This is an electronic reprint of the original article.
This reprint may differ from the original in pagination and typographic detail.

Tabassum, Muhammad Naveed; Ollila, Esa

A Compressive Classification Framework for High-Dimensional Data

Published in:
IEEE Open journal of Signal Processing

DOI:
[10.1109/OJSP.2020.3037825](https://doi.org/10.1109/OJSP.2020.3037825)

Published: 01/01/2020

Document Version
Publisher's PDF, also known as Version of record

Published under the following license:
CC BY

Please cite the original version:
Tabassum, M. N., & Ollila, E. (2020). A Compressive Classification Framework for High-Dimensional Data. *IEEE Open journal of Signal Processing*, 1, 177-186. <https://doi.org/10.1109/OJSP.2020.3037825>

This material is protected by copyright and other intellectual property rights, and duplication or sale of all or part of any of the repository collections is not permitted, except that material may be duplicated by you for your research use or educational purposes in electronic or print form. You must obtain permission for any other use. Electronic or print copies may not be offered, whether for sale or otherwise to anyone who is not an authorised user.

A Compressive Classification Framework for High-Dimensional Data

MUHAMMAD NAVEED TABASSUM (Member, IEEE), AND ESA OLLILA (Member, IEEE)

Department of Signal Processing and Acoustics, Aalto University, FI-00076 Aalto, Finland

CORRESPONDING AUTHOR: ESA OLLILA (e-mail: esa.ollila@aalto.fi)

This work was supported by the Academy of Finland under Grant 298118.

ABSTRACT We propose a compressive classification framework for settings where the data dimensionality is significantly larger than the sample size. The proposed method, referred to as compressive regularized discriminant analysis (CRDA), is based on linear discriminant analysis and has the ability to select significant features by using joint-sparsity promoting hard thresholding in the discriminant rule. Since the number of features is larger than the sample size, the method also uses state-of-the-art regularized sample covariance matrix estimators. Several analysis examples on real data sets, including image, speech signal and gene expression data illustrate the promising improvements offered by the proposed CRDA classifier in practise. Overall, the proposed method gives fewer misclassification errors than its competitors, while at the same time achieving accurate feature selection results. The open-source R package and MATLAB toolbox of the proposed method (named compressiveRDA) is freely available.

INDEX TERMS Classification, covariance matrix estimation, discriminant analysis, feature selection, high-dimensional statistics, joint-sparse recovery.

I. INTRODUCTION

High-dimensional (HD) classification is at the core of numerous contemporary statistical studies. Commonly large amounts of information may be collected on each individual sample point, yet the number of sample points themselves may remain relatively small. Typical examples of such cases are gene expression and protein mass spectrometry data, and other areas of computational biology. Regularization and shrinkage are commonly used tools in many applications such as regression or classification to overcome significant statistical challenges posed particularly due to the high dimension low-sample-size (HDLSS) data settings in which the number of features, p , is often several magnitudes larger than the sample size, n (i.e., $p \gg n$). In this paper, we propose a novel classification framework for such HD data.

In many applications, the ideal classifier of an HD input vector should have three characteristics: (i) the proportion of correctly classified observations is high, (ii) the method is computationally feasible, allowing the training of the classifier in real-time with a personal computer, and (iii) the classifier selects only the significant features, thus yielding interpretable models. The last property is important for

example in gene expression studies, where it is important to identify which genes influence the trait (e.g., disease). This is obviously important for the biological understanding of the disease process but can help in the development of clinical tests for early diagnosis [1]. We propose a HD classification method that takes into account all of the aforementioned three key aspects.

Some of the standard off-the-shelf classification methods either lack feature elimination or require that it is done beforehand. For example, the basic support vector machine (SVM) classifier uses all features. Some other classifiers on the other hand have higher computational cost, e.g., Logit-ASPLS approach (with R-package: *plsgenomics*) [2] depends on three hyperparameters. This increases the computational complexity of the method, especially when finer grids of hyperparameters are used and tuned using cross-validation (CV). In this paper, we propose a classification framework referred to as *compressive regularized discriminant analysis* (CRDA) method, which is computationally efficient, performs feature selection during the learning process and includes the shrunken centroids (SC)RDA method [3] as a special case. Our simulation studies and real data analysis examples

illustrate that CRDA often achieves superior classification performance compared to off-the-shelf classification approaches in the HDLSS settings.

Many classification techniques assign a p -dimensional observation $\mathbf{x}$ to one of the G classes (groups or populations) based on the following rule:

$$\mathbf{x} \in \text{group} \left[\tilde{g} = \arg \max_g d_g(\mathbf{x}) \right], \quad (1)$$

where $d_g(\mathbf{x})$ is called the *discriminant function* for population $g \in \{1, \dots, G\}$. In linear discriminant analysis (LDA), $d_g(\mathbf{x})$ is a linear function of $\mathbf{x}$, $d_g(\mathbf{x}) = \mathbf{x}^\top \mathbf{b}_g + c_g$, for some constant $c_g \in \mathbb{R}$ and vector $\mathbf{b}_g \in \mathbb{R}^p$, $g = 1, \dots, G$. The vector $\mathbf{b}_g = \mathbf{b}_g(\mathbf{\Sigma})$ depends on the unknown covariance matrix $\mathbf{\Sigma}$ of the populations (via its inverse matrix) which is commonly estimated using the pooled sample covariance matrix (SCM). In the HD settings, the SCM is no-longer invertible, and regularized SCM (RSCM) $\hat{\mathbf{\Sigma}}$ is used for constructing an estimated discriminant function $\hat{d}_g(\mathbf{x})$. Such approaches are commonly referred to as regularized LDA methods, which we refer shortly as RDA. See e.g., [1], [3]–[6].

Next note that if the i -th entry of $\mathbf{b}_g$ is zero, then the i -th feature does not contribute in the classification of g -th population. To eliminate unnecessary features, one can shrink $\mathbf{b}_g$ using element-wise soft-thresholding as in the shrunken centroids SCRDA method [3]. These methods are often difficult to tune because the shrinkage threshold parameter is the same across all groups, but different populations would often benefit from different shrinkage intensity. Consequently, they tend to yield higher false-positive rates (FPRs).

Element-wise shrinkage does not achieve simultaneous feature elimination because the eliminated feature from group i may still affect the discriminant function of group j . We propose CRDA classification framework which generalizes the SCRDA method by allowing more general sparsifying transformation functions to be used in the linear discriminant rule. In particular, we promote the usage of *joint-sparsity* via $\ell_{q,1}$ -norm based hard-thresholding. Another advantage is of having a shrinkage parameter that is much easier to tune and interpret: namely, the *joint-sparsity level* $K \in \{1, 2, \dots, p\}$ instead of a shrinkage threshold $\Delta \in [0, \infty)$ used in the soft-thresholding function of SCRDA. We also promote the usage of two recent state-of-the-art regularized covariance matrix estimators proposed in [7] and [8], [9], respectively. The first estimator, called the Ell-RSCM estimator [7], is designed to attain the minimum mean squared error and is fast to compute due to automatic data-adaptive tuning of the involved regularization parameter. The second estimator is a penalized SCM (PSCM) proposed by [8], which we refer to as Rie-PSCM, that minimizes a penalized negative Gaussian likelihood function with a Riemannian distance penalty. Although a closed-form solution of the optimization problem is not possible to find, the estimator can be computed using a relatively fast iterative Newton-Raphson algorithm developed in [9].

