)$ in $\mathbb{R}_+^{m \times m}$ is also a standard result that follows directly from its definition in Boyd and Vandenberghe (2004, Sec. 3). We nevertheless keep our brief derivation here for further reference.

On the other hand, it is clear from their expressions that the sets Ω of MA(1) (1.3), CS (1.4), and Toeplitz (1.6) are convex. Therefore when Ω is the set of positive definite matrices having one of the three structures the problem (2.7) is convex and so has a unique solution. When Ω is the set of AR(1) matrices, the problem is not convex because Ω is not convex. We will show later that only a local minimum of the problem can be expected to be found in this case.

We mention in passing that when $\Omega = \mathbb{R}_+^{m \times m}$, the minimum of $f(B)$ in (2.7) is obtained at $\nabla_B f = 0$, i.e., $B = A$. Moreover, provided A is symmetric positive definite, the entropy loss function $L(A, B)$ at the boundary of the set of symmetric positive definite matrices is $+\infty$, that is, $L(A, B) \rightarrow +\infty$ as $\det(B) \rightarrow 0$.

3. Two-parameter problems

We begin by considering the two-parameter matrices (1.3)–(1.5), for which the problems reduces to computing the zeros of a nonlinear function of a single variable.

3.1. Tridiagonal Toeplitz matrices

Recall that the tridiagonal matrix (1.3)

$$B(c, \sigma) = \sigma^2 \begin{bmatrix} 1 & c & 0 & \cdots & 0 \\ c & 1 & c & \ddots & \vdots \\ 0 & \ddots & \ddots & \ddots & 0 \\ \vdots & \ddots & \ddots & 1 & c \\ 0 & \cdots & 0 & c & 1 \end{bmatrix},$$

can be rewritten

$$B(c, \sigma) = \sigma^2(I + cT_1), \quad (3.10)$$

where T_1 is a symmetric matrix with the first superdiagonal and subdiagonal equal to 1 and all other elements equal to 0. Note that the eigenvalues of $B(c, \sigma)$ are (Higham, 2002, Sec. 28.5)

$$\lambda_j = \sigma^2(1 + 2cs_j), \quad j = 1 : m,$$

where $s_j = \cos(\pi j/(m+1))$. Assuming $m \geq 2$, we have $s_1 > \dots \geq 0 \geq \dots > s_m$, $s_j = -s_{m+1-j}$ and hence $B(c, \sigma)$ is positive definite if and only if $\lambda_1 > 0$ and $\lambda_m > 0$, which gives

$$-\frac{1}{2s_1} < c < \frac{1}{2s_1}. \quad (3.11)$$

Given a positive definite covariance matrix A , the loss function is now

$$f(c, \sigma) := \sigma^2 \text{tr}(A^{-1}) + c\sigma^2 \text{tr}(A^{-1}T_1) + \log(\det(A)) - m \log \sigma^2 - \sum_{j=1}^m \log(1 + 2cs_j) - m. \quad (3.12)$$

It follows that

$$\nabla f := \begin{bmatrix} \frac{\partial f}{\partial c} \\ \frac{\partial f}{\partial \sigma} \end{bmatrix} = \begin{bmatrix} \sigma^2 \text{tr}(A^{-1}T_1) - \sum_{j=1}^m \frac{2s_j}{1 + 2cs_j} \\ 2\sigma \text{tr}(A^{-1}) - 2m/\sigma + 2c\sigma \text{tr}(A^{-1}T_1) \end{bmatrix}, \quad (3.13)$$

and

$$\nabla^2 f := \begin{bmatrix} \frac{\partial^2 f}{\partial c^2} & \frac{\partial^2 f}{\partial c \partial \sigma} \\ \frac{\partial^2 f}{\partial c \partial \sigma} & \frac{\partial^2 f}{\partial \sigma^2} \end{bmatrix} = \begin{bmatrix} \sum_{j=1}^m \frac{(2s_j)^2}{(1 + 2cs_j)^2} & 2\sigma \text{tr}(A^{-1}T_1) \\ 2\sigma \text{tr}(A^{-1}T_1) & 2\text{tr}(A^{-1}) + \frac{2m}{\sigma^2} + 2c \text{tr}(A^{-1}T_1) \end{bmatrix}. \quad (3.14)$$

The stationary points (c, σ) following from $\nabla f = 0$ in (3.13) satisfy the equations

$$\begin{cases} \sigma^2 = \sum_{j=1}^m \frac{2s_j}{1 + 2cs_j} / \text{tr}(A^{-1}T_1), \\ h(c) := \sum_{j=1}^m \frac{2s_j}{1 + 2cs_j} - \frac{m \text{tr}(A^{-1}T_1)}{\text{tr}(A^{-1}) + c \text{tr}(A^{-1}T_1)} = 0. \end{cases} \quad (3.15)$$

Recalling (3.11), as $c \rightarrow -1/(2s_1)$ we have $2s_1/(1 + 2cs_1) \rightarrow +\infty$, so $h(c) \rightarrow +\infty$, while as $c \rightarrow 1/(2s_1) = -1/(2s_m)$ we have $2s_m/(1 + 2cs_m) \rightarrow -\infty$ and so $h(c) \rightarrow -\infty$. Therefore there exists at least one zero of $h(c)$ on $[-1/(2s_1), 1/(2s_1)]$. Since for every c satisfying $h(c) = 0$,

$$\begin{aligned} h'(c) &= -\sum_{j=1}^m \frac{(2s_j)^2}{(1 + 2cs_j)^2} + \frac{m(\text{tr}(A^{-1}T_1))^2}{(\text{tr}(A^{-1}) + c \text{tr}(A^{-1}T_1))^2} \\ &= -\sum_{j=1}^m \frac{(2s_j)^2}{(1 + 2cs_j)^2} + \frac{1}{m} \left(\sum_{j=1}^m \frac{2s_j}{1 + 2cs_j} \right)^2 < 0, \end{aligned}$$

where the last inequality is from the Cauchy-Schwarz inequality $|x^T e| \leq \|x\|_2 \|e\|_2$ with $e = [1, 1, \dots, 1]^T$, there exists only one zero of $h(c)$ on $[-1/(2s_1), 1/(2s_1)]$ and thus a unique point (c, σ^2) satisfying (3.15).

