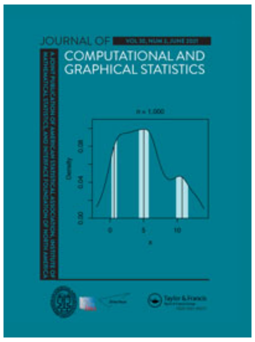

Journal of Computational and Graphical Statistics

ISSN: (Print) (Online) Journal homepage: https://www.tandfonline.com/loi/ucgs20

Shrinking the Covariance Matrix Using Convex Penalties on the Matrix-Log Transformation

Mengxi Yi & David E. Tyler

To cite this article: Mengxi Yi & David E. Tyler (2021) Shrinking the Covariance Matrix Using Convex Penalties on the Matrix-Log Transformation, Journal of Computational and Graphical Statistics, 30:2, 442-451, DOI: 10.1080/10618600.2020.1814788

To link to this article: https://doi.org/10.1080/10618600.2020.1814788

JOURNAL OF COMPUTATIONAL AND GRAPHICAL STATISTICS
2021, VOL. 30, NO. 2, 442–451
https://doi.org/10.1080/10618600.2020.1814788

Taylor & Francis
Taylor & Francis Group

Check for updates

Shrinking the Covariance Matrix Using Convex Penalties on the Matrix-Log Transformation

Mengxi Yi[a] and David E. Tyler[b]

[a]School of Statistics, University of International Business and Economics, Beijing, China; [b]Department of Statistics, Rutgers, The State University of New Jersey, Piscataway, NJ

ABSTRACT
For q-dimensional data, penalized versions of the sample covariance matrix are important when the sample size is small or modest relative to q. Since the negative log-likelihood under multivariate normal sampling is convex in Σ^{-1}, the inverse of the covariance matrix, it is common to consider additive penalties which are also convex in Σ^{-1}. More recently, Deng and Tsui and Yu et al. have proposed penalties which are strictly functions of the roots of Σ and are convex in $\log \Sigma$, but not in Σ^{-1}. The resulting penalized optimization problems, though, are neither convex in $\log \Sigma$ nor in Σ^{-1}. In this article, however, we show these penalized optimization problems to be geodesically convex in Σ. This allows us to establish the existence and uniqueness of the corresponding penalized covariance matrices. More generally, we show that geodesic convexity in Σ is equivalent to convexity in $\log \Sigma$ for penalties which are functions of the roots of Σ. In addition, when using such penalties, the resulting penalized optimization problem reduces to a q-dimensional convex optimization problem on the logs of the roots of Σ, which can then be readily solved via Newton's algorithm. Supplementary materials for this article are available online.

ARTICLE HISTORY
Received December 2018
Revised August 2020

KEYWORDS
Geodesic convexity; Newton–Raphson algorithm; Penalized likelihood

1. Introduction and Motivation

For a q dimensional sample $x_1, \ldots, x_n$, the sample covariance matrix $S_n = n^{-1}\sum_{i=1}^{n}(x_i - \bar{x})(x_i - \bar{x})^{\mathrm{T}}$ is not well-conditioned and can be highly variable when q is of the same order as n. In such cases, one may wish to consider a penalized version of the sample covariance matrix. Since the loss function obtained from the negative log-likelihood under multivariate normal sampling

$$l(\Sigma; S_n) = \mathrm{tr}(\Sigma^{-1}S_n) + \log\det\Sigma, \qquad (1)$$

is convex in Σ^{-1}, it is natural to consider additive penalties which are also convex in Σ^{-1} such as the graphical lasso penalty $\sum_{i\neq j}|\{\Sigma^{-1}\}_{ij}|$ (Yuan and Lin 2007; Friedman, Hastie, and Tibshirani 2008). Minimizing the penalized loss function

$$L(\Sigma; S_n, \eta) = l(\Sigma; S_n) + \eta\Pi(\Sigma), \qquad (2)$$

over the set of symmetric positive definite matrices $\Sigma > 0$, with $\Pi(\Sigma)$ being a nonnegative penalty function and $\eta \geq 0$ being a tuning parameter, is then a convex optimization problem.

More recently, Deng and Tsui (2013) considered the penalty $\Pi_R(\Sigma) \equiv \|\log\Sigma\|_F^2$, where the norm refers to the Frobenius norm. This penalty is strictly convex in $\log\Sigma$ but not in Σ^{-1}. By letting $A = \log\Sigma$, they observe that, when using this penalty, (2) can be expressed in terms of a penalized loss function over the set of symmetric matrices of order q, namely

$$L_o(A; S_n, \eta) = \mathrm{tr}(e^{-A}S_n) + \mathrm{tr}A + \eta\|A\|_F^2, \qquad (3)$$

with the penalty $\|A\|_F^2 = \mathrm{tr}(A^2)$ being strictly convex in A. As noted in Section 2, the function $\mathrm{tr}(e^{-A}S_n)$ is not in general a convex function of A, and consequently minimizing (3) over A does not correspond to a convex optimization problem. Hence, there is no assurance as to the existence and uniqueness of a minimum to (3).

One of our objectives in this article is to argue that rather than using the concept of convexity in $\log\Sigma$ in problem (3), a more appropriate setting is based on the notion of geodesic convexity, or g-convexity for short. The function $\|\log\Sigma\|_F$ has been well studied within Riemannian geometry and corresponds to the Riemannian or geodesic distance between Σ and the identity matrix (Moakher 2005; Bhatia 2009), and is known to be strictly g-convex in Σ, as well as in Σ^{-1}. In general, unlike convexity in Σ or in Σ^{-1}, g-convexity in Σ and g-convexity in Σ^{-1} are equivalent, and so the term g-convexity is used to refer to either or both. For $S_n \neq 0$, the loss function (1) is also strictly g-convex, and consequently the penalized loss function (2) is strictly g-convex, when choosing $\Pi = \Pi_R$. Moreover, (2), with $\Pi = \Pi_R$, can be shown to be g-coercive, which implies it has a unique critical point, with this unique critical point corresponding to its global minimum; see Lemmas 2.2 and 2.3.

The concept of g-convexity can be mathematically challenging, and in practice, it can be difficult to prove that a given function is g-convex. A further contribution of this article is to show that for an orthogonally invariant penalties, that is, $\Pi(\Sigma) = \Pi(P\Sigma P^{\mathrm{T}})$ for any orthogonal P and hence strictly

CONTACT David E. Tyler dtyler@stat.rutgers.edu Department of Statistics, Rutgers, The State University of New Jersey, Hill Center, Busch Campus, Piscataway, NJ 08854.
Supplementary materials for this article are available online. Please go to www.tandfonline.com/r/JCGS.

functions of the eigenvalues of Σ, (strict) g-convexity in Σ and (strict) convexity in $\log \Sigma$ are equivalent. Furthermore, it is shown that (strict) g-convexity for such function reduces to the simpler task of establishing (strict) convexity when viewed as function on the logs of the eigenvalues; see Theorem 3.1. For example, if we express $\Pi_R(\Sigma) = \sum_{j=1}^{q} a_j^2$, where $a_j = \log \lambda_j$ with $\lambda_1 \geq \cdots \geq \lambda_q > 0$ being the eigenvalues of Σ, then it is strictly convex as a function of $a \in \mathbb{R}^q$, and hence $\Pi_R(\Sigma)$ is strictly convex in $\log \Sigma$ and strictly g-convex in Σ.

Deng and Tsui (2013) also proposed an iterative quadratic programming algorithm over the class of symmetric matrices A of order q for finding the minimum of (3). We show in Theorem 5.1, though, that the solution to this problem has the same eigenvectors as S_n. This leads to a simpler algorithm based on finding the minimum of a strictly convex univariate function for each eigenvalue, namely $g(a; d) = de^{-a} + a + \eta a^2$, with d corresponding to an eigenvalue of S_n and a being the corresponding log eigenvalue of Σ. The solution to this univariate convex optimization problem can be readily obtained via a Newton–Raphson algorithm.

As recently noted by Yu, Wang, and Zhu (2017), the penalty $\|A\|_F^2$ shrinks the sample covariance matrix toward the identity matrix. They propose using the alternative penalty $\|A - \widehat{m}I_q\|_F^2$, with $\widehat{m}$ being an estimate of the mean of the log of the eigenvalues of Σ, that is, of $m(A) = \mathrm{tr}A/q$. Since $\widehat{m}$ is first determined from the data, this does not correspond to a pure penalty function for Σ. Rather than using a preliminary estimate of m, we propose replacing $\widehat{m}$ with $m(A)$. This approach yields an estimate of $m(A)$ consistent with the penalized estimate of Σ, that is, $m(\widehat{A}) = \mathrm{tr}\widehat{A}/q$. The resulting penalized objective function (2), when using the penalty $\|A - m(A)I_q\|_F^2 = \sum_{j=1}^{q}(a_j - \bar{a})^2$ is shown, within Section 4, to also be strictly g-convex and g-coercive. Consequently, the global minimum of (2) corresponds to the unique critical point. The solution to this optimization problem reduces to finding the minimum of a strictly convex function on $\mathbb{R}^q$.

Summarizing, this article is organized as follows. In Section 2, the concept of geodesic convexity is briefly reviewed, and results on the existence and uniqueness of penalized sample covariance matrices based on g-convex penalty functions in general are presented. Results on the relationship between convexity in $\log \Sigma$ and g-convexity in Σ are given in Section 3. In Section 4, convexity results for (2) are given when applying the penalty to the shape matrix $\Sigma/(\det \Sigma)^{1/q}$ rather than to Σ itself, with $\|A - m(A)I_q\|_F^2$ being a special case of such a shape penalty. Algorithms for computing the penalized sample covariance matrices, based on orthogonally invariant g-convex penalties are given in Section 5. We emphasize that this article treats g-convex penalties in general, with applications to Π_R treated as a special case. The results of a simulation study discussed in Section 6, together with an example given in Section 7, demonstrate the advantages of penalizing shape. Proofs and some technical details are given in Appendix A.

2. Geodesic Convexity

The notion of geodesic distance between multivariate normal distributions, or equivalently the geodesic distance between their covariance matrices, has been a topic of interest at least as early as Skovgaard (1984). However, the realization the multivariate normal negative log-likelihood $l(\Sigma; S_n)$ is g-convex, and strictly g-convex when $S_n \neq 0$, which follows as a special case of Theorem 1 in Zhang, Wiesel, and Greco (2013), is relatively recent.