We test the effectiveness of the proposed CRDA methods using partially synthetic genomic data as well as on nine real

data sets and compare the obtained results against seven state-of-the-art HD classifiers that also perform feature selection. The results illustrate that overall, the proposed CRDA method gives fewer misclassification errors than its competitors while achieving at the same time accurate feature selection with minimal computational efforts. Preliminary results of CRDA approach has appeared in [10].

The paper is structured as follows. Section II describes the regularized LDA procedure and the regularized covariance matrix estimators that will be used by the CRDA classifiers. In Section III, CRDA is introduced as a compressive version of regularized LDA performing feature selection simultaneously with classification. Classification results for 3 different types (image, speech, genomics) of real data sets are also presented. Section IV illustrates the promising performance of CRDA in feature selection for partially synthetic data set while Section V provides analysis results for six real gene expression data sets against a large number of benchmark classifiers. Section VI provides discussion and concludes the paper. MATLAB and R toolboxes are available at <http://users.spa.aalto.fi/esollila/crda> and <https://github.com/mntabassm/compressiveRDA>.

Notations: We use $\|\cdot\|_F$ to denote the Frobenius (matrix) norm, defined as $\|\mathbf{A}\|_F^2 = \text{tr}(\mathbf{A}^\top \mathbf{A})$ for all matrices $\mathbf{A}$. An ℓ_q -norm is defined as

$$\|\mathbf{a}\|_q = \begin{cases} (|a_1|^q + \dots + |a_n|^q)^{1/q}, & \text{for } q \in [1, \infty) \\ \max(|a_i| : i = 1, \dots, n), & \text{for } q = \infty \end{cases}$$

for any vector $\mathbf{a} \in \mathbb{R}^n$. The soft-thresholding (ST)-operator is defined as

$$\mathcal{S}_\Delta(b) = \text{sign}(b)(|b| - \Delta)_+,$$

where $(x)_+ = \max(x, 0)$ for $x \in \mathbb{R}$ and $\Delta > 0$ is a threshold constant. Finally, $\mathbf{I}(\cdot)$ denotes the indicator function which takes value 1 when true and value 0 otherwise. $\mathbb{E}[\cdot]$ denotes expectation.

II. REGULARIZED LDA

We are given a p -variate random vector $\mathbf{x}$ which we need to classify into one of G groups or populations. In LDA, one assumes that the class populations are p -variate multivariate normal (MVN) with a common positive definite symmetric covariance matrix $\mathbf{\Sigma}$ over each class but having distinct class mean vectors $\boldsymbol{\mu}_g \in \mathbb{R}^p$, $g = 1, \dots, G$. The problem is then to classify $\mathbf{x}$ to one of the MVN populations, $\mathcal{N}_p(\boldsymbol{\mu}_g, \mathbf{\Sigma})$, $g = 1, \dots, G$. Sometimes prior knowledge is available on proportions of each population and we denote by p_g the prior probabilities of the classes ($\sum_{g=1}^G p_g = 1$). LDA uses the rule (1) with discriminant function

$$d_g(\mathbf{x}) = \mathbf{x}^\top \mathbf{b}_g - \frac{1}{2} \boldsymbol{\mu}_g^\top \mathbf{b}_g + \ln p_g,$$

where $\mathbf{b}_g = \mathbf{\Sigma}^{-1} \boldsymbol{\mu}_g$ for $g = 1, \dots, G$.

The unknown parameters in LDA are the class mean vectors $\{\boldsymbol{\mu}_g\}_{g=1}^G$ and the covariance matrix $\mathbf{\Sigma}$. These are estimated from the *training data set* $\mathbf{X} = (\mathbf{x}_1 \dots \mathbf{x}_n)$ which consists

of n_g observations from each class, $g = 1, \dots, G$. Let y_i denote the class label associated with the i -th observation, so $y_i \in \{1, \dots, G\}$. Then $n_g = \sum_{i=1}^n \mathbf{I}(y_i = g)$ is the number of observations belonging to g -th population, and we denote by $\pi_g = n_g/n$ the relative sample proportions. We assume observations in the training data set $\mathbf{X}$ are centered by the sample mean vectors of the classes,

$$\hat{\boldsymbol{\mu}}_g = \bar{\mathbf{x}}_g = \frac{1}{n_g} \sum_{y_i=g} \mathbf{x}_i. \quad (2)$$

Since the data matrix $\mathbf{X}$ is centered group-wise, the pooled (over groups) SCM can be written simply as

$$\mathbf{S} = \frac{1}{n} \mathbf{X} \mathbf{X}^\top. \quad (3)$$

In practice, an observation $\mathbf{x}$ is classified using an estimated discriminant function,

$$\hat{d}_g(\mathbf{x}) = \mathbf{x}^\top \hat{\mathbf{b}}_g - \frac{1}{2} \hat{\boldsymbol{\mu}}_g^\top \hat{\mathbf{b}}_g + \ln \pi_g, \quad (4)$$

where $\hat{\mathbf{b}}_g = \hat{\boldsymbol{\Sigma}}^{-1} \hat{\boldsymbol{\mu}}_g$, $g = 1, \dots, G$ and $\hat{\boldsymbol{\Sigma}}$ is an estimator of $\boldsymbol{\Sigma}$. Note that in (4) the prior probabilities p_g -s are replaced by their estimates, π_g -s. Commonly, the pooled SCM $\mathbf{S}$ in (3) is used as an estimator $\hat{\boldsymbol{\Sigma}}$ of the covariance matrix $\boldsymbol{\Sigma}$.

Since $p \gg n$, the pooled SCM is no longer invertible. To avoid the singularity of the estimated covariance matrix, a commonly used approach (cf. [11]) is to use a *regularized SCM (RSCM)* that shrinks the estimator towards a target matrix. One such method is reviewed in Section II-A. An alternative approach is to use penalized likelihood approach wherein a penalty term is added to the likelihood function. The idea is to penalize covariance matrices that are far from the desired target matrix. One prominent estimator of this type is reviewed in Section II-B. There are many other covariance matrix estimators for high-dimensional data, including Bayes estimators [12], structured covariance matrix estimators [13], [14], robust shrinkage estimators [15]–[18], etc, but most of these are not applicable in HDLSS setting due to incurred high computational cost.

A. REGULARIZED SCM ESTIMATORS

A commonly used RSCM estimator (see e.g., [7], [11], [19]) is of the form:

$$\hat{\boldsymbol{\Sigma}}(\alpha) = \alpha \mathbf{S} + (1 - \alpha) [\text{tr}(\mathbf{S})/p] \mathbf{I}, \quad (5)$$

where $\alpha \in [0, 1)$ denotes the shrinkage (or regularization) parameter and $\mathbf{I}$ denotes an identity matrix of appropriate size. We use the automated method for choosing the tuning parameter α proposed in [7] and the respective estimators, called Ell1-RSCM and Ell2-RSCM. These estimators use $\hat{\alpha}$ that is a consistent estimator of α_o minimizing the mean squared error,

$$\alpha_o = \arg \min_{\alpha \in [0, 1)} \mathbb{E} [\|\hat{\boldsymbol{\Sigma}}(\alpha) - \boldsymbol{\Sigma}\|_F^2]. \quad (6)$$

See Appendix A for details.

