

Bounds on the Number of Measurements for Reliable Compressive Classification

Hugo Reboredo, *Student Member, IEEE*, Francesco Renna, *Member, IEEE*, Robert Calderbank, *Fellow, IEEE*, and Miguel R. D. Rodrigues, *Senior Member, IEEE*

Abstract—This paper studies the classification of high-dimensional Gaussian signals from low-dimensional noisy, linear measurements. In particular, it provides upper bounds (sufficient conditions) on the number of measurements required to drive the probability of misclassification to zero in the low-noise regime, both for random measurements and designed ones. Such bounds reveal two important operational regimes that are a function of the characteristics of the source: 1) when the number of classes is less than or equal to the dimension of the space spanned by signals in each class, reliable classification is possible in the low-noise regime by using a one-vs-all measurement design; 2) when the dimension of the spaces spanned by signals in each class is lower than the number of classes, reliable classification is guaranteed in the low-noise regime by using a simple random measurement design. Simulation results both with synthetic and real data show that our analysis is sharp, in the sense that it is able to gauge the number of measurements required to drive the misclassification probability to zero in the low-noise regime.

Index Terms—Compressed sensing, compressive classification, classification, dimensionality reduction, Gaussian mixture models, measurement design, phase transitions, random measurements.

I. INTRODUCTION

COMPRESSIVE SENSING (CS) is an emerging paradigm that offers the means to simultaneously sense and compress a signal without significant loss of information [3]–[5]

Manuscript received January 14, 2016; revised May 27, 2016 and July 10, 2016; accepted July 18, 2016. Date of publication August 11, 2016; date of current version September 21, 2016. The associate editor coordinating the review of this manuscript and approving it for publication was Dr. John McAllister. This work was supported by Royal Society International Exchanges Scheme IE120996. The work of H. Reboredo was supported by the Fundação para a Ciência e Tecnologia, Portugal, under Doctoral Grant SFRH/BD/81543/2011. The work of F. Renna was carried out in part when he was in the Department of Electronic and Electrical Engineering of University College London, and was supported by the European Union's Horizon 2020 Research and Innovation Programme under the Marie Skłodowska-Curie Grant 655282. The work of R. Calderbank was supported in part by AFOSR under Award FA 9550-13-1-0076. The work of M. R. D. Rodrigues was supported by the EPSRC under Research Grant EP/K033166/1. This paper was presented in part at the 2013 IEEE International Symposium on Information Theory [1] and the 2013 IEEE Global Conference on Signal and Information Processing [2].

H. Reboredo is with the Instituto de Telecomunicações and the Departamento de Ciência de Computadores da Faculdade de Ciências da Universidade do Porto, Porto 4169-007, Portugal (e-mail: hugoreboredo@dcc.fc.up.pt).

F. Renna is with the Department of Applied Mathematics and Theoretical Physics, University of Cambridge, Cambridge CB3 0WA, U.K. (e-mail: fr330@cam.ac.uk).

R. Calderbank is with the Department of Electrical and Computer Engineering, Duke University, Durham, NC 27708 USA (e-mail: robert.calderbank@duke.edu).

M. R. D. Rodrigues is with the Department of Electronic and Electrical Engineering, University College London, London WC1E 7JE, U.K. (e-mail: m.rodrigues@ucl.ac.uk).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TSP.2016.2599496

(under appropriate conditions on the signal model and measurement process). The sensing process is based on computing the inner product of the signal of interest with a set of vectors, which are typically constituted randomly [3]–[5], and the recovery process is based on the resolution of an inverse problem. The result that has captured the imagination of the signal and information processing community is that it is possible to perfectly reconstruct an n -dimensional s -sparse signal (sparse in some orthonormal dictionary or frame) with overwhelming probability with only $\mathcal{O}(s \log(n/s))$ linear random measurements [3], [5], [6] using tractable ℓ_1 minimization methods [4] or iterative methods, like greedy matching pursuit [7].

The focus of compressive sensing has been primarily on exact or near-exact signal reconstruction from a set of linear signal measurements. However, it is also natural to leverage the paradigm to perform other relevant information processing tasks, such as detection, classification and estimation of certain parameters, from the set of compressive measurements. One could in fact argue that the paradigm is a better fit to decision support tasks such as signal detection, signal classification or pattern recognition rather than signal reconstruction, since it may be easier to discriminate between signal classes than reconstruct an entire signal using only partial information about the source signal.