The set of symmetric positive definite matrices of order q can be viewed as a Riemannian manifold with the geodesic path from $\Sigma_0 > 0$ to $\Sigma_1 > 0$ being given by $\Sigma_t = \Sigma_0^{1/2}\{\Sigma_0^{-1/2}\Sigma_1\Sigma_0^{-1/2}\}^t\Sigma_0^{1/2}$ for $0 \leq t \leq 1$, see Bhatia (2009) or Wiesel and Zhang (2015) for more details. An alternative representation for this path is given by $\Sigma_t = Be^{t\Delta}B^{\mathrm{T}}$, where $\Sigma_0 = BB^{\mathrm{T}}$ and $\Sigma_1 = Be^{\Delta}B^{\mathrm{T}}$ with Δ being a diagonal matrix of order q. A function $f(\Sigma)$ is said to be g-convex if and only if $f(\Sigma_t) \leq (1-t)f(\Sigma_0) + tf(\Sigma_1)$ for $0 < t < 1$, and it is strictly g-convex if strict inequality holds for $\Sigma_0 \neq \Sigma_1$. Analogous to convexity in $\log \Sigma$, for which convexity in $\log \Sigma$ holds if an only if convexity in $\log \Sigma^{-1}$ holds, g-convexity in Σ holds if and only if g-convexity in Σ^{-1} holds.

As with convexity, any local minimum of a g-convex function is a global minimum, and when differentiable any critical point is a global minimum, with the set of all minima being g-convex. In addition, if a minimum exists, then the minimum is unique when the function is strictly g-convex. Finally, the sum of two g-convex functions is g-convex, and the sum is strictly g-convex if either of the two g-convex summands is strictly g-convex. Consequently, the following lemma holds.

Lemma 2.1. If $\Pi(\Sigma)$ is g-convex and $S_n \neq 0$, then $L(\Sigma; S_n, \eta)$ is strictly g-convex on $\Sigma > 0$, and the set of all local minima $\mathcal{A}_\eta$ is either empty or contains a single element. That is, if there exists a minimizer $\widehat{\Sigma}_\eta > 0$ to $L(\Sigma; S_n, \eta)$, then it is unique.

The existence of a minimum for a g-convex function requires some additional conditions, with a necessary and sufficient condition being that it be geodesic coercive (Dümbgen and Tyler 2016). A g-convex function $F(\Sigma)$ is said to be g-coercive if and only if $F(\Sigma) \to \infty$ as $\|\log \Sigma\|_F \to \infty$. For $S_n > 0$, $l(\Sigma; S_n)$ is g-coercive and so, since $\Pi(\Sigma)$ is bounded below, $L(\Sigma; S_n, \eta)$ is g-coercive and hence has a unique minimizer. Moreover, as shown in Appendix A, the solution is a continuous function of $\eta \geq 0$. This is summarized in the following lemma.

Lemma 2.2. Under the conditions of Lemma 2.1, if $S_n > 0$, then there exists a unique critical point $\widehat{\Sigma}_\eta > 0$ to $L(\Sigma; S_n, \eta)$, with $\widehat{\Sigma}_\eta$ being its unique minimizer. Furthermore, $\widehat{\Sigma}_\eta$ is a continuous function of $\eta \geq 0$.

For singular S_n, some conditions on the penalty function are needed since it is possible for $\mathrm{tr}(\Sigma^{-1}S_n)$ to be bounded as $\log\det\Sigma \to -\infty$, and hence $l(\Sigma; S_n)$ is not g-coercive in this case. A sufficient condition for $L(\Sigma; S_n, \eta)$ to be g-coercive when S_n is singular is that $\Pi(\Sigma)$ be g-coercive and $\eta > 0$. This condition, however, is too strong, and does not hold for the scale invariant or shape penalties discussed in Section 4. Some weaker conditions are given in the following lemma, with these conditions holding when $\Pi(\Sigma)$ is g-coercive. Note that under each of the three conditions below, $\Pi(\Sigma) \to \infty$.

Lemma 2.3. Under the conditions of Lemma 2.1, if

(i) $\Pi(\Sigma) \to \infty$ whenever $|\log\det\Sigma|$ is bounded above and $\|\log\Sigma\|_F \to \infty$,
(ii) $(\log\det\Sigma)/\Pi(\Sigma) \to 0$ whenever $\log\det\Sigma \to -\infty$ but with λ_1 bounded away from 0, and
(iii) $\{\log(\lambda_1/\lambda_q)\}/\Pi(\Sigma)$ is bounded above whenever $\lambda_1 \to 0$ but with λ_1/λ_q bounded away from 1,

then the conclusions stated in Lemma 2.2 hold when $S_n \neq 0$ is singular and $\eta > 0$.

3. Geodesic Convexity and Convexity in Log

In the following, we show that for orthogonally invariant functions, g-convexity in Σ is equivalent to convexity in $\log\Sigma$. We say a function F on the set of symmetric positive definite matrices is orthogonally invariant if and only if $F(\Sigma) = F(H\Sigma H^{\mathrm{T}})$ for any orthogonal matrix H of order q. It is straightforward to show that such functions can then be expressed in terms of a symmetric function of its eigenvalues $\lambda_1 \geq \cdots \geq \lambda_q > 0$.

Lemma 3.1. The function $F(\Sigma)$ is orthogonally invariant if and only if for some symmetric, that is, permutation invariant, function $f : \mathbb{R}^q \to \mathbb{R}, F(\Sigma) \equiv f(a_1, \ldots, a_q)$ where $a_j = \log\lambda_j$, $j = 1, \ldots, q$.

Theorem 3.1. For an orthogonally invariant function $F(\Sigma)$, the following three conditions are equivalent:

(i) $F(\Sigma)$ is (strictly) g-convex.
(ii) $F(\Sigma)$ is (strictly) convex in $\log\Sigma$.
(iii) The corresponding function f, as defined in Lemma 3.1, is (strictly) convex.

A clarifying point regarding Theorem 3.1 may be helpful. It should be noted, for example, that the corresponding function on $\mathbb{R}^q$ for the log concave function $F(\Sigma) = \log\lambda_q$ is not $f(a_1, \ldots, a_q) = a_q$ which is linear, hence convex, but not symmetric. Rather, its corresponding function is $f(a_1, \ldots, a_q) = \min\{a_1, \ldots, a_q\}$ which is symmetric but concave.

Outside of orthogonally invariant functions, g-convexity and convexity in log do not necessarily coincide. For example, as previously noted, $l(\Sigma; S_n)$ is strictly g-convex, but not necessarily convex in log. In particular, although $\log\det\Sigma = \mathrm{tr}(A)$ is linear and hence convex in A, whether or not the term $\mathrm{tr}(\Sigma^{-1}S_n) = \mathrm{tr}(e^{-A}S_n)$ is convex in A depends on the value of S_n. For example, when $S_n = I$, the convexity of $\mathrm{tr}(e^{-A})$ follows from Theorem 3.1 since $\sum_{j=1}^q e^{-a_j}$ is convex. As far as we are aware, general conditions on S_n needed for $\mathrm{tr}(e^{-A}S_n)$ to be convex have not been formally addressed in the literature. An example of S_n for which $\mathrm{tr}(e^{-A}S_n)$ is not convex in A is given in Appendix A. On the other hand, an example of a function which is convex in log but not g-convex is also presented in Appendix A.

We now apply these results to the penalty studied by Deng and Tsui (2013), that is, $\Pi_R(\Sigma) = \|\log\Sigma\|_F^2$. This penalty is orthogonally invariant and can be expressed as $\Pi_R(\Sigma) = \sum_{j=1}^q a_i^2$, which is symmetric and strictly convex as a function of $a \in \mathbb{R}^q$. Hence, by Theorem 3.1, Π_R is strictly g-convex, and so Lemma 2.2 holds. Furthermore, Lemma 2.3 also holds since $\log\det\Sigma/\Pi_R(\Sigma) = \{\sum_{j=1}^q a_i\}/\{\sum_{j=1}^q a_i^2\} \to 0$ as $\sum_{j=1}^q a_i \to -\infty$.

The geodesic convexity of $\|\log\Sigma\|_F^2$ has been previously established using more involved proofs, see Bhatia (2009) for comparison. The importance of Theorem 3.1 is that it completely characterizes g-convexity for functions which are strictly functions of the eigenvalues of Σ, as well as provides a simple condition for verifying g-convexity. For example, it readily follows that the geodesic distance between Σ and identity, that is, $\|\log\Sigma\|_F$, is g-convex. The condition number penalty λ_1/λ_q and the penalty $\mathrm{tr}(\Sigma) + \mathrm{tr}(\Sigma^{-1})$, among others considered by Wiesel (2012) and Dümbgen and Tyler (2016), are also seen to be g-convex.

4. Penalizing the Shape Matrix

Any penalty on $\Sigma > 0$ can also be applied to its shape matrix $V(\Sigma) = \Sigma/\det(\Sigma)^{1/q}$. Here $\det V(\Sigma) = 1$, with the orbits of $V(\Sigma)$ form equivalence classes over $\Sigma > 0$; see Paindaveine (2008). This then generates the new penalty $\Pi_s(\Sigma) \equiv \Pi(V(\Sigma))$. If the original penalty is minimized, for example, at $\Sigma = I_q$, then the new penalty is minimized at any $\Sigma \propto I_q$, that is, when Σ is proportional to I_q. Applying the penalty studied by Deng and Tsui (2013) to the shape matrix yields

$$\Pi_{R,s}(\Sigma) \equiv \Pi_R\{V(\Sigma)\} = \|\log\Sigma - q^{-1}\{\log\det\Sigma\}I_q\|_F^2 = \|A - mI_q\|_F^2,$$

where $m = q^{-1}\{\log\det\Sigma\} = \mathrm{tr}A/q$. Since $\Pi_{R,s}$ is orthogonally invariant, with $\Pi_{R,s}(\Sigma) = \sum_{j=1}^q (a_i - \overline{a})^2$ being convex, it follows from Theorem 3.1 that $\Pi_{R,s}$ is convex in log as well as g-convex, although the convexity is not strict in this case. Thus, for nonsingular S_n, Lemma 2.2 on existence and uniqueness applies when using $\Pi_{R,s}$ as the penalty term. Also, as shown in Appendix A, the additional conditions given in Lemma 2.3 needed to assure existence and uniqueness when S_n is singular also hold when using this penalty.