Note that SCRDA [3] uses a different estimator, $\hat{\boldsymbol{\Sigma}} = \alpha \mathbf{S} + (1 - \alpha) \mathbf{I}$, where the regularization parameter α is chosen using cross-validation (CV). The benefit of Ell1-RSCM and Ell2-RSCM estimators are that they avoid computationally intensive CV scheme. In the HD settings, another computational burden is related to inverting the matrix $\hat{\boldsymbol{\Sigma}}$ in (5). The matrix inversion can be done efficiently using the SVD-trick detailed in Appendix B.

B. PENALIZED SCM ESTIMATOR USING RIEMANNIAN DISTANCE

Assume that $\{\mathbf{x}_i\}_{i=1}^n$ is a random sample from a MVN distribution. Then the maximum-likelihood estimator of $\boldsymbol{\Sigma}$ is found by minimizing the $-(2/n) \times$ Gaussian log-likelihood function,

$$L(\boldsymbol{\Sigma}; \mathbf{S}) = \text{tr}(\boldsymbol{\Sigma}^{-1} \mathbf{S}) + \log |\boldsymbol{\Sigma}|$$

over the set of positive definite symmetric $p \times p$ matrices, denoted by $\mathbb{S}_{++}^{p \times p}$. When $n > p$ and the SCM $\mathbf{S}$ is non-singular, the minimum is attained at $\hat{\boldsymbol{\Sigma}} = \mathbf{S}$. Since $p \gg n$ and $\mathbf{S}$ is singular, it is common to add a penalty function $\Pi(\boldsymbol{\Sigma})$ to the likelihood function to enforce a certain structure to the covariance matrix and to guarantee that the objective function has a unique (positive definite matrix) solution. One such penalized SCM (PSCM) proposed in [8] uses a Riemannian distance of $\boldsymbol{\Sigma}$ from a scaled identity matrix, $\Pi(\boldsymbol{\Sigma}; m) = \|\log(\boldsymbol{\Sigma}) - \log(m) \mathbf{I}\|_F^2$, where $m > 0$ is a shrinkage constant. One then solves the penalized likelihood function:

$$\hat{\boldsymbol{\Sigma}} = \arg \min_{\boldsymbol{\Sigma} \in \mathbb{S}_{++}^{p \times p}} L(\boldsymbol{\Sigma}; \mathbf{S}) + \eta \|\log(\boldsymbol{\Sigma}) - \log(m) \mathbf{I}\|_F^2, \quad (7)$$

where $\eta > 0$ is the penalty parameter. As in [9], we choose m as the mean of the eigenvalues ($d_1 \geq \dots \geq d_p \geq 0$) of the SCM, i.e.,

$$m = \frac{\text{tr}(\mathbf{S})}{p} = \frac{1}{p} \sum_{i=1}^p d_i.$$

The eigenvalues of $\hat{\boldsymbol{\Sigma}}$ are then shrunk towards m , and $\hat{\boldsymbol{\Sigma}} \rightarrow m \mathbf{I}$ in the limit when $\eta \rightarrow \infty$ [8]. The solution to (7), which we refer to as *Rie-PSCM estimator*, does not have a closed-form solution, but a relatively efficient Newton-Raphson algorithm [9] can be used to compute the solution. We choose the associated penalty parameter η using the cross-validation procedure proposed in [9].

III. COMPRESSIVE RDA

In order to explain the proposed compressive RDA (CRDA) approach, we first write the discriminant rule in (4) in vector form as

$$\begin{aligned} \hat{\mathbf{d}}(\mathbf{x}) &= (\hat{d}_1(\mathbf{x}), \dots, \hat{d}_G(\mathbf{x})) \\ &= \mathbf{x}^\top \hat{\mathbf{B}} - \frac{1}{2} \text{diag}(\hat{\mathbf{M}}^\top \hat{\mathbf{B}}) + \ln \boldsymbol{\pi}, \end{aligned} \quad (8)$$

FIGURE 1. Computing the K -row-sparse coefficient matrix $\hat{\mathbf{B}}$ in (13).

where $\ln \pi = (\ln \pi_1, \dots, \ln \pi_G)$, $\hat{\mathbf{M}} = (\hat{\mu}_1 \dots \hat{\mu}_G)$ and the coefficient matrix $\hat{\mathbf{B}}$ is defined as,

$$\hat{\mathbf{B}} = (\hat{\mathbf{b}}_1 \dots \hat{\mathbf{b}}_G) = \hat{\Sigma}^{-1} \hat{\mathbf{M}}. \quad (9)$$

The discriminant function in (8) is linear in $\mathbf{x}$ with coefficient matrix $\hat{\mathbf{B}} \in \mathbb{R}^{p \times G}$. Hence, if the i -th row of the coefficient matrix $\hat{\mathbf{B}}$ is a zero vector $\mathbf{0}$, then the i -th feature does not contribute to the classification rule and hence can be eliminated. If the coefficient matrix $\hat{\mathbf{B}}$ is row-sparse, then the discriminant rule achieves simultaneous feature elimination. The compressive RDA discriminant function is then defined as

$$\begin{aligned} \hat{\mathbf{d}}(\mathbf{x}) &= (\hat{d}_1(\mathbf{x}), \dots, \hat{d}_G(\mathbf{x})) \\ &= \mathbf{x}^\top \hat{\mathbf{B}} - \frac{1}{2} \text{diag}(\hat{\mathbf{M}}^\top \hat{\mathbf{B}}) + \ln \pi, \end{aligned} \quad (10)$$

where

$$\hat{\mathbf{B}} = \mathcal{H}(\hat{\mathbf{B}}) \quad (11)$$

is a sparsifying transformation of $\hat{\mathbf{B}}$. CRDA framework thus requires specifying the sparsifying transformation $\mathcal{H} : \mathbb{R}^{p \times G} \rightarrow \mathbb{R}^{p \times G}$ and the covariance matrix estimator $\hat{\Sigma}$.

The popular Shrunken Centroids Regularized Discriminant Analysis (SCRDA) [3] is then an example of CRDA method. It uses the element-wise soft-shrinkage operator $\mathcal{S}_\Delta(\cdot)$ as the sparsifying transform $\mathcal{H}(\cdot)$. In other words, it computes (11) as

$$\hat{\mathbf{B}} = \mathcal{S}_\Delta(\hat{\mathbf{B}}), \quad (12)$$

where the soft-thresholding function is applied element-wise to $\hat{\mathbf{B}}$. Furthermore, $\hat{\Sigma} = \alpha \mathbf{S} + (1 - \alpha) \mathbf{I}$, is chosen as the covariance matrix estimator. A disadvantage of this approach is that the parameter $\Delta \in [0, \infty)$ is the same across all groups, and different populations would often benefit from different shrinkage intensity. Another drawback is that an upper bound for shrinkage threshold Δ is data dependent. Two-dimensional CV search is then used to find the tuning parameter pair α and Δ .