It is clear that the (1, 1) element of the Hessian matrix $\nabla^2 f$ in (3.14) is positive and it can be easily verified that for any (c, σ) satisfying (3.15), we have

$$\begin{aligned} \det(\nabla^2 f(c, \sigma)) &= \frac{4m}{\sigma^2} \sum_{j=1}^m \frac{(2s_j)^2}{(1 + 2cs_j)^2} - 4\sigma^2(\text{tr}(A^{-1}T_1))^2 \\ &= \frac{4}{\sigma^2} \left(m \sum_{j=1}^m \frac{(2s_j)^2}{(1 + 2cs_j)^2} - \sigma^4(\text{tr}(A^{-1}T_1))^2 \right) \\ &= \frac{4}{\sigma^2} \left(m \sum_{j=1}^m \frac{(2s_j)^2}{(1 + 2cs_j)^2} - \left(\sum_{j=1}^m \frac{2s_j}{1 + 2cs_j} \right)^2 \right) \\ &> 0, \end{aligned}$$

where the last inequality is from the Cauchy–Schwarz inequality. Therefore the Hessian matrix $\nabla^2 f$ is positive definite and so the stationary point is a minimum point. Note that $f(c, \sigma)$ is defined on the open set $\{(c, \sigma) : \sigma > 0 \text{ and (3.11) holds}\}$ on which $B(c, \sigma)$ is positive definite. As (c, σ) approaches the boundary of the set or as $\sigma \rightarrow \infty$, we have $f(c, \sigma) \rightarrow +\infty$. It follows immediately that the global minimum of $f(c, \sigma)$ is obtained at the unique stationary point. We summarize the discussion above in the following theorem.

Theorem 3.1 *Given a positive definite covariance matrix A , there exists a unique tridiagonal positive definite matrix $B(c, \sigma)$ of the form (3.10) that minimizes the loss function $f(c, \sigma) := L(A, B(c, \sigma))$ given by (3.12). Furthermore, the minimum is attained at (c, σ) satisfying (3.15).*

We note that by setting $a = \sigma^2$ and $b = \sigma^2 c$ the tridiagonal matrix (1.3) can be cast into the equivalent matrix

$$B(a, b) = \begin{bmatrix} a & b & \dots & 0 \\ b & a & \dots & 0 \\ \vdots & \ddots & \ddots & b \\ 0 & \dots & b & 1 \end{bmatrix}.$$

In this case, the matrix $B(a, b)$ is a linear function in a and b and the function $f(B)$ is a convex function at (a, b) , so that the uniqueness of the global minimum with respect to (a, b) becomes straightforward. However, an explicit expression of the stationary point in (a, b) is not available. In contrast, based on the (c, σ) -parametrization the optimization problem for $f(c, \sigma)$ reduces to finding the root of a one-dimensional nonlinear function $h(c)$ since the optimal value of σ^2 has an explicit solution. More importantly, the (c, σ) -parametrization form in (1.3) is a meaningful representation in statistics, that is, σ^2 is the variance of the response variable and c is the correlation coefficient of any adjacent pair of observations.

3.2. Compound symmetry

The matrix in (1.4) can be rewritten as

$$B(c, \sigma) = \sigma^2 \begin{bmatrix} 1 & c & \cdots & c \\ c & 1 & \cdots & c \\ \vdots & \ddots & \ddots & \vdots \\ c & \cdots & c & 1 \end{bmatrix} = \sigma^2(I + c(ee^T - I)), \quad (3.16)$$

where $e = [1, \dots, 1]^T \in \mathbb{R}^m$. The eigenvalues of $B(c, \sigma)$ are $\sigma^2(1 + (m-1)c)$ and $\sigma^2(1 - c)$ of multiplicities 1 and $m-1$, respectively, so $B(c, \sigma)$ is a positive definite matrix if and only if (Borsdorf et al., 2010, Lem. 2.1)

$$-\frac{1}{m-1} < c < 1.$$

Given A , we define $f(c, \sigma) := L(A, B(c, \sigma))$, where $L(A, \cdot)$ is the entropy loss function in (1.1). We want to find an explicit solution to the corresponding optimization problem

$$\min_{\substack{\sigma > 0 \\ -1/(m-1) < c < 1}} f(c, \sigma). \quad (3.17)$$

First, it is clear that $\det(B(c, \sigma)) = \sigma^{2m}(1 - c)^{m-1}(1 + (m-1)c)$. Denoting $t := \text{tr}(A^{-1}(ee^T - I))$, we have

$$\begin{aligned} f(c, \sigma) &= \sigma^2 \text{tr}(A^{-1}) + c\sigma^2 t + \log(\det(A)) - m \log(\sigma^2) \\ &\quad - (m-1) \log(1 - c) - \log(1 + (m-1)c) - m. \end{aligned}$$

Second, we have the gradient of f

$$\nabla f := \begin{bmatrix} \frac{\partial f}{\partial c} \\ \frac{\partial f}{\partial \sigma} \end{bmatrix} = \begin{bmatrix} \sigma^2 t + \frac{m-1}{1-c} - \frac{m-1}{1+(m-1)c} \\ 2\sigma \text{tr}(A^{-1}) + 2\sigma ct - 2m/\sigma \end{bmatrix}, \quad (3.18)$$

and the Hessian matrix

$$\nabla^2 f := \begin{bmatrix} \frac{\partial^2 f}{\partial c^2} & \frac{\partial^2 f}{\partial c \partial \sigma} \\ \frac{\partial^2 f}{\partial c \partial \sigma} & \frac{\partial^2 f}{\partial \sigma^2} \end{bmatrix} = \begin{bmatrix} \frac{m-1}{(1-c)^2} + \frac{(m-1)^2}{(1+(m-1)c)^2} & 2\sigma t \\ 2\sigma t & 2(\text{tr}(A^{-1}) + ct) + \frac{2m}{\sigma^2} \end{bmatrix}. \quad (3.19)$$

Therefore, the stationary points (c, σ) of $f(c, \sigma)$ must satisfy

$$\begin{cases} h(c) := \sigma^2 t + \frac{m-1}{1-c} - \frac{m-1}{1+(m-1)c} = 0 \\ \frac{m}{\sigma^2} = \text{tr}(A^{-1}) + ct. \end{cases} \quad (3.20)$$

Next, it is clear that $h(c)$ is continuous in the interval $(-1/(m-1), 1)$. Since $h(c) \rightarrow +\infty$ as $c \rightarrow 1$ and $h(c) \rightarrow -\infty$ as $c \rightarrow -1/(m-1)$, there exists at least one solution to $h(c) = 0$. The stationary points are obtained immediately by solving $h(c) = 0$, which gives $c = -t/((m-1)\text{tr}(A^{-1}) + (m-2)t)$.