More generally, applying any g-convex penalty or penalty which is convex in $\log\Sigma$ to the shape matrix of Σ, yields, respectively, a new g-convex penalty or penalty convex in $\log\Sigma$. The following theorem applies to any such penalties and does not presume Π is orthogonally invariant.

Theorem 4.1.

(i) If $\Pi(\Sigma)$ is g-convex, then $\Pi_s(\Sigma)$ is also g-convex.
(ii) If $\Pi(\Sigma)$ is convex in $\log\Sigma$, then $\Pi_s(\Sigma)$ is also convex in $\log\Sigma$.

Thus, for g-convex $\Pi(\Sigma)$, Lemma 2.2 on existence and uniqueness for the case when S_n is nonsingular still applies when the penalty term Π is replaced by Π_s. As another example, if we apply the Kullback–Leibler divergence from the identity to the shape matrix of Σ, one obtains the penalty $\Pi_s(\Sigma) = \mathrm{tr}\{V(\Sigma)^{-1}\} = q\overline{\lambda}_g/\overline{\lambda}_h$, where $\overline{\lambda}_g$ and $\overline{\lambda}_h$ are, respectively, the geometric mean and the harmonic mean of the eigenvalues of Σ. This ratio represents a measure of eccentricity for Σ, and is minimized at any $\Sigma \propto I_q$. By the previous theorem, this new penalty is also g-convex, and hence Lemma 2.2 applies. It can be verified that Lemma 2.3 also applies for this case.

5. Optimizing the Penalized Loss Function

As noted in the introduction, Deng and Tsui (2013) proposed a quadratic iterative programming algorithm over $A = \log \Sigma$. The algorithm is derived by a repeated application of the Volterra integral equation for e^{tA} to obtain a second-order expansion. Although they state in their introduction that some other previously proposed "methods have retained the use of the eigenvectors of S_n in estimating Σ or Σ^{-1}," it is not clear if they recognize that the minimum $\widehat{A}_\eta$ to (3), and hence $\widehat{\Sigma}_\eta = e^{\widehat{A}_\eta}$, also retains the same eigenvectors as S_n. As shown by the following theorem and corollary, this is true for any orthogonally invariant penalty.

Theorem 5.1. Suppose $\Pi(\Sigma)$ is orthogonally invariant and $\Sigma > 0$. Using the spectral value decomposition, express $S_n = P_n D_n P_n^{\mathrm{T}}$ with P_n being an orthogonal matrix of order q and $D_n = \mathrm{diag}\{d_1, \ldots, d_q\}$. Then

$$L(\Sigma; S_n, \eta) \geq L(P_n \Lambda P_n^{\mathrm{T}}; S_n, \eta),$$

where $\Lambda = \mathrm{diag}\{\lambda_1, \ldots, \lambda_q\}$ with $\lambda_1 \geq \cdots \geq \lambda_q > 0$ being the ordered eigenvalues of Σ.

Corollary 5.1. If $\Pi(\Sigma)$ is orthogonally invariant, and there exists a global minimum to $L(\Sigma; S_n, \eta)$ over $\Sigma > 0$, then it has a global minimizer of the form $\widehat{\Sigma}_\eta = P_n \widehat{\Lambda}_\eta P_n^{\mathrm{T}}$, where $\widehat{\Lambda}_\eta$ is a diagonal matrix which minimizes $L(P_n \Lambda P_n^{\mathrm{T}}; S_n, \eta)$ over all $\Lambda \in \mathcal{D}_+ = \{\mathrm{diag}\{\lambda_1, \ldots, \lambda_q\} \mid \lambda_1 \geq \cdots \geq \lambda_q > 0\}$. Consequently, if $L(\Sigma; S_n, \eta)$ has a unique global minimizer over $\Sigma > 0$, then it is given by $\widehat{\Sigma}_\eta$, as defined above, with $\widehat{\Lambda}_\eta$ being the unique global minimizer of $L(P_n \Lambda P_n^{\mathrm{T}}; S_n, \eta)$ over $\Lambda \in \mathcal{D}_+$.

In the above corollary, if we express $\widehat{\Lambda}_\eta = \mathrm{diag}\{e^{\widehat{a}_{\eta,1}}, \ldots, e^{\widehat{a}_{\eta,q}}\}$, then $\widehat{a}_\eta$ is the minimizer of

$$L_q(a; d; \eta) = \sum_{j=1}^{q} (d_j e^{-a_j} + a_j) + \eta \pi(a_1, \ldots, a_q), \tag{4}$$

over $a \in \mathbb{R}^q$ with $a_1 \geq \cdots \geq a_q$. Here π is the function on $\mathbb{R}^q$ corresponding to the function Π as defined in Lemma 3.1, and $d_1 \geq \cdots \geq d_q$ are the eigenvalues of S_n. The function $L_q(a; d; \eta)$ is strictly convex whenever $\pi(a)$ is convex in $a \in \mathbb{R}^q$, which by Theorem 3.1 holds whenever $\Pi(\Sigma)$ is g-convex, or equivalently convex in $\log \Sigma$. Thus, for orthogonally invariant g-convex penalties, rather than using a quadratic iterative programming algorithm over the set of symmetric matrices of order q, as proposed in Deng and Tsui (2013), or using a g-convex optimization method over $\Sigma > 0$, the minimization problem (2) reduces to the simpler and numerically well studied problem of minimizing a strictly convex function on $\mathbb{R}^q$. Since the function π is symmetric and $\sum_{j=1}^{q} d_j e^{-a_{(j)}} \leq \sum_{j=1}^{q} d_j e^{-a_j}$, where $a_{(1)} \geq \cdots \geq a_{(q)}$ are the ordered values of $a_1, \ldots, a_q$, it follows that minimizing $L_q(a; d; \eta)$ over $a \in \mathbb{R}^q$ yields $\widehat{a}_{\eta,1} \geq \cdots \geq \widehat{a}_{\eta,q}$, with the inequalities being strict whenever the corresponding inequalities $d_1 \geq \cdots \geq d_q$ are strict.

Summarizing, for a general g-convex and orthogonally invariant penalty $\Pi(\Sigma)$, the algorithm for computing $\widehat{\Sigma}_\eta$, that is, the minimum of (2) over $\Sigma > 0$, can be summarized as follows:

Step 1: Compute the spectral value decomposition of the sample covariance matrix $S_n = P_n D_n P_n^{\mathrm{T}}$.

Step 2: Obtain $\widehat{a}_\eta \in \mathbb{R}^q$ by minimizing the strictly convex function $L_q(a; d; \eta)$, given by (4), over $a \in \mathbb{R}^q$.

Step 3: Finally, compute $\widehat{\Sigma}_\eta = P_n \widehat{\Lambda}_\eta P_n^{\mathrm{T}}$, where $\widehat{\Lambda}_\eta = \mathrm{diag}\{e^{\widehat{a}_{\eta,1}}, \ldots, e^{\widehat{a}_{\eta,q}}\}$.

Step 2 can be done via Newton's method or some other convex optimization method.

5.1. Newton's Method When using Riemannian Penalties

The penalties proposed by Deng and Tsui (2013) and Yu, Wang, and Zhu (2017) have the form $\Pi_R(\Sigma; c) = \|A - cI_q\|_F^2$ with c not dependent on Σ. Here, $\pi_R(a; c) = \sum_{j=1}^{q} (a_j - c)^2$ and so $L_q(a; d; \eta) = \sum_{j=1}^{q} \{d_j e^{-a_j} + a_j + \eta(a_j - c)^2\}$. Minimizing $L_q(a; d; \eta)$ reduces to solving q univariate strictly convex optimization problems, namely $\min\{d_j e^{-a_j} + a_j + \eta(a_j - c)^2\}$ for $j = 1, \ldots, q$. Newton's algorithm for this problem is given by

$$a_{j,k+1} = a_{j,k} + \frac{d_j e^{-a_{j,k}} - 2\eta(a_{j,k} - c) - 1}{d_j e^{-a_{j,k}} + 2\eta}. \tag{5}$$

If S_n has rank $r < q$, then for $j > r$ one obtains the closed form $\widehat{a}_{\eta,j} = c - 1/\eta$.

For the shape version of this penalty, that is, for $\Pi_{R,s}$, we have $\pi_{R,s}(a) = \sum_{j=1}^{q} (a_j - \bar{a})^2$ and hence $L_q(a; d; \eta) = \sum_{j=1}^{q} \{d_j e^{-a_j} + a_j + \eta(a_j - \bar{a})^2\}$. For $\mathrm{rank}(S_n) = r$, Newton's algorithm is given by

$$a_{j,k+1} = a_{j,k} + \delta_{j,k}\, g_{j,k} + \beta_k\, \delta_{j,k} \sum_{i=1}^{q} \delta_{i,k}\, g_{i,k}, \tag{6}$$

where $g_{j,k} = d_j e^{-a_{j,k}} - 1 - 2\eta(a_{j,k} - \bar{a}_k)$, $\delta_{i,k} = 1/(d_i e^{-a_{i,k}} + 2\eta)$, $\beta_k = 2\eta/\{r - 2\eta \sum_{i=1}^{r} \delta_{i,k}\}$, and $\bar{a}_k = \sum_{i=1}^{q} a_{i,k}$. If $r < q$, then $\widehat{a}_{\eta,r+1} = \cdots = \widehat{a}_{\eta,q}$ and (6) can be implemented as an $r + 1$ dimensional, rather than a q dimensional, algorithm after replacing $\sum_{i=1}^{q} \delta_{i,k}\, g_{i,k}$ with $\sum_{i=1}^{r} \delta_{i,k}\, g_{i,k} + (q - r)\delta_{r+1,k}\, g_{r+1,k}$. Here, $\bar{a}_k = (\sum_{i=1}^{r} a_{i,k} + (q - r)a_{r+1,k})/q$, and $\delta_{r+1,k}\, g_{r+1,k} = -\{(2\eta)^{-1} + (a_{r+1,k} - \bar{a}_k)\}$.