We promote using different sparsifying transform $\mathcal{H}(\cdot)$. Namely, a row-sparse $\hat{\mathbf{B}}$ can be obtained by using hard-thresholding (HT) operator:

$$\hat{\mathbf{B}} = H_K(\hat{\mathbf{B}}, \varphi). \quad (13)$$

Above $H_K(\hat{\mathbf{B}}, \varphi)$ is defined as a transform that retains only the elements of K rows of $\hat{\mathbf{B}}$ for which HT-selector function $\varphi : \mathbb{R}^G \rightarrow [0, \infty)$ produces largest values and set the elements of other rows to zero. See Fig. 1 for an illustration. Such CRDA classifier thus uses only K features and is able to do feature selection simultaneously with classification. The tuning parameter K then determines how many features are selected. As HT-selector function φ , one may use ℓ_q -norm,

$$\varphi_q(\mathbf{b}) = \|\mathbf{b}\|_q, \quad \text{for } q \in \{1, 2, \infty\},$$

or the sample variance,

$$\varphi_0(\mathbf{b}) = s^2(\mathbf{b}) = \frac{1}{G-1} \sum_{i=1}^G |b_i - \bar{b}|^2,$$

where $\bar{b} = (1/G) \sum_{i=1}^G b_i$. Two dimensional CV search described in Section III-A can then be used for choosing an appropriate level of joint-sparsity K and the best HT-selector function.

A. CROSS-VALIDATION SCHEME

The joint-sparsity order $K \in \{1, 2, \dots, p\}$ and the HT-selector function $\varphi \in [\varphi] = \{\varphi_0, \varphi_1, \varphi_2, \varphi_\infty\}$ are chosen using a cross-validation (CV) scheme. For K we use a logarithmically spaced grid of K -values, where the upper bound for largest element in the grid is defined as follows. For each φ , we calculate

$$K_{\text{UB}}(\varphi) = \sum_{i=1}^p \mathbf{I}(\varphi(\hat{\mathbf{b}}_{i:}) \geq \text{ave}_j \{\varphi(\hat{\mathbf{b}}_{j:})\}), \quad (14)$$

where $\hat{\mathbf{b}}_{i:}$ denotes the i -th row vector of $\hat{\mathbf{B}}$. Thus $K_{\text{UB}}(\varphi_1)$ counts the number of row vectors of $\hat{\mathbf{B}}$ that has ℓ_1 -norm larger than the average of ℓ_1 -norms of row vectors of $\hat{\mathbf{B}}$. As final upper bound K_{UB} we use the minimum value, $K_{\text{UB}} = \min_{\varphi \in [\varphi]} K_{\text{UB}}(\varphi)$. We then use a logarithmically spaced grid of $J = 10$ values $[K] = \{K_1, K_2, \dots, K_J\}$, with $1 < K_1 < K_2 < \dots < K_J = K_{\text{UB}}$ starting from $K_1 = \lfloor 0.05 \cdot p \rfloor$. Note that our choice for K_{UB} adapts to the distribution of $\varphi(\hat{\mathbf{b}}_{i:})$. For example, choosing $K_{\text{UB}} < \lfloor p \cdot \frac{1}{2} \rfloor$ is recommended when the distribution of $\{\varphi(\hat{\mathbf{b}}_{i:})\}_{i=1}^p$ is such that majority of $\varphi(\hat{\mathbf{b}}_{i:})$ values are small, i.e., the empirical distribution of $\varphi(\hat{\mathbf{b}}_{i:})$ has a shape that is similar to an exponential distribution. In such cases, it is advisable to choose a value for K_{UB} that is smaller than $\lfloor p \cdot \frac{1}{2} \rfloor$. This would happen with our choice of K_{UB} with high

probability since the mean value of the exponential distribution is always larger than its median value.

Apart from the way the lower and upper bounds for the grid of K values are chosen, the used CV scheme is conventional. We split the training data set $\mathbf{X}$ into Q folds and use $Q - 1$ folds as a training set and one fold as the validation set. We fit the model (i.e., the CRDA classifier as defined by (10) and (13)) on the training set using $\varphi \in [\varphi]$ and $K \in [K]$ and compute the misclassification error rate on the validation set. This procedure is repeated so that each of the Q folds is left out in turn as a validation set, and misclassification rates are accumulated. We then choose the pair in the grid $[\varphi] \times [K]$ that yielded smallest misclassification error rate. In case of ties, we choose the one having smaller value of K , i.e., we prefer a model with smaller number of features.

B. CRDA VARIANTS

We consider the following three CRDA variants:

- **CRDA1** uses Ell1-RSCM estimator to compute the coefficient matrix $\hat{\mathbf{B}}$. The tuning parameters, K and φ , are chosen using the CV procedure of Section III-A.
- **CRDA2** is as CRDA1 except that $\hat{\mathbf{B}}$ is computed using the Ell2-RSCM estimator.
- **CRDA3** uses Rie-PSCM to compute the coefficient matrix $\hat{\mathbf{B}}$, $\varphi_\infty(\cdot) = \|\cdot\|_\infty$ is used as the HT-selector and $K = K_{UB}(\varphi_\infty)$ as the joint-sparsity level.

Uniform priors ($\pi_g = 1/G$, $g = 1, \dots, G$) are used in all analysis results.

The main computational complexity of CRDA methods are related to computing the chosen covariance matrix estimator $\hat{\Sigma}$ and the matrix inversion of $\hat{\Sigma}$. First, we note that the SVD-trick that is described in Appendix B reduce the complexity of matrix inversion from $\mathcal{O}(p^3)$ to $\mathcal{O}(pn^2)$ which is a significant saving in $p \gg n$ case. The complexity of computing the RSCM-Ell1 and RSCM-Ell2 estimators is of the same order as computing the SCM $\mathbf{S}$, which is of order $\mathcal{O}(np^2)$. Hence the complexity of CRDA methods are equivalent to complexity of computing the underlying covariance matrix estimator in high-dimensional problems.

C. ANALYSIS OF REAL DATA SETS

Next, we explore the performance of CRDA variants in classification task of three different types of data: an image data (data set #1), speech data (data set #2) and gene expression data (data set #3) described in Table 2. The first data set is constructed for classification of five English vowel letters spoken twice by thirty different subjects while the second data set consists of modified observations from the MNIST data set of handwritten digits (namely, digits 4 and 9) and used in NIPS 2003 feature selection challenge. These data sets are available at UC Irvine machine learning repository [20] under the names, Isolette and Gisette data set, respectively. The 3rd data set is from Khan *et al.* [21] and consists of small round blue cell tumors (SRBCT-s) found in children, and are classified as BL (Burkitt lymphoma), EWS (Ewing's sarcoma), NB

TABLE 1. Classification Results of CRDA Variants and Their Competitors (see Table 3) for Data Sets #1-#3 in Terms of Test Error Rate (TER) and Feature Selection Rates (FSR) Reported in Percentages

	#1 Isolet-vowels		#2 Gisette		#3 Khan <i>et al.</i>	
	TER	FSR	TER	FSR	TER	FSR
CRDA1	2.8	21.8	7.0	18.8	0	5.0
CRDA2	2.7	16.3	14.4	8.2	0	5.0
CRDA3	1.3	45.1	6.3	36.7	1.2	35.5
PLDA	3.6	100	40.5	50.8	18.8	82.0
SCRDA	1.6	88.2	6.3	87.0	2.8	81.5
varSelRF	3.0	28.9	7.9	18.8	2.4	78.8
Logit-ASPLS	26.4	98.6	27.8	90.0	43.2	100
SPCALDA	19.1	73.7	6.4	47.4	3.2	46.9
SVM	1.3	100.0	6.0	100.0	0.8	100.0
LogitBoost	3.0	100.0	7.7	100.0	10.4	100.0
RF	1.8	100.0	9.5	100.0	1.2	100.0

(neuroblastoma) and RMS (rhabdomyosarcoma). The available data is split randomly into training and test sets of fixed length, n and n_t , respectively, such that relative proportions of observations in the classes are preserved as accurately as possible. Then $L = 10$ such random splits are performed and averages of error rates are reported.