Finally, since

$$(\nabla^2 f)_{11} = \frac{m-1}{(1-c)^2} + \frac{(m-1)^2}{(1+(m-1)c)^2} > 0$$

at the stationary points (c, σ) satisfying (3.20), and we also have

$$\begin{aligned} \det(\nabla^2 f) &= \frac{4m}{\sigma^2} \left(\frac{m-1}{(1-c)^2} + \frac{(m-1)^2}{(1+(m-1)c)^2} \right) - 4\sigma^2 t^2 \\ &= \frac{4}{\sigma^2} (m-1) \left(\frac{m-1}{1+(m-1)c} + \frac{1}{1-c} \right)^2 \\ &> 0, \end{aligned}$$

it follows that $\nabla^2 f$ is positive definite. Thus every stationary point is a minimum point. Since $f(c, \sigma)$ is defined on an open set such that $B(c, \sigma) > 0$ and as (c, σ) approaches the boundary or as $\sigma \rightarrow \infty$, we have $f(c, \sigma) \rightarrow +\infty$, the global minimum of $f(c, \sigma)$ is obtained at the unique stationary point.

We summarize the above discussion in the following theorem.

Theorem 3.2 *Given a positive definite covariance matrix $A \in \mathbb{R}^{m \times m}$, define $f(c, \sigma) := L(A, B(c, \sigma))$ where $B(c, \sigma)$ is a positive definite covariance matrix of compound symmetry in (3.16). Then the global minimum of $f(c, \sigma)$ over $\sigma > 0$ and $c \in (-1/(m-1), 1)$ is attained at*

$$\begin{cases} c = -\frac{t}{(m-1)\text{tr}(A^{-1}) + (m-2)t}, \\ \frac{m}{\sigma^2} = \text{tr}(A^{-1}) + ct, \end{cases} \quad (3.21)$$

where $t = \text{tr}(A^{-1}(ee^T - I))$.

3.3. AR(1)

We rewrite B in (1.5) as

$$B(c, \sigma) = \sigma^2 \begin{bmatrix} 1 & c & c^2 & \dots & c^{m-1} \\ c & 1 & c & \dots & c^{m-2} \\ c^2 & c & 1 & \dots & c^{m-3} \\ \vdots & \ddots & \ddots & \ddots & \vdots \\ c^{m-1} & c^{m-2} & \dots & c & 1 \end{bmatrix} = \sigma^2 \sum_{i=0}^{m-1} c^i T_i, \quad (3.22)$$

where $T_0 = I$ and T_i is a symmetric matrix with ones on the i th superdiagonal and subdiagonal and zeros elsewhere. It can be shown that the $k \times k$ leading principal

minor of $B(c, \sigma)$ is $\sigma^{2k}(1-c^2)^{k-1}$, $k = 2 : m$ (Horn and Johnson, 2013, Prob. 7.2.P12). Therefore, $B(c, \sigma)$ is a positive definite covariance matrix if and only if

$$-1 < c < 1. \quad (3.23)$$

The entropy loss function is now

$$f(c, \sigma) := \sigma^2 \sum_{i=0}^{m-1} c^i \text{tr}(A^{-1}T_i) + \log \det(A) - m \log \sigma^2 - (m-1) \log(1-c^2) - m.$$

We find that

$$\nabla f := \begin{bmatrix} \frac{\partial f}{\partial c} \\ \frac{\partial f}{\partial \sigma} \end{bmatrix} = \begin{bmatrix} \sigma^2 \sum_{i=1}^{m-1} i c^{i-1} \text{tr}(A^{-1}T_i) + \frac{2(m-1)c}{1-c^2} \\ 2\sigma \sum_{i=0}^{m-1} c^i \text{tr}(A^{-1}T_i) - 2m/\sigma \end{bmatrix}. \quad (3.24)$$

So the stationary points (c, σ) of $f(c, \sigma)$ satisfy

$$\begin{cases} \frac{m \sum_{i=1}^{m-1} i c^{i-1} \text{tr}(A^{-1}T_i)}{\sum_{i=0}^{m-1} c^i \text{tr}(A^{-1}T_i)} + \frac{2(m-1)c}{1-c^2} = 0 \\ \frac{m}{\sigma^2} = \sum_{i=0}^{m-1} c^i \text{tr}(A^{-1}T_i). \end{cases} \quad (3.25)$$

Since $\sum_{i=0}^{m-1} c^i \text{tr}(A^{-1}T_i) = \text{tr}(A^{-1}B)/\sigma^2 > 0$ and $1-c^2 \neq 0$ for $c \in (-1, 1)$, by rearranging the first equality in (3.25) we have

$$h(c) := m \sum_{i=1}^{m-1} i c^{i-1} \text{tr}(A^{-1}T_i) - m \sum_{i=1}^{m-1} i c^{i+1} \text{tr}(A^{-1}T_i) + 2(m-1) \sum_{i=0}^{m-1} c^{i+1} \text{tr}(A^{-1}T_i) = 0.$$

Since $h(c)$ is continuous in $[-1, 1]$, $h(-1) = -2(m-1)\text{tr}(A^{-1}B(-1, 1)) < 0$ and $h(1) = 2(m-1)e^T A^{-1}e > 0$, where the first inequality is from the positive semidefiniteness of $B(-1, 1)$ as in (3.22), there exists at least one root of $h(c)$ in $(-1, 1)$. Numerical experiments show that in some cases there exists more than one solution to $h(c) = 0$. We then can only expect to find a local minimum in general.