The derivations of (5) and (6) are given in Appendix A. Also, as shown in Appendix A, the gradient of $L_q(a; d; \eta)$ is concave in these cases. This last property allows us to prove the following convergence result.

Theorem 5.2. Given any initial value $(a_{1,0}, \ldots, a_{q,0})$, the algorithms (5) and (6) are decreasing. That is, for $j = 1, \ldots, q$, $a_{j,k} \geq a_{j,k+1} \geq \widehat{a}_{j,\eta}$ for $k \geq 1$, where $(\widehat{a}_{1,\eta}, \ldots, \widehat{a}_{q,\eta})$ denotes the unique global minimizer of (4) when using $\pi_R(a; c)$ and $\pi_{R,s}(a)$, respectively. Furthermore, $a_{j,k} \to \widehat{a}_{j,\eta}$ quadratically as $k \to \infty$.

6. Simulation Study

In this section, we conduct a simulation study to compare the performance of the following five covariance estimators:

- S: the sample covariance matrix,
- LogF: the penalized covariance matrix proposed by Deng and Tsui (2013) with penalty $\|A\|_F^2$, where $A = \log \Sigma$,
- sLogF: our proposed shape penalized covariance matrix based on $\|A - \{\mathrm{tr}(A)/q\} I_q\|_F^2$,
- mLogF: the adjusted penalized covariance matrix proposed by Yu, Wang, and Zhu (2017) based on $\|A - \widehat{m} I_q\|_F^2$, with $\widehat{m}$ being an estimate of $m(A) = tr(A)/q$, and

Table 1. For $q = 100$, simulation results for the performance of the covariance estimators under Models 1 and 2 using four different measures.

		Model 1					Model 2				
	Method	Fnorm	L_1	op-norm	Δ_1	η	Fnorm	L_1	op-norm	Δ_1	η
$n = 50$	sLogF	1.51	0.51∗	0.33	0.25	5.80	2.99	1.00∗	0.71	0.59	9.89
		(0.02)	(0.03)	(0.00)	(0.01)	(1.22)	(0.03)	(0.07)	(0.01)	(0.03)	(3.08)
	dLogF	1.37∗	0.86	0.32∗	0.06∗	1.44	2.82∗	1.59	0.67∗	0.29∗	2.53
		(0.01)	(0.06)	(0.01)	(0.03)	(0.17)	(0.02)	(0.13)	(0.02)	(0.06)	(0.39)
	mLogF	2.09	0.68	0.45	0.35	2.69	4.83	1.53	1.04	0.84	3.29
		(0.02)	(0.01)	0.00	(0.01)	(0.12)	(0.03)	(0.02)	(0.01)	(0.02)	(0.14)
	LogF	2.91	2.60	0.88	0.76	0.30	4.27	3.90	1.14	0.74	0.59
		(0.05)	(0.14)	(0.05)	(0.05)	(0.00)	(0.07)	(0.19)	(0.06)	(0.08)	(0.03)
	S	2.86	3.18	1.05	0.94	NA	7.13	7.86	2.52	2.19	NA
		(0.08)	(0.25)	(0.09)	(0.10)	NA	(0.18)	(0.54)	(0.20)	(0.21)	NA
$n = 100$	sLogF	1.41	0.64∗	0.33	0.20	1.81	2.85	1.22∗	0.71	0.51	3.15
		(0.02)	(0.02)	(0.01)	(0.01)	(0.15)	(0.03)	(0.05)	(0.01)	(0.02)	(0.35)
	dLogF	1.23∗	0.88	0.30∗	0.03∗	0.77	2.63∗	1.65	0.65∗	0.24∗	1.37
		(0.01)	(0.04)	(0.01)	(0.02)	(0.05)	(0.02)	(0.07)	(0.02)	(0.04)	(0.11)
	mLogF	2.07	0.72	0.44	0.33	0.76	4.75	1.69	1.02	0.76	0.81
		(0.02)	(0.02)	(0.00)	(0.01)	(0.05)	(0.04)	(0.03)	(0.01)	(0.02)	(0.04)
	LogF	2.21	2.04	0.66	0.55	0.20	3.64	3.30	0.92	0.55	0.42
		(0.04)	(0.10)	(0.04)	(0.04)	(0.00)	(0.05)	(0.14)	(0.04)	(0.06)	(0.04)
	S	2.02	2.15	0.68	0.58	NA	5.04	5.31	1.61	1.32	NA
		(0.05)	(0.13)	(0.05)	(0.06)	NA	(0.11)	(0.31)	(0.12)	(0.14)	NA
$n = 150$	sLogF	1.30	0.69∗	0.32	0.16	1.01	2.71	1.35∗	0.69	0.45	1.77
		(0.02)	(0.02)	(0.01)	(0.01)	(0.06)	(0.03)	(0.05)	(0.01)	(0.02)	(0.14)
	dLogF	1.13∗	0.85	0.29∗	0.02∗	0.54	2.47∗	1.66	0.62∗	0.21∗	0.97
		(0.01)	(0.04)	(0.01)	(0.02)	(0.05)	(0.03)	(0.06)	(0.02)	(0.03)	(0.06)
	mLogF	1.23	0.77	0.32	0.11	0.73	2.59	1.54	0.68	0.34	1.20
		(0.02)	(0.03)	(0.01)	(0.01)	(0.04)	(0.03)	(0.05)	(0.01)	(0.03)	(0.08)
	LogF	1.96	1.68	0.55	0.45	0.20	3.24	2.82	0.77	0.42	0.39
		(0.03)	(0.07)	(0.02)	(0.03)	0.00	(0.04)	(0.12)	(0.03)	(0.05)	(0.03)
	S	1.65	1.72	0.53	0.43	NA	4.11	4.24	1.25	0.98	NA
		(0.04)	(0.10)	(0.04)	(0.04)	NA	(0.08)	(0.22)	(0.08)	(0.09)	NA

NOTE: Means and standard deviations (in parentheses) are based on 100 runs.

dLogF: an adjusted penalized covariance matrix based on $\|A - \log(\bar{d})I_q\|_F^2$, where $\bar{d} = \mathrm{tr}(S_n)/q$, that is, the average of the sample eigenvalues.

Comparisons of LogF and mLogF to other penalized covariance estimators are given in Deng and Tsui (2013) and Yu, Wang, and Zhu (2017).

As the tuning constant $\eta \to \infty$, the estimator LogF goes to the identity matrix and so one would anticipate its performance would be poor whenever $\bar{\lambda} = \mathrm{tr}\Sigma/q$ is far from one. This would be particularly problematic when heavy tuning is desirable, as would be the case whenever the roots of Σ are not well separated or in general when S_n is singular. As noted by Yu, Wang, and Zhu (2017), this weakness can be alleviated by using the estimator mLogF. Alternatively, the estimators sLogF or dLogF can be considered. As shown in Appendix A, the estimator sLogF goes to $\bar{d}I_q$ as $\eta \to \infty$. On the other hand, an adjusted estimator, that is, one using a penalty of the form $\|A - cI_q\|_F^2$, goes to $e^c I_q$ as $\eta \to \infty$. Consequently, the estimators mLogF and dLogF go to $e^{\widehat{m}}I_q$ and $\bar{d}I_q$, respectively, as $\eta \to \infty$.

The performance of the estimator mLogF depends on the definition of $\widehat{m}$. Yu, Wang, and Zhu (2017) observed that the simple choice $\widehat{m}_0 = m(\log S_n) = \mathrm{tr}(\log S_n)/q$ is known to underestimate $m(A)$. They proposed using a bias corrected estimator of the form $\widehat{m}_1 = m(\log S_n) + b_{n,q}$ when $q < n$, and a Bayesian estimator for $\widehat{m}_3$ when $q \geq n$; see Yu, Wang, and Zhu (2017) for details. We use their proposed choices of $\widehat{m}$ in our simulation study. When using $\widehat{m}_1$ the estimator mLogF shrinks the eigenvalues of S_n toward $e^{\widehat{m}_1} \propto (\det S_n)^{1/q} = \bar{d}_g$, the geometric mean of the eigenvalues of S_n. We surmise it would be better to shrink them toward the arithmetic mean since $\bar{d}$ is the minimum variance unbiased estimator of λ when random sampling from a spherical multivariate normal distribution with $\Sigma = \lambda I_q$. In particular, we anticipate our proposed estimators sLogF and dLogF, which both shrink the eigenvalues of S_n toward $\bar{d}$, will have a better performance in settings where it is desirable for the tuning parameter to be large. The results of our simulation study, reported in Tables 1 and 2, support this heuristic argument.

For the simulations, we consider $q = 100$ dimensional data arising as a random sample from a multivariate normal distribution with mean $\mu = 0$ and covariance matrix $\Sigma = \{\sigma_{ij}\}$. The four different covariance models used in the simulations are listed below. For each of these covariance models, the arithmetic mean $\bar{\lambda}$ and the geometric mean $\bar{\lambda}_g$ of their eigenvalues are reported. Note that the ratio $\gamma = \bar{\lambda}/\bar{\lambda}_g \geq 1$ represents a measure of the eccentricity of Σ, that is, a measure of how much Σ deviates from proportionality to I_q, with $\gamma = 1$ if and only if $\Sigma \propto I_q$. The models are listed in decreasing order of γ.

Model 1: An MA(2) model for which $\sigma_{ii} = 0.20, \sigma_{i,i-1} = \sigma_{i-1,i} = 0.10, \sigma_{i,i-2} = \sigma_{i-2,i} = 0.05$, and $\sigma_{ij} = 0$ otherwise. Here, $\bar{\lambda} = 0.200$ and $\bar{\lambda}_g = 0.145$.

Model 2: An AR(1) model for which $\sigma_{ij} = 0.5\rho^{|i-j|}$ and $\rho = 0.4$. Here, $\bar{\lambda} = 0.500$ and $\bar{\lambda}_g = 0.421$.

Model 3: $\Sigma^{-1} = \{\sigma^{ij}\}$ where $\sigma^{ii} = 1$ and $\sigma^{ij} = 0.6$ for $i \neq j$. Here, $\bar{\lambda} = 2.475$ and $\bar{\lambda}_g = 2.378$.