The basic metrics used are the test error rate (TER), feature selection rate (FSR) and the average computational time (ACT) in seconds. TER is the misclassification rate calculated by classifying the observations from the test set using the classification rule estimated from the training set. FSR is defined as $\text{FSR} = S/p$, where S denotes the number of features used by the classifier, and it is a measure of interpretability of the chosen model. The FSR of CRDA3 will naturally always be higher than FSR of CRDA1 or CRDA2.

The results of the study reported in Table 1 is divided in three groups. The first group consists of the proposed CRDA classification methods and the middle group consists of five state-of-the-art classifiers described in Table 3. The third group consists of three standard off-the-shelf classifiers commonly used in machine learning, but which do not perform feature selection. The methods in the third group are *RF* which refers to random forest using decision trees (where maximum number of splits in tree was 9 and number of learning cycles was 150), *LogitBoost* which refers to LogitBoost method [22] using stumps (i.e., a two terminal node decision tree) as base learners and 120 as the number of boosting iterations, and *LinSVM* which refers to a linear support vector machine (SVM) using one versus one (OVO) coding design and linear kernel. Note that since $p \gg n$, one can always find a separating hyperplane and hence a linear kernel suffices for SVM.

We first discuss the results of CRDA variants when compared against one another. In the case of data set #2, CRDA1 and CRDA3 obtained notably better misclassification rates than CRDA2. In the case of data set #3, CRDA1 and CRDA2 obtained perfect 0% misclassification results while using only 5% out of all $p = 2308$ features. This is a remarkable achievement, especially when compared to TER/FSR rates of the competing estimators. In the case of data set #1, CRDA3 had

TABLE 2. Summary of the Real Data Sets Used in This Paper

data set	p	n	n_t	$\{n_1, \dots, n_G\}$	G	Description
#1 Isolet-vowels [20]	617	100	200	{20,20,20,20,20}	5	Spoken English vowel letters
#2 Gisette [20]	5000	350	6650	{175,175}	2	OCR digit recognition ('4' vs '9')
#3 Khan <i>et al.</i> [21]	2308	38	25	{5, 14, 7, 12}	4	SRBCTs
#4 Su <i>et al.</i> [23]	5565	62	40	{15, 16, 17, 14}	4	Multiple mammalian tissues
#5 Gordon <i>et al.</i> [24]	12533	109	72	{90, 19}	2	Lung Cancer
#6 Burczynski <i>et al.</i> [25]	22283	76	51	{35, 25, 16}	3	Inflammatory bowel diseases (IBD)
#7 Yeoh <i>et al.</i> [26]	12625	148	100	{9,16, 38, 12, 26, 47}	6	Acute lymphoblastic leukemia (ALL)
#8 Tirosh <i>et al.</i> [27]	13056	114	77	{23, 22, 23, 46}	4	Yeast species
#9 Ramaswamy <i>et al.</i> [28]	16063	117	73	{7, 6, 7, 7, 13, 7, 6, 6, 18, 7, 7, 7, 7, 12}	14	Cancer

TABLE 3. State-of-the-art Methods Included in our Comparative Study. All of the Included Methods are Designed With an Aim of Performing Feature Selection During the Learning Process. Tuning Parameters are Chosen Using the Default Method Implemented in the R-packages, Including the Grids for Parameter tuning. When CV is Used, then a Five-Fold CV is Selected

Method	Reference	R-package or other	Version
PLDA	Penalized LDA [4]	penalizedLDA	1.1
SCRDA	Shrunken centroids RDA [3]	rda	1.0.2.2
varSelRF	Random forest (RF) with variable selection [30]	varSelRF	0.7.8
Logit-ASPLS	Adaptive sparse partial least squares based logistic regression [2]	plsgenomics	1.5.2
SPCALDA	Supervised principal component analysis based LDA [31]	SPCALDA	1.0
Lasso	Regularized multinomial regression with ℓ_1 (Lasso) penalty [32]	GLMNet	3.0-2
EN	Regularized multinomial regression with Elastic Net (EN) penalty [32]	GLMNet	3.0-2

better performance than CRDA1 and CRDA2 but it attained this result using more features than CRDA1 and CRDA2.

Next, we compare CRDA against the methods in the middle group. We notice from Table 1 that one of the CRDA variants is always the best performing method in terms of both accuracy (TER rates) and feature selection rates (FSR). Only in the case of data set #2, SCRDA is able to achieve equally good error rate (6.3%) as CRDA3.

Finally, we compare the CRDA methods against the methods in the third group. First we notice that the linear SVM is the best performing method among the third group, and for data set #2 it achieves the best error rate among all methods. However, it chooses *all* features.

As noted earlier, for data set #3 (Khan *et al.*), CRDA1 and CRDA2 attained perfect classification (0% error rate) while choosing only 5% out of 2308 genes. The heat map of *randomly selected* 115 genes is shown in Fig. 2(a) while Fig. 2(b) shows the heat map of 115 genes picked by CRDA1. While the former shows random patterns, clearly visible patterns can be seen in the latter heat map. Besides for prognostic classification, the set of genes found important by CRDA1 can be used in developing understanding of genes that influence SRBCT-s.

IV. FEATURE SELECTION PERFORMANCE

Next we create a partially synthetic gene expression data set using the 3rd data set in Table 2 to study how well the different classifiers are able to detect differentially expressed (DE) genes. The data set #3 consists of expression levels of $p = 2308$ genes and we randomly kept $p_1 = 115$ (i.e., $\approx 5\%$ of p) covariates of gene expression levels as significant features, i.e., differentially expressed, and the rest $p_0 = 2193$ irrelevant (non-DE) features are generated from $\mathcal{N}_n(\mathbf{0}, 0.01 \times \mathbf{I})$.

FIGURE 2. For the data set #3 (Khan *et al.*), the figure displays the heat map of (a) randomly selected 115 genes and (b) 115 genes that were selected by CRDA1 classifier. We note that the CRDA1 classifier obtained 0% error rate on all 10 Monte Carlo splits of the data to training and test set, and outperformed all other classifiers.

Since we now have a ground truth of differentially expressed genes available, we may evaluate the estimated models using the false-positive rate (FPR) and the false-negative rate (FNR). A false-positive is a feature that is picked by the classifier but is not DE in any of the classes in the sense that $\mu_{g,j} = 0$ for all $g = 1, \dots, G$. A false-negative is a feature that is DE in the true model in the sense that $\mu_{g,j} \neq 0$ for all $g = 1, \dots, G$, but is not picked by the classifier. Formally,

$$\text{FPR} = \frac{F}{p_0} \quad \text{and} \quad \text{FNR} = \frac{p_1 - T}{p_1},$$

where p_1 is the number of truly positive (DE) features, p_0 is the number of truly negative (non-DE) features, F is the number of false positives and T is the number of true positives,

FIGURE 3. Feature selection results for the partially synthetic genomic data. (a) Volcano plot illustrating that 19 (9 up-regulated and 10 down-regulated) genes are selected when comparing the third class against the others. When each class are compared in turn against the others, then 60 out of 115 DE genes are picked. (b) Plot of gene index, i , versus the $\log_{10}$ of the length of the i th row vector of the coefficient matrix $\hat{\mathbf{B}}$ produced by HT-selector φ_0 , i.e., $\text{length} = \varphi_0(\hat{\mathbf{b}}_i)$. CRDA2 has selected 107 differentially expressed genes out of total 115.

i.e., $S = F + T$ gives the total number of features picked by the classifier. Again, recall that the sparsity level K used in the hard-thresholding operator is the number S in the case of CRDA methods. Note that both FPR and FNR should be as small as possible for exact and accurate selection (and elimination) of genes.