This paper concentrates on the classification of signals from a set of compressive linear and noisy measurements. In particular, we consider the case where signals associated to different classes lie on low-dimensional linear subspaces. This problem is fundamental to the broad fields of signal and image processing [8]–[10], computer vision [11], [12] and machine learning [13], [14], as pre-processing often relies on dimension reduction to increase the speed and reliability of classification as well as reduce the complexity and cost of data processing and computation.

Compressive classification appears in the machine learning literature as feature extraction or supervised dimensionality reduction. For example, linear dimensionality reduction methods based on geometrical characterizations of the source have been developed, with linear discriminant analysis (LDA) [15] and principal component analysis (PCA) [15] just depending on second order statistics. In particular, LDA, which is one of the most well-known supervised dimensionality reduction methods [16], addresses simultaneously the between-class scattering and the within-class scattering of the measured data. Linear dimensionality reduction methods based on higher-order statistics of the data have therefore also been developed [14], [17]–[23]. In particular, an information-theoretic supervised dimensionality reduction inspired approach, which uses the mutual information

between the data class labels and the data measurements [14] or approximations of the mutual information via the quadratic Rényi entropy [18], [23], [24] as a criterion to linearly reduce dimensionality, have been shown to lead to state-of-the-art classification results. More recently, learning methods for linear dimensionality reduction based on nuclear norm optimization have also been proposed [25], which have been shown to lead to state-of-the-art results for face clustering, face recognition and motion segmentation applications. Low-dimensional random linear measurements have also been used in conjunction with linear classifiers in scenarios where the number of training samples is smaller than the data dimension [26]. In particular, [26] derives bounds on the generalization error of a binary Fisher linear discriminant (FLD) classifier with linear random measurements.

Compressive classification also appears in the compressive information processing literature in view of recent advances in compressive sensing [13], [27]–[33]. Reference [27] presents algorithms for signal detection, classification, estimation and filtering from random compressive measurements. References [28]–[30] and [31] study the performance of compressive detection and compressive classification for the case of random measurements. References [32] and [33] consider the problem of detection of spectral targets based on noisy incoherent projections. Reference [13] notes that a small number of random measurements captures sufficient information to allow robust face recognition. The common thread in this line of research relates to the demonstration that the detection and classification problems can be solved directly in the measurement domain, without requiring the transformation of the data from the compressive to the original data domain, i.e. without requiring the reconstruction of the data.

Other works associated with compressive classification that have arisen in the computational imaging literature, and developed under the rubric of task-specific sensing, include [8]–[10], [34]–[36]. In particular, task-specific sensing, which advocates that the sensing procedure has to be matched to the task-specific nature of the sensing application, has been shown to lead to substantial gains in performance over compressive sensing in applications such as localization [34], target detection [8], (face) recognition [9], [10], and reconstruction [35].

The majority of the contributions in the literature to date has focused on the proposal of linear measurement design algorithms for two- and multiple-class classification problems (e.g. [14], [15], [17]–[23], [37], [38]). Such algorithms – with the exception of two-class problems [37], [38] – do not typically lead to closed-form measurement designs thereby not providing a clear insight about the geometry of the measurements and preventing us to understand how classification performance behaves as a function of the number of measurements. This paper attempts to fill in this gap by asking the question:

What is the number of measurements that guarantees reliable classification in compressive classification applications?

We answer this question both for the scenario where the measurements are random and the more challenging scenario where the measurements are designed, when the distribution of the signal conditioned on the class is multivariate Gaussian with

zero mean and a certain (rank-deficient) covariance matrix. In addition, our answer to this question also leads to simple and insightful closed-form measurements designs both for two-class and multi-class classification problems.¹ Analytical bounds on the number of measurements required for reliable classification are derived in this work for the asymptotic regime of low noise. On the other hand, the validity of such predictions also for positive noise levels is showcased by numerical results.