We summarize the discussion above in the following theorem.

Theorem 3.3 *Given a positive definite covariance matrix $A \in \mathbb{R}^{m \times m}$, define $f(c, \sigma) := L(A, B(c, \sigma))$ where $B(c, \sigma)$ is a positive definite covariance matrix of the AR(1) model as in (3.22). Then the local minima of $f(c, \sigma)$ are attained at the points (c, σ) satisfying (3.25).*

4. Toeplitz problems

Now we consider the problem for banded Toeplitz matrices, for which

$$B = \sigma^2 \begin{bmatrix} 1 & c_1 & \cdots & c_p & \cdots & 0 \\ c_1 & 1 & c_1 & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & \ddots & c_p \\ c_p & \ddots & \ddots & \ddots & \ddots & \vdots \\ \vdots & \ddots & \ddots & \ddots & 1 & c_1 \\ 0 & \cdots & c_p & \cdots & c_1 & 1 \end{bmatrix}. \quad (4.26)$$

Define $q(t) = 1 + 2 \sum_{k=1}^p c_k \cos(kt)$. Then B is positive-definite if and only if $q(t) \geq 0$, $q(t) \not\equiv 0$, for all $t \in \mathbb{R}$ (Parter, 1962, Remark II).

Now let $x_0 = \sigma^2$ and $x_i = \sigma^2 c_i$, $i = 1 : p$. The matrix B in (4.26) can be rewritten as

$$B(x) = x_0 I + \sum_{i=1}^p x_i T_i,$$

where $x = [x_0, x_1, \dots, x_p]^T \in \mathbb{R}^{p+1}$ and T_i is a symmetric matrix with the i th superdiagonal and subdiagonal elements equal to 1 and zeros elsewhere.

We define $\Omega \subset \mathbb{R}^{p+1}$ by

$$\Omega := \left\{ x \in \mathbb{R}^{p+1} : B(x) = x_0 I + \sum_{i=1}^p x_i T_i \text{ is positive definite} \right\}, \quad (4.27)$$

and $f(x) : \mathbb{R}^{p+1} \mapsto \mathbb{R}$,

$$f(x) := L(A, B(x)) = \text{tr}(A^{-1}B(x)) - \log(\det(A^{-1}B(x))) - m. \quad (4.28)$$

Since Ω is isomorphic to the set of all positive definite matrices of structure (4.26), our problem of minimizing $f(B)$ over positive definite B of structure (4.26) is equivalent to

$$\min \quad f(x) \text{ in (4.28)} \quad (4.29a)$$

$$\text{subject to } x \in \Omega \text{ in (4.27)}. \quad (4.29b)$$

We note that this problem was mentioned in a recent work by Ning et al. (2012), but no solution method was developed there.

Since $f(B) := L(A, B)$ is a strictly convex function of B (see section 2) and $B(x) = x_0 I + \sum_{i=1}^p x_i T_i$ is an affine map of x , by the fact that composition with an affine mapping preserves convexity (Boyd and Vandenberghe, 2004, Sec. 3.2.2), $f(x) := f(B(x))$ is strictly convex in x . On the other hand, the set of all positive definite Toeplitz matrices is a convex set and so is Ω . Therefore, (4.29) is a convex optimization problem and so has a unique minimizer.

We now explore further the properties of the objective function. For notational simplicity we define $T_0 = I$. From (2.8) and (2.9) and that $\nabla_{x_i} B = T_i$, by applying the chain rule, we have the gradient of f

$$\nabla_{x_i} f = \text{tr}(T_i(A^{-1} - B^{-1})), \quad i = 0 : p, \quad (4.30)$$

and the Hessian $H = [h_{ij}] \in \mathbb{R}^{(p+1) \times (p+1)}$ of f

$$h_{ij} = \nabla_{x_i x_j}^2 f = \text{tr}(T_i B^{-1} T_j B^{-1}), \quad i, j = 0 : p. \quad (4.31)$$

Since $f(x)$ is strictly convex in Ω and in general Newton's method works very well for strictly convex objective functions, we apply Newton's method with backtracking line search to problem (4.29). In the implementation, we choose an initial point $x^{(0)}$ such that $B(x^{(0)})$ is positive definite. To ensure that the iterates remain in Ω , in the backtracking line search to choose the step size t , we first multiply the initial guess $t = 1$ by a constant $\beta \in (0, 1)$ until $B(x + t\Delta x_{nt})$ is positive definite and then continue backtracking until a sufficient decrease condition is satisfied. We outline the method in Algorithm 4.1.

Algorithm 4.1 (Newton's method for solving problem (4.29)) *Given a starting point $x \in \Omega$ in (4.27) and tolerance ϵ , repeat:*

- 1 Compute the Newton step and decrement:
 evaluate the gradient g (4.30) and Hessian H (4.31) at x ;
 $\Delta x_{nt} := -H^{-1}g$; $\lambda^2 := g^T H g$.
- 2 Stopping criterion:
 quit if $\lambda^2/2 \leq \epsilon$.
- 3 Backtracking line search: given parameters $\alpha \in (0, 0.5)$ and $\beta \in (0, 1)$,
 $t := 1$;
 while $x + t\Delta x_{nt} \notin \Omega$ in (4.27),
 $t := \beta t$;
 while $f(x + t\Delta x_{nt}) > f(x) + \alpha t g^T \Delta x_{nt}$,
 $t := \beta t$.
- 4 Update: $x = x + t\Delta x_{nt}$.

The classical analysis of Newton's method can also be used here to get a complexity bound. Assume that $f(x)$ has the minimum value p^* and the absolute convergence tolerance of the problem is set to $\epsilon = 10^{-10}$. Then a bound on the number of Newton iterations required is $375(f(x^{(0)} - p^*)) + 6$ (Boyd and Vandenberghe, 2004, (9.57)). Note that this is a pessimistic bound and in our numerical experiments in the next section, for matrices of size $m = 100$ and 200 the number of iterations needed is at most 17. For more details on the analysis of Newton's method, we refer the reader to section 9.6 in Boyd and Vandenberghe (2004).