First, we compare the gene selection performance of CRDA with the conventional gene selection approach [29] which performs a t-test and computes the p-values using permutation analysis with 10,000 permutations. A gene is considered to be differentially expressed when it shows both statistical and biological significance. Statistical significance is attested when the obtained p-value of the test statistics is below the significance level (false alarm rate), say $p_{FA} = 0.05$. Biological significance is attested when the log ratio of fold change of mean gene expressions is outside the interval $[-\log_2(c), \log_2(c)]$ for some cutoff limit c , say $c = 1.2$. Fig. 3(a) display such a volcano plot which displays the $-\log_{10}$ of p-values of t -tests for each genes against the $\log_2$ fold change. The used statistical significance level is shown as the dotted horizontal line and the biological significance bounds as dotted vertical lines.

The plot illustrates that there are 19 out of 115 differentially expressed genes which are both statistically and biologically significant (genes on the upper left and upper right corners of the plot) along with a few which are biologically significant but not statistically significant (lower left and lower right corners). When considering each class in turn against the rest, then 60 out of a total of 115 DE genes are picked as being DE. The results for CRDA2 based gene selection is shown in Fig. 3(b) for the same data set where we have plotted the gene index, i , against the sample variance, s_i , of the i th row vector of the coefficient matrix $\hat{\mathbf{B}}$. From 115 DE genes, CRDA2 has selected 107 DE genes.

As a baseline, we further compute the test error rate of the *naive classifier* that assigns all observations in the test set to the class that has the most observations in the training set. Fig. 4(a) displays the box plots of FSR and TER of

FIGURE 4. Classification accuracy results for the partially synthetic gene expression data. (a) Boxplots of test error rates (TER) and feature selection rates (FSR). (b) Boxplots of false negative rates (FNR) and false positive rates (FPR) reported in percentages. Results are from simulation study of 10 random splits of the data to training set and test set.

CRDA methods and three state-of-the-art classifiers. Again the CRDA classifiers obtained the smallest average TER while choosing the smallest number of features. Fig. 4(b) provides the box plots of FPR and FNR for all methods. The proposed CRDA methods offer the best feature selection accuracy. On the contrary, SCRDA and varSelRF methods produce large false-positive rates, meaning that they have chosen a large set of genes as DE while they are not. PLDA is performing equally well with CRDA in terms of FPR/FNR rates, but its misclassification rates are considerably larger. Furthermore, CRDA methods exhibit the smallest amount of variation (i.e., standard error) in all measures (TER/FSR/FPR/FNR). We also notice that varSelRF does not perform very well compared to CRDA or other regularized LDA techniques like SCRDA.

V. ANALYSIS OF REAL GENE EXPRESSION DATA

We compare the performance of the proposed CRDA methods to seven state-of-the-art classifiers on a diverse set of real gene expression (microarray) data sets. These data sets #4 – #9 in Table 2 constitute a representative set of gene expression data sets available in the open literature, having a varying number of features (p), sample sizes (n_g) or classes (G). As earlier, in each Monte-Carlo trial, the data is divided randomly into training sets and test sets of fixed length, n and n_t , respectively, such that relative proportions of observations in the classes are preserved as accurately as possible. We do not include CRDA3 in this study due to incurred high computational cost related to computing the Rie-PSCM covariance matrix estimator. Furthermore, since the performance of LogitBoost and Random Forest (RF) were inferior to LinSVM, we do not report the performance of these methods in the study.

The Gene Expression Omnibus (GEO) data repository contains the data sets #6 and #8 with accession numbers GDS1615 and GDS2910, respectively. The 9th data set (also known as the global cancer map (GCM) data set) is available

FIGURE 5. Test error rates (TER) and feature selection rates (FSR) in percentages, and average computational times (ACT) in seconds are reported for the last six real data sets given in Table 2. Results are averages over $L = 10$ random splits of data into training and test sets.

online [33]. The rest of the data sets, namely data sets 4, 5, and 7 are included in the R-package *datamicroarray*.¹

Fig. 5 displays the average TER and FSR values. CRDA method obtained the lowest error rates for 5 out of 6 data sets considered; only for data set #5 (Gordon), Logit-ASPLS had better performance than CRDA. Furthermore, CRDA achieved the top performance by picking a largely reduced number of genes: the number of genes selected ranged from 5.1% to 23.4%. Only Lasso and EN selected fewer number of genes than CRDA methods but their error rates were much higher in all cases except for data set #5 (Gordon).

For some data sets, such as for the seventh (Yeoh) and eighth (Tirosh), SCRDA and CRDA obtained comparable TER values, but the former fails to eliminate as many genes. Again we observe that many of the competing classifiers, such as varSelRF, PLDA or SPCALDA are performing rather poorly for some data sets; this is especially the case for the sixth (Burczynski) and ninth (Ramaswamy) data sets.

Data set #9 (Ramaswamy) presents a very difficult data set for any any classification method: it contains $G = 14$ classes, $p = 16063$ genes and only 117 observations in the training set. Results obtained by CRDA clearly illustrates its top performance compared to others in very demanding HDLSS data scenarios. CRDA classifiers outperformed most of the methods with a significant margin: CRDA gives smallest error

rates (25.2%) while at the same time reducing the number of features (genes) used in the classification procedure. The next best method is varSelRF, but it is still far behind (32.6%). We note that Logit-ASPLS method failed for this data set due to computational and memory issues and hence it is excluded.

Finally, bottom panel of Fig. 5 displays the ACT values (in seconds) of the methods. Variation in ACT between different data sets is due to different number of features (p) and observations (n) of data sets. In terms of computational complexity, CRDA loses to Lasso, EN and LinSVM. However, compared to other methods that are designed to perform feature selection, CRDA is among the most efficient method. Particularly, one can note that Logit-ASPLS and varSelRF are computationally very intensive.

VI. DISCUSSIONS AND CONCLUSIONS

We proposed a novel compressive regularized discriminant analysis (CRDA) method for high-dimensional data sets, where data dimensionality p is much larger than the sample size n (i.e., $p \gg n$, where p is very large, often several thousands).

The method showcase accurate classification results and feature selection ability while being fast to compute in real-time. Feature selection is important for post-selection inference in various problems as it provides insights on which features (e.g., genes) are particularly important in expressing a particular trait (e.g., cancer type).

¹[Online]. Available: <https://github.com/ramhiser/datamicroarray>

Our simulation studies and real data analysis examples illustrate that CRDA offers top performance in all three considered facets: classification accuracy, feature selection, and computational complexity. Software (both MATLAB and R) toolboxes to reproduce the results are available: see Section I for links to the packages.