We adopt this data model for three main reasons: first, our classification problem corresponds to the Bayesian counterpart of low-dimensional subspace classification problems that are ubiquitous in practice; then, this model often leads to state-of-the-art results in the compressive classification applications such as character and digit recognition as well as image classification [14], [23], [24]; in addition, this model – which entails that the source distribution is a Gaussian mixture model (GMM) – also relates to various well-known models in the literature including union of sub-spaces [39], [40], wavelet trees [39], [41] or manifolds [42], [43], that aim to capture additional signal structure beyond primitive sparsity in order to yield further gains. The framework based on GMM priors has also been used for the problem of signal recovery. In such case, the objective is not to determine from which Gaussian distribution the observed signal was drawn, but to reconstruct its value from compressive, noisy, linear measurements. Analytical bounds on the number of measurements needed for reliable signal reconstruction in the low-noise regime have been derived in [44], [45].

The remainder of this paper is organized as follows: Section II defines the problem, including the measurement model, source model, and performance metrics. Section III presents an upper bound to the misclassification probability and its expansion at low noise that is the basis of our analysis. Sections IV and V derive upper bounds on the number of measurements sufficient for reliable classification. In Section VI we report numerical results that validate the theoretical analysis with both synthetic data and real data from video segmentation and face recognition applications. Section VII contains a discussion on the impact of model mismatch in real data scenarios and, finally, we draw conclusions in Section VIII. The proofs of some of the results are relegated to the Appendices.

The article adopts the following notation: boldface upper-case letters denote matrices (\mathbf{X}), boldface lower-case letters denote column vectors (\mathbf{x}) and italics denote scalars (x); the context defines whether the quantities are deterministic or random. \mathbf{I}_N represents the $N \times N$ identity matrix, $\mathbf{0}_{M \times N}$ represents the $M \times N$ zero matrix (the subscripts that refer to the dimensions of such matrices will be dropped when evident from the context) and $\text{diag}(a_1, a_2, \dots, a_N)$ represents an $N \times N$ diagonal matrix with diagonal elements a_1, a_2, \dots, a_N . The operators $(\cdot)^T$, $\text{rank}(\cdot)$, $\det(\cdot)$, $\text{pdet}(\cdot)$ and $\text{tr}(\cdot)$ represent the transpose operator, the rank operator, the determinant operator,

¹Note that the problem of compressive classification of signals drawn from Gaussian distributions has been also considered in the preliminary papers [1], [2], where the behavior of the misclassification probability in the low-noise regime was studied for the case of random measurements [1] and for the case of designed measurements [2], but offering a closed-form characterization only for binary classifiers.

the pseudo-determinant operator and the trace operator, respectively. $\text{Null}(\cdot)$ and $\text{Im}(\cdot)$ denote the null space and the (column) image of a matrix, respectively, and $\dim(\cdot)$ denotes the dimension of a linear subspace. We also use the symbol $(\cdot)^\perp$ to denote the orthogonal complement of a linear space. The multivariate Gaussian distribution with mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$ is denoted by $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ and the symbol $\mathbb{P}[E]$ is used to denote the probability of the event E . $\log(\cdot)$ denotes the natural logarithm. For the sake of a compact notation, we also use the symbols $\mathcal{N}_i = \text{Null}(\boldsymbol{\Sigma}_i)$ and $\mathcal{R}_i = \text{Im}(\boldsymbol{\Sigma}_i)$, as well as $\mathcal{N}_{ij} = \text{Null}(\boldsymbol{\Sigma}_i + \boldsymbol{\Sigma}_j) = \mathcal{N}_i \cap \mathcal{N}_j$ and $\mathcal{R}_{ij} = \text{Im}(\boldsymbol{\Sigma}_i + \boldsymbol{\Sigma}_j) = \mathcal{R}_i + \mathcal{R}_j$, where $+$ denotes the sum of linear subspaces. The article also uses the symbol $[x]^+ = \max\{x, 0\}$, the floor operator $\lfloor x \rfloor$, which represents the larger integer less than or equal to x , and the little o notation where $g(x) = o(f(x))$ if $\lim_{x \rightarrow \infty} g(x)/f(x) = 0$.