5. Numerical Experiments

In this section, we illustrate numerically how the techniques discussed above can be used in regularizing the underlying covariance structure. We first carry out simulation studies and then apply our techniques to real data analysis. All computations were performed with MATLAB 2012b. The root-finding problem in section 3 is solved with MATLAB `fzero` and the Newton method in section 4 is coded by authors.

5.1. Simulation studies

Recall that our idea is, given a covariance matrix and a class of possible candidate covariance structures, to find for each structure a covariance matrix that minimizes the entropy loss function. The structure of the minimizer that has the smallest entropy loss function value among the class is considered to be the most likely underlying covariance structure for the given covariance matrix. To examine the idea, our simulation experiments were carried out as follows. Let m be the dimension of the covariance matrices we test. We first generate an $m \times n$ data matrix R with columns randomly drawn from the multivariate normal distribution $\mathcal{N}(\mu, \Sigma)$ with a common mean vector $\mu = \sigma^2 e \in \mathbb{R}^m$ (recall that e is the vector of ones) and a common covariance matrix Σ . We then compute the sample covariance matrix A with the generated data R : $A = n^{-1} \sum_{i=1}^n (r_i - \bar{r})(r_i - \bar{r})^T$, where r_i is the i th column of R and $\bar{r} = n^{-1} \sum_{i=1}^n r_i$ is the sample mean. We test with the true covariance matrix Σ of various dimensions m , being either unstructured or having structures as discussed in the previous sections, where for each structure we consider several different values for σ^2 and c . The sample size is chosen as $n = 1000$. We summarize the experimental results in Tables 1–3, which are for the experiments with covariance matrix size $m = 100$, and Tables 4–6, which are for $m = 200$. We choose $c \in \{0.2, 0.5, 0.75\}$ and $\sigma^2 \in \{2, 4, 8\}$ for Σ having MA(1), CS, and AR(1) structures. For Σ having a general Toeplitz matrix we use the above σ^2 but randomly assign the correlation coefficients. In Tables 1–6 each row stands for one experiment and for each experiment we report the results averaged over 100 repeated simulations. The first column gives the true underlying covariance structure and the second column presents the discrepancy between the true covariance matrix Σ and the sample covariance matrix A under the measure of entropy loss function. The rest of the columns report the results from the computed matrix B with different structures. Note that we do not include a row for tridiagonal Σ with $c = 0.75$ because there does not exist such a positive definite covariance matrix in this case. The notation and abbreviation for the results reported in the tables are summarized

- Σ : true covariance matrix.
- A : sample covariance matrix.
- B : the computed covariance matrix that has a certain structure and minimizes the entropy loss function $L(A, B)$ in (1.1).

Table 1: Simulation results with $m = 100$; $\sigma^2 = 2$.

$c = 0.20$		B								
Σ	$L_{\Sigma,A}$	MA(1)		CS		AR(1)		Toep		L_{diff}
		$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	
MA(1)	5.22	5.40	0.55	9.74	4.88	5.58	0.73	5.29	0.64	0.11
CS	5.22	7.70	2.83	5.41	0.55	7.70	2.83	5.30	0.64	0.11
AR(1)	5.23	5.56	0.70	9.20	4.34	5.42	0.55	5.31	0.64	0.11
$c = 0.5$		B								
Σ	$L_{\Sigma,A}$	MA(1)		CS		AR(1)		Toep		L_{diff}
		$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	
MA(1)	5.23	5.42	0.55	290.45	285.33	186.02	180.93	5.37	0.70	0.16
CS	5.23	9.03	4.16	5.43	0.55	9.03	4.16	5.32	0.64	0.11
AR(1)	5.24	10.25	5.37	26.85	21.96	5.43	0.55	5.32	0.64	0.11
$c = 0.75$		B								
Σ	$L_{\Sigma,A}$	MA(1)		CS		AR(1)		Toep		L_{diff}
		$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	
CS	5.24	10.13	5.25	5.43	0.55	10.13	5.25	5.32	0.64	0.11
AR(1)	5.23	23.86	19.00	47.99	43.09	5.43	0.55	5.32	0.64	0.11
c not assigned		B								
Σ	$L_{\Sigma,A}$	MA(1)		CS		AR(1)		Toep		L_{diff}
		$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	
UnStr	5.23	255.24	250.32	255.27	250.36	255.22	250.30	241.07	236.39	
Toep	5.23	21.26	16.35	21.35	16.45	21.27	16.36	5.32	0.64	

- $L_{\Sigma,A}$, $L_{A,B}$ and $L_{\Sigma,B}$: the entropy loss function $L(\Sigma, A)$, $L(A, B)$ and $L(\Sigma, B)$, respectively.
- $L_{\text{diff}} = L(B_1, B_2)$: the entropy loss function between two best estimators for each structure.

Table 2: Simulation results with $m = 100$; $\sigma^2 = 4$.

$c = 0.20$		B								
Σ	$L_{\Sigma,A}$	MA(1)		CS		AR(1)		Toep		L_{diff}
		$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	
MA(1)	5.22	5.41	0.54	9.73	4.87	5.59	0.72	5.30	0.63	0.11
CS	5.23	7.70	2.83	5.41	0.55	7.70	2.83	5.30	0.63	0.11
AR(1)	5.25	5.58	0.70	9.22	4.33	5.43	0.55	5.32	0.63	0.11
$c = 0.5$		B								
Σ	$L_{\Sigma,A}$	MA(1)		CS		AR(1)		Toep		L_{diff}
		$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	
MA(1)	5.23	5.42	0.56	290.09	285.30	185.71	180.90	5.37	0.70	0.16
CS	5.24	9.05	4.16	5.43	0.55	9.05	4.16	5.32	0.64	0.11
AR(1)	5.24	10.26	5.37	26.82	21.95	5.43	0.55	5.32	0.63	0.11
$c = 0.75$		B								
Σ	$L_{\Sigma,A}$	MA(1)		CS		AR(1)		Toep		L_{diff}
		$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	
CS	5.23	10.12	5.24	5.42	0.55	10.12	5.24	5.32	0.63	0.11
AR(1)	5.22	23.82	19.00	47.87	43.08	5.41	0.56	5.30	0.64	0.11
c not assigned		B								
Σ	$L_{\Sigma,A}$	MA(1)		CS		AR(1)		Toep		L_{diff}
		$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	
UnStr	5.23	272.78	267.82	272.68	267.72	272.78	267.82	256.90	252.14	
Toep	5.24	23.05	18.17	22.48	17.60	23.03	18.15	5.32	0.63	

Table 3: Simulation results with $m = 100$; $\sigma^2 = 8$.