APPENDIX A TUNING PARAMETER SELECTION OF ELL1- AND ELL2-RSCM ESTIMATOR

ELL1-RSCM and ELL2-RSCM estimators [7] computes an estimate $\hat{\alpha}$ of the optimal shrinkage parameter α_o in [7] as

$$\hat{\alpha} = \frac{(\hat{\gamma} - 1)}{(\hat{\gamma} - 1) + \hat{\kappa}(2\hat{\gamma} + p)/n + (\hat{\gamma} + p)/(n - 1)}, \quad (15)$$

where $\hat{\gamma}$ is an estimator of the sphericity parameter $\gamma = \text{ptr}(\Sigma^2)/\text{tr}(\Sigma)^2$ and $\hat{\kappa}$ is an estimate of the elliptical kurtosis parameter $\kappa = (1/3)\text{kurt}(x_i)$, calculated as the average sample kurtosis of the marginal variables and scaled by 1/3. The formula above is derived under the assumption that the data is a random sample from an unspecified elliptically symmetric distribution with mean vector μ and covariance matrix Σ .

ELL1-RSCM and ELL2-RSCM estimators differ only in the used estimate $\hat{\gamma}$ of the sphericity parameter in (15). Estimate of sphericity used by ELL1-RSCM is

$$\hat{\gamma} = \min \left(p, \max \left(1, \frac{n}{n-1} \left(\text{ptr}(\tilde{S}^2) - \frac{p}{n} \right) \right) \right) \quad (16)$$

where $\tilde{S}$ is the *sample spatial sign covariance matrix*, defined as

$$\tilde{S} = \frac{1}{n} \sum_{i=1}^n \frac{(\mathbf{x}_i - \hat{\mu})(\mathbf{x}_i - \hat{\mu})^\top}{\|\mathbf{x}_i - \hat{\mu}\|^2},$$

where $\hat{\mu} = \arg \min_{\mu} \sum_{i=1}^n \|\mathbf{x}_i - \mu\|$ is the spatial median. In ELL2-RSCM, the estimate of sphericity $\hat{\gamma}$ is computed as

$$\hat{\gamma} = \min \left(p, \max \left(1, \hat{b}_n \left(\frac{\text{ptr}(\mathbf{S}^2)}{\text{tr}(\mathbf{S})^2} - \hat{a}_n \frac{p}{n} \right) \right) \right), \quad (17)$$

where

$$\hat{a}_n = \left(\frac{n}{n + \hat{\kappa}} \right) \left(\frac{n}{n-1} + \hat{\kappa} \right)$$

$$\hat{b}_n = \frac{(\hat{\kappa} + n)(n-1)^2}{(n-2)(3\hat{\kappa}(n-1) + n(n+1))}$$

and $\mathbf{S}$ is the pooled SCM in (3).

APPENDIX B ON COMPUTING $\hat{\Sigma}^{-1}$

The matrix inversion can be performed efficiently using the SVD-trick [3]. The SVD of $\mathbf{X}$ is

$$\mathbf{X} = \mathbf{U}\mathbf{D}\mathbf{V}^\top,$$

where $\mathbf{U} \in \mathbb{R}^{p \times m}$, $\mathbf{D} \in \mathbb{R}^{m \times m}$, $\mathbf{V} \in \mathbb{R}^{n \times m}$ and $m = \text{rank}(\mathbf{X})$. Direct computation of SVD is time consuming and the trick is that $\mathbf{V}$ and $\mathbf{D}$ can be computed first from SVD of $\mathbf{X}^\top \mathbf{X} =$

$\tilde{\mathbf{V}}\tilde{\mathbf{D}}^2\tilde{\mathbf{V}}^\top$, which is only an $n \times n$ matrix. Here $\tilde{\mathbf{V}}$ is an orthogonal $n \times n$ matrix whose first m -columns are $\mathbf{V}$ and $\tilde{\mathbf{D}}$ is an $n \times n$ diagonal matrix whose upper left corner $m \times m$ matrix is $\mathbf{D}$. After we compute $\mathbf{V}$ and $\mathbf{D}$ from SVD of $\mathbf{X}^\top \mathbf{X}$, we may compute $\mathbf{U}$ from $\mathbf{X}$ by $\mathbf{U} = \mathbf{X}\mathbf{V}\mathbf{D}^{-1}$. Then, using the SVD representation of the SCM, $\mathbf{S} = (1/n)\mathbf{U}\mathbf{D}^2\mathbf{U}^\top$, and simple algebra, one obtains a simple formula for the inverse of the regularized SCM in (5) as:

$$\hat{\Sigma}^{-1} = \mathbf{U} \left[\left(\frac{\alpha}{n} \mathbf{D}^2 + (1 - \alpha)\eta \mathbf{I}_m \right)^{-1} - \frac{1}{(1 - \alpha)\eta} \mathbf{I}_m \right] \mathbf{U}^\top + \frac{1}{(1 - \alpha)\eta} \mathbf{I}_p,$$

where $\eta = \text{tr}(\mathbf{S})/p = \text{tr}(\mathbf{D}^2)/np$. This reduces the complexity from $\mathcal{O}(p^3)$ to $\mathcal{O}(pn^2)$ which is a significant saving in $p \gg n$ case.

REFERENCES

- [1] R. Tibshirani, T. Hastie, B. Narasimhan, and G. Chu, "Class prediction by nearest shrunken centroids, with applications to DNA microarrays," *Statist. Sci.*, vol. 18, no. 1, pp. 104–117, 2003.
- [2] G. Durif, L. Modolo, J. Michaelsson, J. E. Mold, S. Lambert-Lacroix, and F. Picard, "High dimensional classification with combined adaptive sparse PLS and logistic regression," *Bioinformatics*, vol. 34, no. 3, pp. 485–493, Feb. 2018.
- [3] Y. Guo, T. Hastie, and R. Tibshirani, "Regularized linear discriminant analysis and its application in microarrays," *Biostatistics*, vol. 8, no. 1, pp. 86–100, Jan. 2007.
- [4] D. M. Witten and R. Tibshirani, "Penalized classification using Fisher's linear discriminant," *J. Royal Statist. Soc., Ser. B*, vol. 73, no. 5, pp. 753–772, Nov. 2011.
- [5] A. Sharma, K. K. Paliwal, S. Imoto, and S. Miyano, "A feature selection method using improved regularized linear discriminant analysis," *Mach. Vis. Appl.*, vol. 25, no. 3, pp. 775–786, Nov. 2014.
- [6] E. Neto, F. Biessmann, H. Aurlien, H. Nordby, and T. Eichele, "Regularized linear discriminant analysis of EEG features in dementia patients," *Front Aging Neurosci.*, vol. 8, p. 273, 2016.
- [7] E. Ollila and E. Raninen, "Optimal shrinkage covariance matrix estimation under random sampling from elliptical distributions," *IEEE Trans. Signal Process.*, vol. 67, no. 10, pp. 2707–2719, May 2019.
- [8] P. L. Yu, X. Wang, and Y. Zhu, "High dimensional covariance matrix estimation by penalizing the matrix-logarithm transformed likelihood," *Comput. Statist. Data Anal.*, vol. 114, pp. 12–25, Oct. 2017.
- [9] M. Yi and D. E. Tyler, "Shrinking the covariance matrix using convex penalties on the matrix-log transformation," *J. Comput. Graphical Statist.*, to be published, doi: [10.1080/10618600.2020.1814788](https://doi.org/10.1080/10618600.2020.1814788).
- [10] M. N. Tabassum and E. Ollila, "Compressive regularized discriminant analysis of high-dimensional data with applications to microarray studies," in *Proc. IEEE Int. Conf. Acoust., Speech Signal Process.*, 2018, pp. 4204–4208.
- [11] O. Ledoit and M. Wolf, "A well-conditioned estimator for large-dimensional covariance matrices," *J. Mult. Anal.*, vol. 88, no. 2, pp. 365–411, Feb. 2004.
- [12] A. Coluccia, "Regularized covariance matrix estimation via empirical bayes," *IEEE Signal Process. Lett.*, vol. 22, no. 11, pp. 2127–2131, Nov. 2015.
- [13] B. Meriaux, C. Ren, M. N. El Korso, A. Breloy, and P. Forster, "Robust estimation of structured scatter matrices in (mis) matched models," *Signal Process.*, vol. 165, pp. 163–174, Dec. 2019.
- [14] A. Wiesel and T. Zhang, "Structured robust covariance estimation," *Found. Trends Signal Process.*, vol. 8, no. 3, pp. 127–216, 2015.
- [15] E. Ollila and D. E. Tyler, "Regularized M -estimators of scatter matrix," *IEEE Trans. Signal Process.*, vol. 62, no. 22, pp. 6059–6070, Nov. 2014.
- [16] F. Pascal, Y. Chitour, and Y. Quek, "Generalized robust shrinkage estimator and its application to STAP detection problem," *IEEE Trans. Signal Process.*, vol. 62, no. 21, pp. 5640–5651, Nov. 2014.