II. PROBLEM STATEMENT

We consider the standard measurement model given by:

$$\mathbf{y} = \boldsymbol{\Phi}\mathbf{x} + \mathbf{n}, \quad (1)$$

where $\mathbf{y} \in \mathbb{R}^M$ represents the measurement vector, $\mathbf{x} \in \mathbb{R}^N$ represents the source vector, $\boldsymbol{\Phi} \in \mathbb{R}^{M \times N}$ represents the measurement matrix or kernel² and $\mathbf{n} \sim \mathcal{N}(\mathbf{0}, \sigma^2 \mathbf{I}) \in \mathbb{R}^M$ represents white Gaussian noise.³

We also consider that the source model is such that:

- A.1 The source class $C \in \{1, \dots, L\}$ is drawn with probability $p_i, i = 1, \dots, L$.
- A.2 The source signal conditioned on the class $C = i$ is drawn from a multivariate Gaussian distribution with zero mean and (rank deficient) covariance matrix $\boldsymbol{\Sigma}_i \in \mathbb{R}^{N \times N}$.

We should point out that we use a low-rank modeling approach even though many natural signals (e.g. patches extracted from natural images, face images, motion segmentation features, handwritten digits images, etc.) are not always low-rank but rather “approximately” low-rank [43]. The justification for the use of such low-rank modeling approach is two-fold: first, a low-rank representation is often a very good approximation to real scenarios, particularly as the eigenvalues of the class conditioned covariances often decay rapidly; second, it is then standard practice to account for the mismatch between the low-rank and the “approximately” low-rank model by adding extra noise in the measurement model in (1) (see [45]).⁴

It is assumed that the classifier – which infers the true signal class from the signal measurements using a maximum *a posteriori* (MAP) classifier – is provided with the knowledge of the true model parameters, i.e., the prior probabilities $p_i, i = 1, \dots, L$, the source covariance matrices $\boldsymbol{\Sigma}_i, i = 1, \dots, L$, the

measurement matrix $\boldsymbol{\Phi}$ and the noise variance σ^2 .⁵ In particular, the signal class estimate produced by the classifier is given by:

$$\hat{C} = \arg \max_{i \in \{1, \dots, L\}} p(C = i | \mathbf{y}) = \arg \max_{i \in \{1, \dots, L\}} p(\mathbf{y} | C = i) p_i, \quad (2)$$

where $p(C = i | \mathbf{y})$ is the *a posteriori* probability of class $C = i$ given the measurement vector \mathbf{y} and $p(\mathbf{y} | C = i)$ represents the probability density function (pdf) of the measurement vector \mathbf{y} given the class $C = i$, which is zero-mean Gaussian, with covariance matrix $\boldsymbol{\Phi}\boldsymbol{\Sigma}_i\boldsymbol{\Phi}^T + \mathbf{I}\sigma^2$.

Our objective is to characterize the number of measurements sufficient for reliable classification in the asymptotic limit of low noise, i.e. such that

$$\lim_{\sigma^2 \rightarrow 0} P_e = 0, \quad (3)$$

where P_e is the misclassification probability of the MAP classifier. Note also that the asymptotic regime of low-noise plays a fundamental role in various signal and image processing scenarios, e.g., digit recognition and satellite data classification [14]. In particular, by using the law of total probability, we can write

$$P_e = \mathbb{P}[\hat{C} \neq C] = \sum_{i=1}^L p_i \mathbb{P}[\hat{C} \neq i | C = i] \quad (4)$$

$$= \sum_{i=1}^L p_i \int_{\mathbb{R}^M \setminus \mathcal{D}_i} p(\mathbf{y} | C = i) d\mathbf{y}, \quad (5)$$

where \mathcal{D}_i is the decision region associated to class i , that is, the set of values \mathbf{y} corresponding to the output $\hat{C} = i$. Moreover, we can express the set $\mathbb{R}^M \setminus \mathcal{D}_i$ in terms of the unit step function $u(\cdot)$ and, by leveraging the definition of the MAP classifier in (2), we can write the misclassification probability as

$$P_e = \sum_{i=1}^L p_i \int_{-\infty}^{+\infty} p(\mathbf{y} | C = i) u \left(\max_{\substack{j \\ j \neq i}} \log \frac{p_j p(\mathbf{y} | C = j)}{p_i p(\mathbf{y} | C = i)} \right) d\mathbf{y}. \quad (6)$$

The low-noise characterization of the misclassification probability will be carried out both for random measurements and designed measurements.