Model 4: $\Sigma = 5\, I_q$. Here, $\bar{\lambda} = \bar{\lambda}_g = 5.0$.

Table 2. For $q = 100$, simulation results for the performance of the covariance estimators under Models 3 and 4 using four different measures.

		Model 3					Model 4				
	Method	Fnorm	L_1	op-norm	Δ_1	η	Fnorm	L_1	op-norm	Δ_1	η
$n = 50$	sLogF	2.53*	2.50	2.41	0.06*	135.96	0.98	0.52*	0.15*	0.11*	141.54
		(0.08)	(0.03)	(0.05)	(0.05)	(75.34)	(0.61)	(0.45)	(0.11)	(0.12)	(72.26)
	dLogF	2.53*	2.49*	2.43	0.10	118.81	0.97*	0.62	0.20	0.18	137.10
		(0.09)	(0.04)	(0.05)	(0.11)	(81.39)	(0.68)	(0.61)	(0.18)	(0.19)	(78.25)
	mLogF	20.70	5.03	2.25	1.39	4.16	42.64	8.70	4.59	2.95	4.18
		(0.07)	(0.07)	(0.00)	(0.04)	(0.15)	(0.12)	(0.20)	(0.01)	(0.07)	(0.14)
	LogF	14.29	4.20	1.57*	0.81	6.85	37.54	8.67	4.10	2.27	4.58
		(0.04)	(0.09)	(0.01)	(0.05)	(0.54)	(0.09)	(0.26)	(0.00)	(0.09)	(0.22)
	S	35.17	37.82	11.26	11.26	NA	71.09	76.42	22.73	22.73	NA
		(0.71)	(2.05)	(0.60)	(0.60)	NA	(1.43)	(4.36)	(1.25)	(1.25)	NA
$n = 100$	sLogF	2.49*	2.48*	2.41	0.04*	115.62	0.70*	0.37*	0.11*	0.08*	140.36
		(0.04)	(0.02)	(0.04)	(0.04)	(78.64)	(0.39)	(0.30)	(0.07)	(0.08)	(73.24)
	dLogF	2.50	2.47	2.42	0.08	78.61	0.73	0.44	0.13	0.12	130.94
		(0.04)	(0.02)	(0.04)	(0.07)	(74.42)	(0.43)	(0.41)	(0.10)	(0.11)	(76.93)
	mLogF	19.77	6.29	2.45	0.85	0.84	39.89	11.87	4.89	1.77	0.89
		(0.19)	(0.19)	(0.00)	(0.08)	(0.05)	(0.36)	(0.37)	(0.00)	(0.16)	(0.05)
	LogF	13.53	4.71	1.66*	0.54	2.85	34.26	11.11	4.27	1.28	1.57
		(0.05)	(0.13)	(0.01)	(0.05)	(0.16)	(0.15)	(0.31)	(0.01)	(0.10)	(0.06)
	S	24.81	25.66	7.01	7.01	NA	50.13	51.77	14.12	14.12	NA
		(0.39)	(1.21)	(0.33)	(0.33)	NA	(0.80)	(2.39)	(0.68)	(0.68)	NA
$n = 150$	sLogF	2.48*	2.47	2.41	0.04*	88.05	0.56*	0.31*	0.09*	0.07*	136.97
		(0.02)	(0.02)	(0.03)	(0.04)	(74.94)	(0.31)	(0.26)	(0.06)	(0.07)	(74.63)
	dLogF	2.48*	2.45*	2.41	0.08	61.03	0.57	0.32	0.10	0.09	133.76
		(0.02)	(0.02)	(0.04)	(0.06)	(65.00)	(0.33)	(0.27)	(0.07)	(0.08)	(74.68)
	mLogF	2.63	2.53	2.33	0.06	59.34	0.72	0.35	0.12	0.11	131.43
		(0.11)	(0.04)	(0.04)	(0.04)	(62.31)	(0.43)	(0.28)	(0.08)	(0.09)	(75.22)
	LogF	12.69	5.30	1.75*	0.24	1.54	29.36	15.35	4.45	0.61	0.53
		(0.07)	(0.15)	(0.01)	(0.05)	(0.07)	(0.36)	(0.68)	(0.02)	(0.26)	(0.05)
	S	20.29	20.66	5.40	5.40	NA	40.99	41.73	10.88	10.88	NA
		(0.28)	(0.93)	(0.24)	(0.24)	NA	(0.58)	(1.86)	(0.49)	(0.49)	NA

NOTE: Means and standard deviations (in parentheses) are based on 100 runs.

To evaluate the performance of the different estimators under the various models, we use the same four measures of the discrepancy between the estimated and true covariance matrix as in Yu, Wang, and Zhu (2017).

Fnorm: $||\widehat{\Sigma} - \Sigma||_F = \sqrt{\sum_{i,j}(\widehat{\sigma}_{ij} - \sigma_{ij})^2}$.

L_1: $||\widehat{\Sigma} - \Sigma||_1 = \max_j \sum_i |\widehat{\sigma}_{ij} - \sigma_{ij}|$.

op-norm: $||\widehat{\Sigma} - \Sigma||_{op} = \max_j |\tilde{\sigma}_j|$, where $\tilde{\sigma}_j$'s are the singular values of $\widehat{\Sigma} - \Sigma$.

Δ_1: $|\widehat{\lambda}_1 - \lambda_1|$, the absolute difference between the largest eigenvalues of $\widehat{\Sigma}$ and Σ,

Since all five estimators of the covariance matrix are orthogonally equivariant, the results of the simulations for the Fnorm, op-norm, and Δ_1 would be the same if the model Σ were replaced by $P\Sigma P^{\mathrm{T}}$ for any orthogonal P. These results depend only on the eigenvalue of Σ. The results under L_1, though, would differ since, unlike the other measures, it is not orthogonally invariant. Also, since all the estimators except LogF are scale equivariate, if the model Σ were multiplied by a positive constant c, then the simulation result for these estimators would simply be multiplied by c. Hence, the relative comparisons across these estimators would be the same. The relative performance of LogF, though, would depend on the value of c.

We follow the simulation protocol used by both Deng and Tsui (2013) and Yu, Wang, and Zhu (2017). For each covariance model, $2n$ data points are generated, with the first n observations serving as a training set and last n observations serving as a validation set. The tuning parameter for any particular method is then selected via the holdout method of cross-validation. That is, we select the value of η which minimizes the non-penalized loss $l(\widehat{\Sigma}_{\eta,o}; S_{n,1})$ defined by (1), where $\widehat{\Sigma}_{\eta,o}$ is the penalized covariance estimate based on the training set, and $S_{n,1}$ is the sample covariance matrix of the validation set. Since the true mean $\mu = 0$ and interest lies in the performance of the estimators of Σ, the non-centered sample covariance matrices $S_n = n^{-1}\sum_{i=1}^n x_i x_i^{\mathrm{T}}$ are used in the simulations. We consider three values for the sample size $n = 50, 100, 150$, with the dimension being $q = 100$ in each case. The simulations are repeated 100 times and the means and standard deviations (in parenthesis) over the 100 trials for each of the discrepancy measures are reported in Tables 1 and 2. The means and standard deviations over the 100 trials of the value of the selected tuning parameter η are also reported.

In Tables 1 and 2, the estimator which performed the best for a given sample size and given discrepancy measure is noted with an asterisk (*). For all four models, with one exception, either sLogF or dLogF is the best performing estimator, depending on the particular discrepancy measure used. The notable exception is under model 3, for which LogF performs best under the operator norm. Overall, the performance of our two proposed estimators sLogF and dLogF are comparable, with sLogF being slightly better when the model is close to proportionality to the identity and dLogF performing better otherwise. The performance of mLogF is also similar to these two estimators when $n = 150$, but performs considerably worse when $n = 100$ or $n = 50$. The sample covariance matrix performs uniformly worst. Further simulations are given in a supplement to this article.

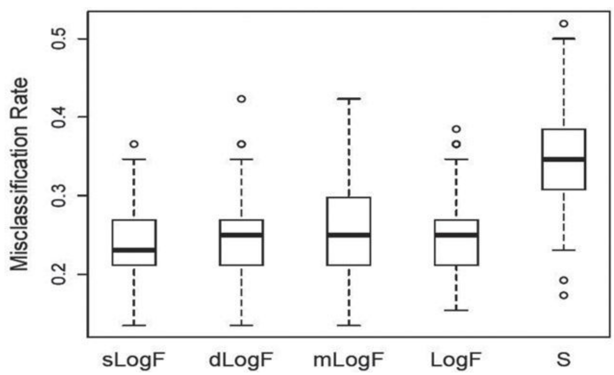

Figure 1. Boxplot of misclassification errors over 100 runs.

Table 3. Means and standard deviations of the misclassification error over 100 runs.

	sLogF	dLogF	mLogF	LogF	S
Mean	0.239	0.247	0.259	0.253	0.349
SD	0.048	0.051	0.062	0.049	0.071

7. An Example: Sonar Data

Datasets for which the sample size is modest relative to the dimension typically arise, for example, in functional data, image data, financial data, and signal processing. As an example, we consider the sonar dataset obtained from University of California Irvine Machine Learning Repository, which was developed and first analyzed by Gorman and Sejnowski (1988). This dataset consists of 208 multivariate observations of dimension $q = 60$. For each observation, the 60 variables correspond to the average energy over a particular frequency band obtained by bouncing sonar signals off of an object under various conditions, with 111 observations labeled M (metal cylinder) and the other 97 observations labeled R (rock).

Our goal here is to study the relative performance of covariance estimators when used within linear discriminant analysis (LDA) to classify an observation as either M or R. The dataset is randomly partitioned into a training set of size 78 for estimating the covariance matrix, a validation set of size 78 for selecting the tuning parameter and a test set of size 52 for computing the misclassification error. The covariance estimators being compared are those considered in the simulation study in Section 6. The role of the sample covariance matrix used in defining all of the estimators is replaced with the pooled sample covariance matrix of the two groups M and R.