$c = 0.20$		B								
Σ	$L_{\Sigma,A}$	MA(1)		CS		AR(1)		Toep		L_{diff}
		$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	
MA(1)	5.22	5.41	0.55	9.74	4.88	5.59	0.73	5.30	0.64	0.11
CS	5.22	7.69	2.83	5.41	0.55	7.69	2.83	5.30	0.63	0.11
AR(1)	5.23	5.57	0.70	9.19	4.33	5.42	0.55	5.31	0.63	0.11
$c = 0.5$		B								
Σ	$L_{\Sigma,A}$	MA(1)		CS		AR(1)		Toep		L_{diff}
		$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	
MA(1)	5.23	5.42	0.54	290.14	285.29	185.75	180.90	5.36	0.68	0.16
CS	5.23	9.04	4.16	5.42	0.55	9.04	4.16	5.31	0.64	0.11
AR(1)	5.23	10.24	5.37	26.81	21.96	5.42	0.55	5.31	0.64	0.11
$c = 0.75$		B								
Σ	$L_{\Sigma,A}$	MA(1)		CS		AR(1)		Toep		L_{diff}
		$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	
CS	5.24	10.12	5.25	5.43	0.55	10.12	5.25	5.32	0.64	0.11
AR(1)	5.22	23.86	19.00	47.94	43.08	5.41	0.55	5.30	0.64	0.11
c not assigned		B								
Σ	$L_{\Sigma,A}$	MA(1)		CS		AR(1)		Toep		L_{diff}
		$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	
UnStr	5.22	266.86	261.89	266.91	261.94	266.86	261.89	252.49	247.76	
Toep	5.23	18.63	13.76	18.92	14.04	18.64	13.77	5.31	0.64	

Table 4: Simulation results with $m = 200$; $\sigma^2 = 2$.

$c = 0.20$		B								
Σ	$L_{\Sigma,A}$	MA(1)		CS		AR(1)		Toep		L_{diff}
		$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	
MA(1)	21.61	23.29	4.67	32.08	13.46	23.65	5.04	23.04	4.82	0.25
CS	21.58	26.19	7.63	23.25	4.69	26.19	7.63	23.00	4.84	0.25
AR(1)	21.61	23.59	4.99	30.98	12.39	23.29	4.69	23.05	4.83	0.25
$c = 0.5$		B								
Σ	$L_{\Sigma,A}$	MA(1)		CS		AR(1)		Toep		L_{diff}
		$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	
MA(1)	21.59	23.29	4.68	728.71	709.60	500.37	481.32	23.41	5.19	0.60
CS	21.61	27.60	8.97	23.30	4.67	27.60	8.97	23.05	4.82	0.25
AR(1)	21.66	33.09	14.43	67.06	48.44	23.36	4.73	23.11	4.87	0.25
$c = 0.75$		B								
Σ	$L_{\Sigma,A}$	MA(1)		CS		AR(1)		Toep		L_{diff}
		$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	
CS	21.61	28.69	10.06	23.30	4.67	28.69	10.06	23.05	4.82	0.25
AR(1)	21.63	60.18	41.60	110.43	91.85	23.32	4.72	23.07	4.87	0.25
c not assigned		B								
Σ	$L_{\Sigma,A}$	MA(1)		CS		AR(1)		Toep		L_{diff}
		$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	
UnStr	21.64	161.47	142.80	161.47	142.80	161.47	142.80	159.64	141.36	
Toep	21.62	24.76	6.12	24.76	6.12	24.76	6.12	23.06	4.82	

Table 5: Simulation results with $m = 200$; $\sigma^2 = 4$.

$c = 0.20$		B								
Σ	$L_{\Sigma,A}$	MA(1)		CS		AR(1)		Toep		L_{diff}
		$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	
MA(1)	21.62	23.33	4.71	32.12	13.49	23.69	5.07	23.08	4.86	0.25
CS	21.61	26.26	7.65	23.30	4.70	26.26	7.65	23.05	4.85	0.25
AR(1)	21.65	23.65	5.02	31.03	12.42	23.34	4.72	23.09	4.87	0.25
$c = 0.5$		B								
Σ	$L_{\Sigma,A}$	MA(1)		CS		AR(1)		Toep		L_{diff}
		$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	
MA(1)	21.59	23.27	4.67	728.26	709.52	499.97	481.25	23.39	5.17	0.59
CS	21.66	27.67	9.00	23.37	4.70	27.67	9.00	23.12	4.85	0.25
AR(1)	21.58	32.95	14.38	66.97	48.39	23.25	4.68	23.00	4.83	0.25
$c = 0.75$		B								
Σ	$L_{\Sigma,A}$	MA(1)		CS		AR(1)		Toep		L_{diff}
		$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	
CS	21.61	28.68	10.08	23.29	4.69	28.68	10.08	23.04	4.84	0.25
AR(1)	21.60	60.14	41.56	110.41	91.80	23.27	4.67	23.03	4.82	0.25
c not assigned		B								
Σ	$L_{\Sigma,A}$	MA(1)		CS		AR(1)		Toep		L_{diff}
		$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	
UnStr	21.66	161.66	143.00	161.63	142.97	161.66	143.00	159.81	141.55	
Toep	21.62	24.87	6.26	24.87	6.26	24.87	6.26	23.06	4.85	

Table 6: Simulation results with $m = 200$; $\sigma^2 = 8$.