- [17] Y. Sun, P. Babu, and D. P. Palomar, "Regularized Tyler's scatter estimator: Existence, uniqueness, and algorithms," *IEEE Trans. Signal Process.*, vol. 62, no. 19, pp. 5143–5156, Oct. 2014.
- [18] Y. Chen, A. Wiesel, Y. C. Eldar, and A. O. Hero, "Shrinkage algorithms for MMSE covariance estimation," *IEEE Trans. Signal Process.*, vol. 58, no. 10, pp. 5016–5029, Oct. 2010.
- [19] E. Ollila, "Optimal high-dimensional shrinkage covariance estimation for elliptical distributions," in *Proc. Eur. Signal Process. Conf.*, Kos, Greece, 2017, pp. 1689–1693.
- [20] D. Dua and C. Graff, "UCI machine learning repository," 2017. [Online]. Available: <http://archive.ics.uci.edu/ml>
- [21] J. Khan *et al.*, "Classification and diagnostic prediction of cancers using gene expression profiling and artificial neural networks," *Nat. Med.*, vol. 7, no. 6, pp. 673–679, Jun. 2001.
- [22] J. Friedman, T. Hastie, and R. Tibshirani, "Additive logistic regression: A statistical view of boosting," *Ann. Statist.*, vol. 28, no. 2, pp. 337–407, Apr. 2000.
- [23] A. I. Su *et al.*, "Large-scale analysis of the human and mouse transcriptomes," *Proc. Nat. Acad. Sci. United States Amer.*, vol. 99, no. 7, pp. 4465–4470, Apr. 2002.
- [24] G. J. Gordon *et al.*, "Translation of microarray data into clinically relevant cancer diagnostic tests using gene expression ratios in lung cancer and mesothelioma," *Cancer Res.*, vol. 62, no. 17, pp. 4963–4967, Sep. 2002.
- [25] M. E. Burczynski *et al.*, "Molecular classification of Crohn's disease and ulcerative colitis patients using transcriptional profiles in peripheral blood mononuclear cells," *J. Mol. Diagn.*, vol. 8, no. 1, pp. 51–61, Feb. 2006.
- [26] E.-J. Yeoh *et al.*, "Classification, subtype discovery, and prediction of outcome in pediatric acute lymphoblastic leukemia by gene expression profiling," *Cancer Cell*, vol. 1, no. 2, pp. 133–143, Mar. 2002.
- [27] I. Tirosh, A. Weinberger, M. Carmi, and N. Barkai, "A genetic signature of interspecies variations in gene expression," *Nat. Genet.*, vol. 38, no. 7, pp. 830–834, 2006.
- [28] S. Ramaswamy *et al.*, "Multiclass cancer diagnosis using tumor gene expression signatures," *Proc. Nat. Acad. Sci.*, vol. 98, no. 26, pp. 15 149–15 154, Dec. 2001.
- [29] X. Cui and G. A. Churchill, "Statistical tests for differential expression in cDNA microarray experiments," *Genome Biol.*, vol. 4, no. 4, 2003, Art. no. 210.
- [30] R. Diaz-Uriarte, "GeneSrf and varSelRF: A web-based tool and R package for gene selection and classification using random forest," *BMC Bioinf.*, vol. 8, no. 1, Sep. 2007 Art. no. 328.
- [31] Y. Niu, N. Hao, and B. Dong, "A new reduced-rank linear discriminant analysis method and its applications," *Statistica Sinica*, vol. 28, no. 1, pp. 189–202, Jan. 2018.
- [32] J. Friedman, T. Hastie, and R. Tibshirani, "Regularization paths for generalized linear models via coordinate descent," *J. Statist. Softw.*, vol. 33, no. 1, pp. 1–22, 2010.
- [33] Cancer Program Legacy Publication Resources - Broad Institute, "Global cancer map dataset," 2001, Accessed: Sep. 18, 2018. [Online]. Available: <http://portals.broadinstitute.org/cgi-bin/cancer/publications/view/61>

MUHAMMAD NAVEED TABASSUM (Member, IEEE) received the M.Sc. degree in electrical engineering from King Saud University, Riyadh, Saudi Arabia, in 2015, and the D.Sc. (Tech.) degree in signal processing and data science from Aalto University, Espoo, Finland, in 2020. He was with multidisciplinary teams on funded research projects with NPST and RFTONICS Research Centers, King Saud University. Earlier, he developed iOS apps while working in agile software development teams with Coeus Solutions GmbH, Lahore, Pakistan. He is currently working as an R&D Engineer with Nokia, Espoo, Finland. His research interests include research and development of advanced solutions to the signal processing and data science based problems related to natural language processing, Internet of Things, wireless communication and sensor networks, and bioinformatics applications, particularly focusing statistical-, machine- and deep-learning approaches.

ESA OLLILA (Member, IEEE) received the M.Sc. degree in mathematics from the University of Oulu, Oulu, Finland, in 1998, the Ph.D. degree in statistics with honors from the University of Jyväskylä, Jyväskylä, Finland, in 2002, and the D.Sc. (Tech) degree with honors in signal processing from Aalto University, Aalto, Finland, in 2010. From 2004 to 2007, he was a Postdoctoral Fellow and from August 2010 to May 2015 an Academy Research Fellow of the Academy of Finland. He has also been a Senior Lecturer with the University of Oulu.

Since June 2015, he has been an Associate Professor of signal processing with Aalto University. He is also an Adjunct Professor (Statistics) with Oulu University. During the Fall-term 2001, he was a Visiting Researcher with the Department of Statistics, Pennsylvania State University, State College, PA, USA, whereas the academic year 2010–2011 he spent as a Visiting Postdoctoral Research Associate with the Department of Electrical Engineering, Princeton University, Princeton, NJ, USA. He is coauthor of a recent textbook *Robust Statistics for Signal Processing* (Cambridge University Press, 2018). His research interests include intersection of the fields of statistical signal processing, high-dimensional statistics, bioinformatics, and machine learning with emphasis on robust statistical methods and modeling. He is a member of EURASIP SAT in Theoretical and Methodological Trends in Signal Processing. Recently, he organized a special session on Robust Statistics in Signal Processing in CAMSAP 2019 and a special issue on *Statistical Signal Processing Solutions and Advances for Data Science in Signal Processing* (Elsevier).