Our characterization will also be based on the following additional assumptions:

- A.3 The linear spaces $\mathcal{R}_i = \text{Im}(\boldsymbol{\Sigma}_i), i = 1, \dots, L$ are of equal dimension, i.e. $\dim(\mathcal{R}_i) = r_\Sigma < N, i = 1, \dots, L$;⁶
- A.4 The linear spaces \mathcal{R}_i are independently drawn from a continuous pdf over the Grassmann manifold of subspaces of dimension r_Σ in \mathbb{R}^N , so that the null spaces $\mathcal{N}_i = \text{Null}(\boldsymbol{\Sigma}_i)$ are also of equal dimension, i.e. $\dim(\mathcal{N}_i) = N - r_\Sigma$, and are also drawn independently

²We refer to $\boldsymbol{\Phi}$ as the measurement or sensing matrix/kernel interchangeably throughout the paper.

³The results presented in the remainder of the paper can be easily generalized to the case when the noise covariance matrix is a positive definite matrix $\boldsymbol{\Sigma}_n$.

⁴We also note that our analysis focuses on zero-mean models, since various datasets (e.g., face images and motion segmentation features) can be well represented via zero-mean classes [43], [46]. However, some of the results in the paper can also be generalized to the case of nonzero-mean classes.

⁵Although our analysis assumes that the classifier is given the true distributions, in Section VI, we also conduct experiments with real datasets to assess scenarios where the classifier does not know the true distributions but rather approximate ones, that are learnt from training data.

⁶Many of the results presented in this work naturally generalize to the case when the linear spaces spanned by signals in different classes have different dimensions.

from a continuous pdf over the Grassmann manifold of subspaces of dimension $N - r_{\Sigma}$ in \mathbb{R}^N .⁷

The assumptions A.3 and A.4 imply that with probability 1

$$\dim(\mathcal{R}_{ij}) = \dim(\mathcal{R}_i + \mathcal{R}_j) = \min\{N, 2r_{\Sigma}\}, \quad (7)$$

and

$$\dim(\mathcal{N}_{ij}) = \dim(\mathcal{N}_i \cap \mathcal{N}_j) = [N - 2r_{\Sigma}]^+. \quad (8)$$

Our characterization will also use the quantities:

$$R = \dim(\mathcal{R}_i) + \dim(\mathcal{R}_j) - 2 \dim(\mathcal{R}_i \cap \mathcal{R}_j) \quad (9)$$

$$= 2 \dim(\mathcal{R}_{ij}) - \dim(\mathcal{R}_i) - \dim(\mathcal{R}_j) \quad (10)$$

$$= 2 \min\{N - r_{\Sigma}, r_{\Sigma}\}, \quad (11)$$

that relates to the difference between the dimension of the subspaces spanned by source signals in classes i or j and the dimension of the intersection of such sub-spaces;

$$r_i = \text{rank}(\Phi \Sigma_i \Phi^T) \quad (12)$$

$$v_i = \text{pdet}(\Phi \Sigma_i \Phi^T), \quad (13)$$

which measure the dimension of the sub-space spanned by the linear transformation of the signals in class i and the volume occupied by those signals in \mathbb{R}^M , respectively, and

$$r_{ij} = \text{rank}(\Phi(\Sigma_i + \Sigma_j)\Phi^T) \quad (14)$$

$$v_{ij} = \text{pdet}(\Phi(\Sigma_i + \Sigma_j)\Phi^T), \quad (15)$$

which measure the dimension of the direct sum of sub-spaces spanned by the linear transformation of the signals in classes i or j and the volume occupied by the measured signals from classes i and j in \mathbb{R}^M , respectively.

III. MISCLASSIFICATION PROBABILITY, BOUNDS AND EXPANSIONS

The basis of our characterization of an upper bound to the number of random or designed measurements sufficient for reliable classification is an asymptotic expansion of an upper bound to the misclassification probability of the MAP classifier in (6). We work with an upper bound to the misclassification probability *in lieu* of the true misclassification probability, in view of the lack of closed-form expressions for the misclassification probability of the MAP classifier.