$c = 0.20$		B								
Σ	$L_{\Sigma,A}$	MA(1)		CS		AR(1)		Toep		L_{diff}
		$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	
MA(1)	21.63	23.32	4.67	32.09	13.45	23.68	5.03	23.07	4.82	0.25
CS	21.61	26.25	7.63	23.31	4.68	26.25	7.63	23.07	4.83	0.24
AR(1)	21.64	23.64	5.02	31.03	12.42	23.33	4.71	23.09	4.86	0.24
$c = 0.5$		B								
Σ	$L_{\Sigma,A}$	MA(1)		CS		AR(1)		Toep		L_{diff}
		$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	
MA(1)	21.61	23.28	4.68	727.68	709.41	499.49	481.16	23.48	5.26	0.67
CS	21.63	27.62	8.99	23.32	4.69	27.62	8.99	23.07	4.84	0.25
AR(1)	21.62	32.98	14.39	67.00	48.41	23.28	4.69	23.03	4.84	0.25
$c = 0.75$		B								
Σ	$L_{\Sigma,A}$	MA(1)		CS		AR(1)		Toep		L_{diff}
		$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	
CS	21.59	28.66	10.05	23.27	4.66	28.66	10.05	23.02	4.81	0.25
AR(1)	21.64	60.21	41.59	110.52	91.85	23.32	4.70	23.08	4.85	0.25
c not assigned		B								
Σ	$L_{\Sigma,A}$	MA(1)		CS		AR(1)		Toep		L_{diff}
		$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	$L_{A,B}$	$L_{\Sigma,B}$	
UnStr	21.61	162.09	143.50	162.09	143.50	162.09	143.50	160.36	142.14	
Toep	21.58	24.79	6.22	24.78	6.20	24.79	6.22	23.02	4.83	

In Tables 1–6, we have the following observations.

1. When Σ is unstructured, for all the structures considered here the covariance matrices B we found have discrepancies $L_{\Sigma,B}$ and $L_{A,B}$ around 50 times as large as the discrepancy $L_{\Sigma,A}$ for $m = 100$ and around 8 times for $m = 200$. This indicates that no regularization is needed for A .
2. When Σ is structured, the matrix B we found having the minimum $L_{\Sigma,B}$ has the same structure as Σ . Moreover, we have $L_{\Sigma,B} < L_{\Sigma,A}$ for all cases as long as the estimated covariance matrix B has the same structure as the true covariance matrix Σ . In other words, the regularized estimator B that has the same structure as Σ is much better than the sample covariance matrix A in terms of the entropy loss function. This shows that regularization of the sample covariance matrix, or any other available estimators of the covariance matrix, is necessary not only for the convenient use of known structure but also for the accuracy of covariance estimation.
3. For Σ having one of the structures of MA(1), CS or AR(1), among different minimizers B , there are two structures clearly winning out in the sense of having smaller $L_{A,B}$: the one having the same structure as Σ and the Toeplitz, the latter always being the best. It is not surprising for the matrix B with Toeplitz structure to win out because all MA(1), CS, and AR(1) are indeed special Toeplitz structures. There is no doubt that minimizing among the larger feasible set will give the smaller minimum. We also point out that with the bandwidth p of the general Toeplitz ranging from 1 to $m - 1$, the smallest minimum is always obtained when $p = m - 1$.

To see how much difference there is between the two best regularized estimators - the one with the same structure as Σ and the one with general banded Toeplitz structure, we measure the discrepancy between these two with the entropy loss function: $L_{\text{diff}} = L(B_T, B_X)$, where B_T denotes the estimator B we find for the Toeplitz structure and B_X denotes the estimator B with the same structure as the underlying Σ . It turns out that L_{diff} is around 2.5% of both $L(A, B_T)$ and $L(A, B_X)$ for $m = 100$ and around 1.1% for $m = 200$. That means that the two best estimators B_X and B_T are very close in terms of the entropy loss function.

Note that the observation made from the discrepancy $L_{A,B}$ is extremely important because in practice the true covariance is usually unknown and so is $L_{\Sigma,B}$. Thus, the discrepancy $L_{A,B}$ can be used to identify the correct covariance structure as long as the class of the candidate structures is broad enough.

4. The observations above are common to all choices of the structure of Σ in the class we considered, the various values of c and σ^2 and the dimension m of the covariance matrix. Therefore, the findings are reliable in this sense.

5.2. Real data analysis

We also did experiments with some real data. Kenward’s (1987) cattle data was analyzed by various statistical methods for longitudinal data in the literature (Pourahmadi, 1999; Pan and Mackenzie, 2003). In the experiment, 60 cattle were assigned randomly to two treatment groups 1 and 2, each of which consists of 30 cattle, and received a certain treatment. The cattle in each group were weighed 11 times over a nineteen-week period. The weighing times for all cattle were the same, so that the cattle data is a balanced longitudinal data set. The aim of Kenward’s study was to investigate treatment effects on intestinal parasites of the cattle. Our experiments were carried out with the cattle data in a similar way as in section 5.1 and the results are reported in Table 7. We also show under the column “Time” the time (in seconds) used to find the optimal matrix B for each structure. L_{diff} shows the discrepancy between the two best estimators measured by the entropy loss function $L(B_1, B_2)$.

Table 7: Results of experiments on Kenward’s cattle data.

	MA(1)		CS		AR(1)		Toep		L_{diff}
	$L_{A,B}$	Time	$L_{A,B}$	Time	$L_{A,B}$	Time	$L_{A,B}$	Time	
Group 1	9.86	2.87e-03	8.55	2.74e-03	5.22	7.93e-03	4.75	4.45e-02	0.47
Group 2	8.05	1.81e-03	5.92	1.76e-03	3.15	4.59e-03	2.08	1.89e-02	1.07

Note that in this real data analysis the true covariance matrix Σ is unknown, so $L_{\Sigma,A}$ and $L_{\Sigma,B}$ are not available, where A is the sample covariance matrix. Instead, we use the discrepancy $L_{A,B}$ to identify the most likely covariance structure among the possible candidate structures, MA(1), CS, AR(1) and general Toeplitz.

From Table 7, it is clear that the underlying covariance structures are very likely to be Toeplitz for both groups, among the four possible candidate structures, as their discrepancy $L_{A,B}$ has smaller values than others. We can also claim, more specifically, that Group 1 tends to have an AR(1) covariance structure, due to the very small value of L_{diff} between the Toeplitz and AR(1). This agrees with the finding by Pourahmadi (1999) and Pan and Mackenzie (2003).