The above procedure is independently repeated for 100 times. A boxplot of the misclassification errors over these 100 trials are displayed in Figure 1, and the mean and standard deviation of the misclassification errors are showed in Table 3. Finally, Table 4 displays the frequency over the 100 trials that a given estimator (row) has a lower classification rate than another estimator (column). For example, sLogF has less misclassification errors than mLogF in 46 of the 100 runs, and more misclassification errors in 22 of the runs, with the two estimators having the same misclassification rate in the other 32 runs. Among the estimators of the covariance matrix considered here, our proposed sLogF estimator performs best.

Table 4. Out of 100 runs, the number of times the estimator indicated in the row has a lower misclassification rate than the estimator indicated in the column.

	sLogF	dLogF	mLogF	LogF	S
sLogF	0	38	46	59	89
dLogF	28	0	31	59	93
mLogF	22	18	0	47	84
LogF	38	33	49	0	90
S	5	6	12	8	0

8. Some Final Remarks

The focus of this article is on the uniqueness and computation for g-convex penalized sample covariance matrices. As is the case with the papers by Deng and Tsui (2013) and Yu, Wang, and Zhu (2017), the current article does not address the statistical properties of the penalized covariance estimators outside of the simulations. For a fixed tuning constant η, fixed dimension q and smooth penalty Π, though, it readily follows from the theory of M-estimation that under suitable regularity conditions the estimator $\widehat{\Sigma}_\eta$ is consistent for its functional version Σ_η, that is, for the unique minimizer of $\mathcal{L}(\Sigma; \Sigma_o, \eta) \equiv \mathrm{tr}(\Sigma^{-1}\Sigma_o) + \log\det\Sigma + \eta\Pi(\Sigma)$ over $\Sigma > 0$, where Σ_o denotes the true population covariance matrix. In addition, $\sqrt{n}(\widehat{\Sigma}_\eta - \Sigma_\eta)$ converges in distribution to a multivariate normal distribution. The consistency and asymptotic normality of $\widehat{\Sigma}_\eta$ for Σ_o itself holds if $\eta = o(1/\sqrt{n})$. When the value of η is chosen via a data driven method, consistency and asymptotic normality still hold provided the estimated tuning parameter $\widehat{\eta} = o_p(1/\sqrt{n})$.

Penalization methods are of primary importance when the sample size n is small relative to the dimension q. In such cases, fixed q asymptotics may not provide reasonable approximations for a given n and q. Consequently, there is a growing interest in understanding the asymptotic behavior of a covariance estimator, say $\widehat{\Sigma}$, as both n and q increase. In particular, for regularized covariance estimators, rates of convergence for the error term $\|\widehat{\Sigma} - \Sigma_o\|$, in terms of n and q, have been studied by Bickel and Levina (2008), Lam and Fan (2009), and Cai, Zhang, and Zhou (2010), among others. A study of the rates of convergence for g-convex penalized sample covariance estimators is a worthwhile topic which we hope to address, or that others may wish to address, in future research. Rates of convergence have primarily been studied when the norm for the error term is taken to be either the Frobenius norm or the operator norm. We are unaware of any results for when the norm is taken to be the Riemannian norm, which would be a natural choice for the g-convex penalized covariance estimators. Finally, we note that in a follow-up paper (Tyler and Yi 2020) the asymptotic behavior of the eigenvalues of $\widehat{\Sigma}_\eta$ as $q \to \infty$ and $q/n \to c \in (0, 1)$ have been studied for a special class of g-convex penalty terms.

Appendix A: Proofs and Some Technical Details

A.1. Counterexamples to the Equivalency of g-Convexity and Convexity in Log

Lemma 1.14 in Wiesel and Zhang (2015) states that $x^{\mathrm{T}}\Sigma^{-1}x$ is a strictly g-convex function of Σ, which implies that $\mathrm{tr}(\Sigma^{-1}S_n)$ is g-convex for $S_n \neq 0$. It is difficult to show analytically whether or not $\mathrm{tr}(\Sigma^{-1}S_n)$ is convex in $\log\Sigma$ for a given S_n, and almost all randomly generated examples tend to suggest that it is true. After extensive trials, though, the following counterexample was found which shows that $\{\Sigma^{-1}\}_{11}$ is not a convex function of $\log \Sigma$, and consequently $\mathrm{tr}(\Sigma^{-1}S_n)$ cannot be convex in $\log \Sigma$ in general. For $q = 2$, let $A = \log \Sigma$ and choose

$$A_0 = \begin{bmatrix} 0 & -1 \\ -1 & 300 \end{bmatrix} \quad \text{and} \quad A_1 = -\begin{bmatrix} 0 & 0.01 \\ 0.01 & 0.01 \end{bmatrix}.$$

This gives $\{e^{-(0.5A_1+0.5A_2)}\}_{11} = 1.001690296 > 1.001688939 = 0.5\{e^{-A_0}\}_{11} + 0.5\{e^{-A_1}\}_{11}$, and so $\{e^{-A}\}_{11} = \{\Sigma^{-1}\}_{11}$ is not convex in A.

On the other hand, a function may be convex in $\log \Sigma$ but not g-convex in Σ. For example, the matrix L_1 norm on the elements of $\log \Sigma$, that is, $H(\Sigma) = \max_{1\le k\le q} \sum_{j=1}^{q} |\{\log \Sigma\}_{jk}|$, is convex in $\log \Sigma$. The following counter-example, though, shows that it is not g-convex. For $q = 3$, choose

$$\Sigma_0 = \begin{bmatrix} 1.00 & 0.30 & 0.09 \\ 0.30 & 1.00 & 0.30 \\ 0.09 & 0.30 & 1.00 \end{bmatrix} \quad \text{and} \quad \Sigma_1 = \begin{bmatrix} 1.00 & 0.90 & 0.81 \\ 0.90 & 1.00 & 0.90 \\ 0.81 & 0.90 & 1.00 \end{bmatrix}.$$

This gives $H(\Sigma_{0.5}) = 2.289438 > 2.284073 = 0.5\, H(\Sigma_0) + 0.5\, H(\Sigma_1)$, and so $H(\Sigma)$ is not g-convex.

A.2. Proof of Lemma 2.2

As already noted, the first part of the lemma follows from Lemma 2.1 and the g-coercivity of $L(\Sigma; S_n, \eta)$. To prove continuity, we first state the following general lemma.

Lemma A.1. Let $\mathcal{D}$ be a closed subset of $\mathbb{R}^p$. Suppose the real-valued functions $f(x)$ and $g(x)$ are continuous on $\mathcal{D}$, with $g(x) > 0$. Furthermore, suppose $h(x; \eta) = f(x) + \eta g(x)$ has a unique minimum in $\mathcal{D}$ for any $0 \le \eta_o \le \eta \le \eta_1$. If the set $\{x \in \mathcal{D} \mid h(x; \eta_o) \le c\}$ is compact for any $c \ge \inf\{h(x; \eta_o) \mid x \in \mathcal{D}\}$, then the function $x(\eta) = \arg\inf\{h(x; \eta) \mid x \in \mathcal{D}\}$ is continuous for $\eta_o \le \eta < \eta_1$.

To prove this lemma, first note that $h(x; \eta)$ is increasing in η, and so the set $\{x(\eta) \mid \eta_o \le \eta < \eta_1\}$ is contained in the compact set $\{x \mid h(x; \eta_o) \le h(x(\eta_1); \eta_1)\}$. So, if $\eta_k \to \eta$, then $x(\eta_k)$ has a convergent subsequence, say $x(\eta_{k'}) \to \tilde{x}$. By definition, $h(x(\eta_{k'}); \eta_{k'}) \le h(x(\eta); \eta_{k'})$. By continuity, the left-hand side converges to $h(\tilde{x}; \eta)$ and the right-hand side converges to $h(x(\eta); \eta)$. By uniqueness, this implies $\tilde{x} = x(\eta)$. Hence, $x(\eta_k) \to x(\eta)$, which establishes Lemma A.1.

This lemma then applies to $h(\Sigma; \eta) = L(\Sigma; S_n, \eta)$, for which $f(\Sigma) = l(\Sigma; S_n)$ and $g(\Sigma) = \Pi(\Sigma)$. By g-convexity, both f and g are continuous. Also, the level sets of $h(\Sigma; \eta)$ are compact since by g-coercivity $h(\Sigma; \eta) \to \infty$ as $\| \log \Sigma \|_F \to \infty$. Hence, $\widehat{\Sigma}_\eta$ is continuous for $\eta > 0$.

A.3. Proof of Lemma 2.3

The lemma follows if $L(\Sigma; S_n, \eta)$ is g-coercive. Consider any sequence in Σ such that $\| \log \Sigma \|_F \to \infty$. Divide the proof into the following three cases: a) $\log\det \Sigma \to \infty$, b) $|\log\det \Sigma|$ is bounded above, and c) $\log\det \Sigma \to -\infty$. For case (a), the result holds since both $\mathrm{tr}(\Sigma^{-1}S_n) \ge 0$ and $\Pi(\Sigma) \ge 0$. For case (b), the result follows from condition (i) since $\mathrm{tr}(\Sigma^{-1}S_n) \ge 0$ and $\Pi(\Sigma) \to \infty$.

When case (c) holds, consider the two sub-cases: (c1) λ_1 is bounded away from zero, and (c2) $\lambda_1 \to 0$. If (c1) holds, condition (ii) implies $(\log\det \Sigma)/\Pi(\Sigma) \to 0$ and so $\Pi(\Sigma) \to \infty$. Hence, for $\eta > 0$,

$$L(\Sigma; S_n, \eta) = \mathrm{tr}(\Sigma^{-1}S_n) + \Pi(\Sigma)\{(\log\det \Sigma)/\Pi(\Sigma) + \eta\} \to \infty.$$

If (c2) holds, since $\mathrm{tr}(\Sigma^{-1}S_n) \ge \mathrm{tr}(S_n)/\lambda_1$ and $\log\det \Sigma \ge \log \lambda_1 - (q-1)\log \lambda_q$, it follows that

$$L(\Sigma; S_n, \eta) \ge \mathrm{tr}(S_n)/\lambda_1 + q \log \lambda_1 + (q-1)\log(\lambda_q/\lambda_1) + \eta \Pi(\Sigma),$$

with $\mathrm{tr}(S_n)/\lambda_1 + q\log\lambda_1 \to \infty$. So, if $\lambda_1/\lambda_q \to 1$, then $L(\Sigma; S_n, \eta) \to \infty$. Whereas, if λ_1/λ_q is bounded away from one, then by condition (iii), $(q-1)\log(\lambda_q/\lambda_1) + \eta\Pi(\Sigma) = \Pi(\Sigma)\{(q-1)\log(\lambda_q/\lambda_1)/\Pi(\Sigma) + \eta\}$ is bounded below and so $L(\Sigma; S_n, \eta) \to \infty$.