In particular, the Bhattacharyya bound [31] represents an upper bound to the misclassification probability associated to the binary MAP classifier which is based on the inequality $\min\{a, b\} \leq \sqrt{ab}$, for $a, b > 0$. Then, it is possible to establish, by using the union-bound in conjunction with the Bhattacharyya bound, that the misclassification probability of the MAP classifier can be upper bounded as follows:

$$\bar{P}_e = \sum_{i=1}^L \sum_{\substack{j=1 \\ j \neq i}}^L \sqrt{p_i p_j} e^{-K_{ij}}, \quad (16)$$

⁷Note that this assumption on the linear spaces occupied by signals in different classes reflects well the behavior of many real data ensembles for various applications as face recognition, video motion segmentation, or digits classification [14], [46].

where

$$K_{ij} = \frac{1}{4} \log \frac{\left(\det \left(\frac{\Phi(\Sigma_i + \Sigma_j)\Phi^T + 2\sigma^2 \mathbf{I}}{2} \right) \right)^2}{\det(\Phi \Sigma_i \Phi^T + \sigma^2 \mathbf{I}) \det(\Phi \Sigma_j \Phi^T + \sigma^2 \mathbf{I})}. \quad (17)$$

Note that the exponent K_{ij} is a function of the ratio between the volume collectively occupied by measured signals belonging to classes i and j and the product of the volumes occupied distinctly by measured signals in class i and measured signals in class j .

The following lemma now provides the low-noise expansion of the upper bound to the probability of error. It defines the probability of error using two quantities: one quantity characterizes the slope of the decay of the upper bound to the misclassification probability (in a $\log \sigma^2$ scale) and the other quantity defines the power offset of the upper bound to the misclassification probability at low-noise levels.

Lemma 1: Consider the measurement model in (1) and the assumptions A.1, A.2 in Section II. Then, in the regime of low noise where $\sigma^2 \rightarrow 0$, the upper bound to the probability of misclassification can be expanded as:

$$\bar{P}_e = g (\sigma^2)^d + o\left((\sigma^2)^d\right), \quad (18)$$

where

$$d = \min_{\substack{i,j \\ j \neq i}} d(i, j), \quad d(i, j) = (2r_{ij} - r_i - r_j) / 4 \quad (19)$$

and

$$g = \sum_{(i,j) \in \mathcal{S}_d} \sqrt{p_i p_j} 2^{r_{ij}/2} \left[\frac{\sqrt{v_i v_j}}{v_{ij}} \right]^{1/2} \quad (20)$$

where $\mathcal{S}_d = \{(i, j) : i \neq j, d(i, j) = d\}$.

Proof: See Appendix A. ■

This lemma leads immediately to the following corollary that provides conditions for $\lim_{\sigma^2 \rightarrow 0} \bar{P}_e = 0$ and hence conditions for $\lim_{\sigma^2 \rightarrow 0} P_e = 0$.

Corollary 1: Consider the measurement model in (1) and the assumptions A.1, A.2 in Section II. We have that

$$\exists(i, j), i \neq j : r_i + r_j = 2r_{ij} \Rightarrow \lim_{\sigma^2 \rightarrow 0} \bar{P}_e = g > 0 \quad (21)$$

and

$$r_i + r_j < 2r_{ij}, \quad \forall(i, j), i \neq j \Rightarrow \lim_{\sigma^2 \rightarrow 0} \bar{P}_e = 0. \quad (22)$$

The conditions that guarantee that $\lim_{\sigma^2 \rightarrow 0} \bar{P}_e = 0$ stem directly from conditions that guarantee $d > 0$. The ensuing analysis then concentrates on how to define the effect of the number of random measurements or designed measurements on the value of the exponent d as a proxy to characterize the phase transition in (3).

IV. RANDOM MEASUREMENTS

We first consider the simpler problem where the measurement matrix Φ is random. In particular, we consider that the

measurement matrix is randomly drawn from a left rotation-invariant distribution.⁸

We consider the following problem:

Determine the minimum number of random measurements needed to guarantee that

$$\lim_{\sigma^2 \rightarrow 0} -\frac{\log P_e}{\log(1/\sigma^2)} > d_0. \quad (23)$$

The following proposition provides a solution for the case $d_0 = 0$ that leads precisely to the minimum number of measurements for $\lim_{\sigma^2 \rightarrow 0} \bar{P}_e = 0$ hence an upper bound on the minimum number of measurements for $\lim_{\sigma^2 \rightarrow 0} P_e = 0$.