6. Discussion

We have proposed a method to regularize the underlying covariance structure with a given covariance matrix A and a class of candidate covariance structures, based on minimizing the entropy loss function between the given covariance matrix and the matrix that has a certain structure. Our simulation studies demonstrate the reliability of the proposed method. Our simulation experiments were carried out with the given matrix A being the sample covariance matrix. In principle, any available estimated covariance matrix using a statistical method can be chosen as the given matrix A . In general, the structure behind the matrix A is not obvious

due to noise in the matrix A , in particular, when the dimension m of the matrix A is large. Our aim is to regularize the matrix A , so as to filter the noise in A and to have a standard structure to characterize the covariance/correlation process of the data studied.

Our proposed method can overcome the difficulties that are met by alternative approaches and can produce a reliable estimator of the covariance matrix even if the dimension of the matrix is as large as 200. For example, it is not easy to directly calculate the maximum likelihood estimator of a covariance matrix that has a certain structure. In contrast, the proposed method does not require any distribution assumption of the data, and can provide a regularized covariance structure estimator as long as an estimator of the covariance matrix is given.

A restriction of our studies here is that the class of candidate covariance structures we have considered comprises just four possible structures: MA(1), CS, AR(1), and banded Toeplitz. In principle, the ideas and the proposed approach are applicable to any structured covariance matrices. We are currently studying some other covariance structures, including linearly structured covariance, factor analytic, Hankel structure, all of which are very useful in statistics. However, with more complicated covariance structures more challenging work is inevitably involved in finding the structured covariance matrix that minimizes the entropy loss function. Note that for the three structures of MA(1), CS and AR(1) considered in section 3, the dimension m of the covariance matrix does not affect the cost of computing the optimal covariance matrix B , because the optimization problem reduces to computing the zeros of a nonlinear function of a single variable. For more complicated covariance structures, the optimization problem is liable to have more variables and the number of variables may increase with the dimension m of the covariance matrix, making the solution of the optimization problem much more challenging.

Acknowledgement

This work was supported by the EPSRC grant EP/I01912X/1 and European Research Council Advanced Grant MATFUN (267526). We gratefully acknowledge very helpful and constructive comments and suggestions made by four anonymous referees, which leads to significant improvements to the paper.

References

- R. Borsdorf, N. J. Higham, and M. Raydan. Computing a nearest correlation matrix with factor structure. *SIAM J. Matrix Anal. Appl.*, 31(5):2603–2622, 2010.
- S. Boyd and L. Vandenberghe. *Convex optimization*. Cambridge University Press, Cambridge, UK, 2004. ISBN 978-0-521-83378-3.
- J. P. Burg, D. G. Luenberger, and D. L. Wenger. Estimation of structured covariance matrices. *Proceedings of the IEEE*, 70(9):963–974, 1982.

- D. K. Dey and C. Srinivasan. Estimation of a covariance matrix under Stein's loss. *The Annals of Statistics*, 13(4):1581–1591, 1985.
- I. S. Dhillon and J. A. Tropp. Matrix nearness problems with Bregman divergences. *SIAM J. Matrix Anal. Appl.*, 29(4):1120–1146, 2007.
- A. Ferrante, M. Pavon, and M. Zorzi. A maximum entropy enhancement for a family of high-resolution spectral estimators. *IEEE Trans. Automat. Control*, 57(2):318–329, 2012.
- N. J. Higham. *Accuracy and Stability of Numerical Algorithms*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, second edition, 2002. ISBN 0-89871-521-0.
- R. A. Horn and C. R. Johnson. *Matrix Analysis*. Cambridge University Press, Cambridge, UK, second edition, 2013. ISBN 978-0-521-83940-2.
- W. James and C. Stein. Estimation with quadratic loss. In J. Neyman, editor, *Proceedings of the Fourth Berkeley Symposium*, volume 1 of *Mathematical Statistics and Probability*, pages 361–379. University of California Press, 1961. The Statistical Laboratory, University of California, June 30–July 30, 1960.
- M. G. Kenward. A method for comparing profiles of repeated measurements. *Appl. Statist.*, pages 296–308, 1987.
- F. Lin and M. R. Jovanović. Least-squares approximation of structured covariances. *IEEE Trans. Automat. Control*, 54(7):1643–1648, 2009.
- E. M. Maboudou-Tchao and V. Agbotu. Monitoring the covariance matrix with fewer observations than variables. *Computational Statistics & Data Analysis*, 64: 99–112, 2013.
- J. R. Magnus and H. Neudecker. *Matrix Differential Calculus with Applications in Statistics and Econometrics*. Wiley, Chichester, UK, revised edition, 1999. ISBN 0-471-98633-X.
- L. Ning, X. Jiang, and T. Georgiou. Geometric methods for structured covariance estimation. In *American Control Conference*, pages 1877–1882. IEEE, 2012.
- J. Pan and K. Fang. *Growth Curve Models and Statistical Diagnostics*. Springer-Verlag, New York, 2002.
- J. Pan and G. Mackenzie. On modelling mean-covariance structures in longitudinal studies. *Biometrika*, 90(1):239–244, 2003.
- S. V. Parter. An observation on the numerical solution of difference equations and a theorem of Szegő. *Numerische Mathematik*, 4:293–295, 1962.

- F. Pascal, Y. Chitour, J. P. Ovarlez, P. Forster, and P. Larzabal. Covariance structure maximum-likelihood estimates in compound Gaussian noise: existence and algorithm analysis. *IEEE Trans. Signal Processing*, 56(1):34–48, 2008.
- M. Pourahmadi. Joint mean-covariance models with applications to longitudinal data: Unconstrained parameterisation. *Biometrika*, 86(3):677–690, 1999. ISSN 0006-3444.
- V. Vinciotti and H. Hashem. Robust methods for inferring sparse network structures. *Computational Statistics & Data Analysis*, 67:84–94, 2013.
- H. Ye and J. Pan. Modelling of covariance structures in generalised estimating equations for longitudinal data. *Biometrika*, 93(4):927–941, 2006. ISSN 0006-3444.