A.4. Proof of Theorem 3.1

First, we show (i)⇒(iii). Suppose that $F(\Sigma)$ is (strictly) g-convex, then by Lemma 3.6 of Dümbgen and Tyler (2016), $F(BD(e^x)B^{\mathrm{T}})$ is (strictly) convex in $x \in \mathbb{R}^q \backslash \{0\}$ for any nonsingular B of order q. Here, for $y \in \mathbb{R}^q$, $D(y)$ represents the diagonal matrix with the elements of y corresponding to its diagonal elements. Thus, by Lemma 3.1, $f(x) = F(D(e^x))$ is (strictly) convex.

Next, we show (iii)⇒(i). Here, the concept of majorization plays an important role. For a vector $v \in \mathbb{R}^q$, denote its ordered values by $v_{(1)} \ge \cdots \ge v_{(q)}$. A vector $y \in \mathbb{R}^q$ is then said to majorize a vector $x \in \mathbb{R}^q$, denoted $x \prec y$ if and only if $\sum_{j=1}^{k} x_{(j)} \le \sum_{j=1}^{k} y_{(j)}$, with equality when $k = q$. As stated in Theorem 1.3 of Ando (1957), $x \prec y$ if an only if x is a convex combination of coordinate permutations of y, that is,

$$x \prec y \Leftrightarrow x = \sum_{j=1}^{q} w_j P_j y, \tag{A.1}$$

where, for $j = 1, \ldots, q$, P_j is a permutation matrix of order q, hence orthogonal, and $w_j \ge 0$ with $\sum_{j=1}^{q} w_j = 1$. As a side note, the Birkhoff–von Neumann theorem notes that Q is a doubly stochastic matrix of order q if and only if it has the representation $Q = \sum_{j=1}^{q} w_j P_j$. For $\Sigma > 0$, let $\lambda(\Sigma) = (\lambda_1(\Sigma), \ldots, \lambda_q(\Sigma))$ denote the vector of the ordered eigenvalues of Σ. An important result given by Lemma 2.17 in Sra and Hosseini (2015) states

$$\log(\lambda(\Sigma_t)) \prec (1-t)\log(\lambda(\Sigma_0)) + t\log(\lambda(\Sigma_1)), \tag{A.2}$$

where Σ_t is the geodesic curve from Σ_0 and Σ_1. So, by (A.1), we can express

$$\log \lambda(\Sigma_t) = Q\{(1-t)\log \lambda(\Sigma_0) + t\log\lambda(\Sigma_1)\}, \tag{A.3}$$

with $Q = \sum_{j=1}^{q} w_j P_j$ being defined as in (A.1). Thus,

$$\begin{aligned} F(\Sigma_t) &= f(\log \lambda(\Sigma_t)) = f\left(Q[(1-t)\log\lambda(\Sigma_0) + t\log\lambda(\Sigma_1)]\right) \\ &\le (1-t)f(Q\log\lambda(\Sigma_0)) + tf(Q\log\lambda(\Sigma_1)) \\ &\le (1-t)\sum_{j=1}^{q} w_j f(P_j \log\lambda(\Sigma_0)) + t\sum_{j=1}^{q} w_j f(P_j\log\lambda(\Sigma_1)) \\ &= (1-t)\sum_{j=1}^{q} w_j f(\log\lambda(\Sigma_0)) + t\sum_{j=1}^{q} w_j f(\log\lambda(\Sigma_1)) \\ &= (1-t)F(\Sigma_0) + tF(\Sigma_1). \end{aligned}$$

The two inequalities above follow from condition (iii), that is, f is convex. Suppose now that f is strictly convex, then the first inequality is strict unless $\lambda(\Sigma_0) = \lambda(\Sigma_1)$, and the second inequality is strict unless $Q = I_q$. Thus, both equality holds if and only if $\lambda(\Sigma_t) = \lambda(\Sigma_0)$ for $0 \le t \le 1$. However, since $\mathrm{tr}(\Sigma)$ is strictly g-convex (see, e.g., Wiesel and Zhang 2015, Lemma 1.15), it follows that $\mathrm{tr}(\Sigma_t) < (1-t)\mathrm{tr}(\Sigma_0) + t\,\mathrm{tr}(\Sigma_1) = \mathrm{tr}(\Sigma_0)$ for $0 < t < 1$, unless $\Sigma_0 = \Sigma_1$. Thus, $F(\Sigma)$ is strictly g-convex.

Finally, we note the statement (ii)⇔(iii) follows from the main theorem in Davis (1957), at least in the convex case. The strictly convex case can be shown to hold by applying arguments analogous to those used in the (i)⇔(iii) case.

A.5. Proof That $\Pi_{R,s}$ Satisfies the Conditions of Lemma 2.3

Again let $a_j = \log \lambda_j$, and so $\bar{a} = q^{-1} \log \det \Sigma$, $\sum_{j=1}^{q} a_j^2 = \| \log \Sigma \|_F^2$ and $Q(a) \equiv \sum_{j=1}^{q} (a_j - \bar{a})^2 = \Pi_{R,s}(\Sigma)$. Condition (i) states that if $|\bar{a}|$ is bounded above and $\sum_{j=1}^{q} a_j^2 \to \infty$, then $Q(a) \to \infty$, which holds since $Q(a) = \sum_{j=1}^{q} a_j^2 - q\bar{a}^2$. Condition (ii) states that if $\bar{a} \to -\infty$ and a_1 is bounded below, then $\bar{a}/Q(a) \to 0$. To show this, express $\bar{a}/Q(a) = \{\bar{a}\, Q(b)\}^{-1}$, where $b_j = a_j/\bar{a}$. Since a_1 is bounded below, $b_1 \to 0$ and $\sum_{j=1}^{q} b_j \to q$. Hence, $Q(b)$ must be bounded away from zero, which implies $\bar{a}/Q(a) \to 0$. Condition (iii) states that if $a_1 \to -\infty$ and $a_1 - a_q > \epsilon > 0$, then $(a_1 - a_q)/Q(a)$ is bounded above. This follows since $Q(a) \geq (a_1 - a_q)^2$ and so $(a_1 - a_q)/Q(a) \leq 1/(a_1 - a_q) \leq 1/\epsilon$.

A.6. Proof of Theorem 4.1

(i) Let $\Sigma_0 \#_t \Sigma_1 := \Sigma_0^{1/2} (\Sigma_o^{-1/2} \Sigma_1 \Sigma_o^{-1/2})^t \Sigma_o^{1/2}, t \in [0, 1]$, and so $\Sigma_t = \Sigma_0 \#_t \Sigma_1$ (Sra and Hosseini 2015). It readily follows that $V(\Sigma_t) = V(\Sigma_0) \#_t V(\Sigma_1)$, and so

$$\begin{aligned}\Pi_s(\Sigma_t) &= \Pi\{V(\Sigma_t)\} = \Pi\{V(\Sigma_0) \#_t V(\Sigma_1)\} \\ &\leq (1-t)\, \Pi\{V(\Sigma_0)\} + t\, \Pi\{V(\Sigma_1)\} \\ &= (1-t)\, \Pi_s(\Sigma_0) + t\, \Pi_s(\Sigma_1).\end{aligned}$$

(ii) Let $A = \log \Sigma$, and define $G(A) = \Pi(e^A)$ and $G_s(A) = \Pi_s(e^A)$. The goal is to show that if $G(A)$ is convex in A, then $G_s(A)$ is also convex in A. Since $G_s(A) = G(\tilde{A})$, where $\tilde{A} \equiv A - (\mathrm{tr}(A)/q) * I_w$, and so

$$\begin{aligned}G_s((1-t)A_0 + tA_1) &= G((1-t)\tilde{A}_0 + t\tilde{A}_1) \\ &\leq (1-t)G(\tilde{A}_0) + tG(\tilde{A}_1) = (1-t)G_s(A_0) + tG_s(A_1).\end{aligned}$$

A.7. Proof of Theorem 5.1

The proof relies on the following property of eigenvalues of symmetric matrices. Let B be a symmetric matrix of order q, and let $C = [c_1 \cdots c_m]$ be of order $q \times m$, $m \leq q$, with orthonormal columns. Then $\mathrm{tr}\{C^{\mathrm{T}} B C\} = \sum_{j=1}^{m} c_j^{\mathrm{T}} B c_j$ is bounded above and below by the sum of the largest m and the sum of the smallest m eigenvalues of B, respectively.

Expressing $\Sigma = P \Lambda P^{\mathrm{T}}$ in terms of its spectral value decomposition, let $H = [h_1 \cdots h_q] = P_n^{\mathrm{T}} P$, which is itself an orthogonal matrix. Define $\kappa_1 = 1/\lambda_1$ and $\kappa_j = 1/\lambda_j - 1/\lambda_{j-1}$ for $j \neq 1$. Inverting this relationship gives $\lambda_j^{-1} = \sum_{k=1}^{j} \kappa_k$. Since $\kappa_j \geq 0$, the above noted property of eigenvalues of a symmetric matrix implies

$$\begin{aligned}\mathrm{tr}\{\Sigma^{-1} S_n\} &= \mathrm{tr}\{\Lambda^{-1} H^{\mathrm{T}} D_n H\} = \sum_{j=1}^{q} \lambda_j^{-1} h_j^{\mathrm{T}} D_n h_j \\ &= \sum_{k=1}^{q} \kappa_k \left\{ \sum_{j=k}^{q} h_j^{\mathrm{T}} D_n h_j \right\} \\ &\geq \sum_{j=1}^{q} \kappa_j \left\{ \sum_{k=j}^{q} d_k \right\} = \sum_{j=1}^{q} d_j/\lambda_j = \mathrm{tr}\{\Lambda^{-1} D_n\},\end{aligned}$$

with equality when $H = I_q$, that is, when $P = P_n$. This gives the inequality in the theorem since $\det \Sigma = \det \Lambda$ and $\Pi(\Sigma) = \Pi(\Lambda)$.

A.8. Proof of Corollary 5.1

Suppose $L(\widetilde{\Sigma}_\eta; S_n, \eta) \leq L(\Sigma; S_n, \eta)$ for all $\Sigma > 0$, with the eigenvalues of $\widetilde{\Sigma}_\eta$ being $\widetilde{\lambda}_{1,\eta} \geq \cdots \geq \widetilde{\lambda}_{q,\eta} > 0$. Let $\widehat{\Sigma}_\eta$ be defined as in the corollary. Theorem 5.1 then implies $L(\widehat{\Sigma}_\eta; S_n, \eta) \leq L(\widetilde{\Sigma}_\eta; S_n, \eta) \leq L(\Sigma; S_n, \eta)$ for all $\Sigma > 0$, and so $\widehat{\Sigma}_\eta$ must also be a minimizer of $L(\Sigma; S_n, \eta)$ over $\Sigma > 0$. Further, $\widehat{\Lambda}_\eta$ must be a global minimizer of $L(P_n \Lambda P_n^{\mathrm{T}}; S_n, \eta)$ over $\Lambda \in \mathcal{D}_+$, otherwise we have the contradiction $L(P_n \Lambda P_n^{\mathrm{T}}; S_n, \eta) < L(\widehat{\Sigma}_\eta; S_n, \eta)$ for some $\Lambda \in \mathcal{D}_+$. If $L(\Sigma; S_n, \eta)$ has a unique global minimizer over $\Sigma > 0$, then $\widehat{\Sigma}_\eta = P_n \widehat{\Lambda}_\eta P_n^{\mathrm{T}}$ is uniquely determined.

A.9. Proof of Theorem 5.2 and Derivation of the Algorithms (5) and (6)

Algorithm (5) finds the unique minimum of the univariate function $f(a) = de^{-a} + a + \eta(a - c)^2$. Newton's method is given by $a_{k+1} = a_k - f'(a)/f''(a)$. Here, we have $f'(a) = -de^{-a} + 1 + 2\eta(a - c)$ and $f''(a) = de^{-a} + 2\eta > 0$. Algorithm (5) then follows. Furthermore, $f'''(a) = -de^{-a} < 0$, which implies that $F(a) = -f'(a)$ is a convex function. The convergence results for algorithm (5) then follow from the monotone Newton theorem and the global Newton theorem as stated by Theorems 13.3.4 and 13.3.7, respectively, in Ortega and Rheinboldt (2000).

Algorithm (6) finds the unique minimum of $f(a) = \sum_{j=1}^{q} \{d_i e^{-a_i} + a_i + \eta(a_i - \bar{a})^2\}$ over $a \in \mathbb{R}^q$. Newton's method is given by $a_{k+1} = a_k - \{\nabla^2 f(a_k)\}^{-1} \nabla f(a_k)$, where $\nabla f(a)$ and $\nabla^2 f(a)$ are the gradient and Hessian of $f(a)$, respectively. The gradient is given by $\partial f(a)/\partial a_i = -d_i e^{-a_i} + 1 + 2\eta(a_i - \bar{a})$. For the hessian, we have $\partial^2 f(a)/\{\partial a_i\}^2 = d_i e^{-a_i} + 2\eta(1 - 1/q)$ and $\partial^2 f(a)/\{\partial a_i \partial a_j\} = -2\eta/q$ for $i \neq j$, which gives $\nabla^2 f(a) = \mathrm{diag}\{d_i e^{-a_i}\} + 2\eta(I_q - q^{-1} 1_q 1_q') > 0$. Algorithm (6) follows after noting $\{\nabla^2 f(a)\}^{-1} = \{m_{ij}\}$, with $m_{ii} = \delta_i + \beta\delta_i^2$ and $m_{ij} = \beta\delta_i\delta_j$ for $i \neq j$, where $\delta_i = 1/(d_i e^{-a_i} + 2\eta)$ and $\beta = 2\eta/(q - 2\eta \sum_{i=1}^{q} \delta_i)$. For the third partial derivatives, we have $\partial^3 f(a)/\{\partial a_i\}^3 = -d_i e^{-a_i} < 0$, with all other third partial derivatives being zero. This is sufficient to establish that $F(a) = -\nabla f(a)$ is a convex function from $\mathbb{R}^q \to \mathbb{R}^q$ as defined by Definition 13.3.1 in Ortega and Rheinboldt (2000). The convergence results for algorithm (6) again follow from the monotone and global Newton theorems.

A.10. Limiting Behavior of the sLogF Estimator as $\eta \to \infty$

As $\eta \to \infty$, the penalty term $\Pi_{R,s}(\widehat{\Sigma}_\eta)$ must go to 0, which implies $\widehat{\Sigma}_\eta$ is proportional to I_q in the limit. The eigenvalues of $\widehat{\Sigma}_\eta$ correspond to the unique critical point of $\sum_{j=1}^{q} \{d_j e^{-a_j} + a_j + \eta(a_j - \bar{a})^2\}$, where again $a_j = \log \lambda_j$, which in turn corresponds to the unique solution to the set of equations $d_j e^{-a_j} = 1 + 2\eta(a_j - \bar{a})$ for $j = 1, \ldots, q$. By taking the sum, we obtain $q = \sum_{j=1}^{q} d_j e^{-a_j} = \sum_{j=1}^{q} d_j/\lambda_j$ for any $\eta \geq 0$. Hence, since the eigenvalues of $\widehat{\Sigma}_\eta$ approach each other as $\eta \to \infty$, it follows that $\widehat{\lambda}_j \to \bar{d}$ or $\widehat{\Sigma}_\eta \to \bar{d} I_q$.

Supplementary Materials

The supplementary materials include the following: a supplementary manuscript reporting further simulation results, an R-package **logconvx** for computing the proposed penalized covariance matrices, an Rdata workspace **Sim.Rdata** for reproducing the simulations reported in the manuscript, and an RData workspace **Sonar.Rdata** for reproducing the results for the example given in section 7.

Funding

Research for both authors was supported in part by the National Science Foundation grants DMS-1407751 and DMS-1812198. Mengxi Yi's research was also supported in part by the Austrian Science Fund (FWF) under grant P31881-N32 and in part by the Scientific Research Starting Foundation of UIBE.

References

Ando, T. (1989), "Majorization, Doubly Stochastic Matrices, and Comparison of Eigenvalues," *Linear Algebra and Its Applications*, 118, 163–248. [449]

Bhatia, R. (2009), *Positive Definite Matrices* (Vol. 24), Princeton, NJ: Princeton University Press. [442,443,444]

Bickel, P. J., and Levina, E. (2008), "Regularized Estimation of Large Covariance Matrices," *The Annals of Statistics*, 36, 199–227. [448]

Cai, T., Zhang, C.-H., and Zhou, H. (2010), "Optimal Rates of Convergence for Covariance Matrix Estimation," *The Annals of Statistics*, 38, 2118–2144. [448]

Davis, C. (1957), "All Convex Invariant Functions of Hermitian Matrices," *Archiv der Mathematik*, 8, 276–278. [449]

Deng, X., and Tsui, K. W. (2013), "Penalized Covariance Matrix Estimation Using a Matrix-Logarithm Transformation," *Journal of Computational and Graphical Statistics*, 22, 494–512. [442,443,444,445,446,447,448]

Dümbgen, L., and Tyler, D. E. (2016), "Geodesic Convexity and Regularized Scatter Estimators," arXiv no. 1607.05455. [443,444,449]

Friedman, J., Hastie, T., and Tibshirani, R. (2008), "Sparse Inverse Covariance Estimation With the Graphical Lasso," *Biostatistics*, 9, 432–441. [442]

Gorman, R. P., and Sejnowski, T. J. (1988), "Analysis of Hidden Units in a Layered Network Trained to Classify Sonar Targets," *Neural Networks*, 1, 75–89. [448]

Lam, C., and Fan, J. (2009), "Sparsistency and Rates of Convergence in Large Covariance Matrix Estimation," *The Annals of Statistics*, 37, 4254–4278. [448]

Moakher, M. (2005), "A Differential Geometric Approach to the Geometric Mean of Symmetric Positive-Definite Matrices," *SIAM Journal on Matrix Analysis and Applications*, 26, 735–747. [442]

Ortega, J. M., and Rheinboldt, W. C. (2000), *Iterative Solution of Nonlinear Equations in Several Variables*, Philadelphia, PA: Society for Industrial and Applied Mathematics. [450]

Paindaveine, D. (2008), "A Canonical Definition of Shape," *Statistics & Probability Letters*, 78, 2240–2247. [444]

Skovgaard, L. T. (1984), "A Riemannian Geometry of the Multivariate Normal Model," *Scandinavian Journal of Statistics*, 11, 211–223. [443]

Sra, S., and Hosseini, R. (2015), "Conic Geometric Optimization on the Manifold of Positive Definite Matrices," *SIAM Journal on Optimization*, 25, 713–739. [449,450]

Tyler, D. E., and Yi, M. (2020), "Lassoing Eigenvalues," *Biometrika*, 107, 397–414. [448]

Wiesel, A. (2012), "Unified Framework to Regularized Covariance Estimation in Scaled Gaussian Models," *IEEE Transactions on Signal Processing*, 60, 29–38. [444]

Wiesel, A., and Zhang, T. (2015), "Structured Robust Covariance Estimation," *Foundations and Trends in Signal Processing*, 8, 127–216. [443,448,449]

Yu, P. L. H., Wang, X., and Zhu, Y. (2017), "High Dimensional Covariance Matrix Estimation by Penalizing the Matrix-Logarithm Transformed Likelihood," *Computational Statistics & Data Analysis*, 114, 12–25. [443,445,446,447,448]

Yuan, M., and Lin, Y., (2007), "Model Selection and Estimation in the Gaussian Graphical Model," *Biometrika*, 94, 19–35. [442]

Zhang, T., Wiesel, A., and Greco, M. S. (2013), "Multivariate Generalized Gaussian Distribution: Convexity and Graphical Models," *IEEE Transactions on Signal Processing*, 61, 4141–4148. [443]