Proposition 1: Consider the measurement model in (1) and the assumptions A.1–A.4 in Section II. Then, an upper bound on the minimum number of measurements for

$$\lim_{\sigma^2 \rightarrow 0} P_e = 0 \quad (24)$$

is

$$M = r_\Sigma + 1. \quad (25)$$

Proof: The proof of this proposition follows immediately from the characterization in Corollary 1 and from the observation that, with probability 1 over the distribution of Φ , $r_i = \min\{M, r_\Sigma\}$ and $r_{ij} = \min\{M, N, 2r_\Sigma\}$. ■

The following proposition provides a generalization of this result from the case $d_0 = 0$ to $d_0 > 0$.

Proposition 2: Consider the measurement model in (1) and the assumptions A.1–A.4 in Section II. Then, if $d_0 < R/4$, an upper bound on the minimum number of measurements for

$$\lim_{\sigma^2 \rightarrow 0} -\frac{\log P_e}{\log(1/\sigma^2)} > d_0 \quad (26)$$

is

$$M = \lfloor 2d_0 + r_\Sigma \rfloor + 1. \quad (27)$$

Proof: The proof of this proposition follows immediately from the characterization of the exponent d in Lemma 1 and from the observation that, with probability 1 over the distribution of Φ , $r_i = \min\{M, r_\Sigma\}$ and $r_{ij} = \min\{M, N, 2r_\Sigma\}$. ■

We note that the result in Proposition 1 implies that reliable classification with random measurements is obtained when the signals are embedded into a linear space with dimension strictly greater than the dimension of the spaces spanned by the class conditioned input signals, i.e., r_Σ ; in fact, when this is not the case, the measured signals occupy the entire space \mathbb{R}^M and, therefore, they are not distinguishable with arbitrarily low misclassification probability when $\sigma^2 \rightarrow 0$.

On the other hand, the results in Proposition 2 unveil the interplay between the decay rate of the upper bound to the misclassification probability, the measurements and the geometry of the source. In particular, the results imply that the decay rate

⁸A random matrix $\mathbf{A} \in \mathbb{R}^{m \times n}$ is said to be (left or right) rotation-invariant if the joint pdf of its entries $p(\mathbf{A})$ satisfies $p(\Theta\mathbf{A}) = p(\mathbf{A})$, or $p(\mathbf{A}\Psi) = p(\mathbf{A})$, respectively, for any orthogonal matrix Θ or Ψ . A special case of (left and right) rotation-invariant random matrices is represented by matrices with independent identically distributed (i.i.d.), zero-mean Gaussian entries with fixed variance, which is common in the CS literature [3], [5].

scales linearly with the number of measurements up to the maximum decay rate associated with the upper bound in (16), i.e., $R/4$, which is achieved when signals are embedded into a linear space with dimension equal to $\dim(\mathcal{R}_{ij}) = \min\{N, 2r_\Sigma\}$, i.e., the dimension of the sum of any pair of spaces spanned by signals in a given class.

V. DESIGNED MEASUREMENTS

We now consider the more challenging problem where the measurement matrix Φ is designed. In particular, we also want to consider the following problem:

Determine the minimum number of designed measurements needed to guarantee that

$$\lim_{\sigma^2 \rightarrow 0} -\frac{\log P_e}{\log(1/\sigma^2)} > d_0. \quad (28)$$

Note once again that by setting $d_0 = 0$ one obtains an upper bound to the minimum number of measurements for $\lim_{\sigma^2 \rightarrow 0} P_e = 0$, thereby guaranteeing a phase transition in the misclassification probability; and by setting $d_0 > 0$ one obtains an upper bound to the minimum number of measurements for $\lim_{\sigma^2 \rightarrow 0} -\frac{\log P_e}{\log(1/\sigma^2)} > d_0$, thereby guaranteeing a certain decay in the misclassification probability.

We will consider separately the case of two classes and the multiple classes scenario.

A. Two Classes

The following propositions provide an upper bound to the minimum number of measurements required to drive the misclassification probability to zero at a rate higher than a given value d_0 .

Proposition 3: Consider the measurement model in (1) where the assumptions A.1–A.4 in Section II are verified and $L = 2$. Then, an upper bound on the minimum number of measurements for

$$\lim_{\sigma^2 \rightarrow 0} P_e = 0 \quad (29)$$

is

$$M = 1, \quad (30)$$

and a possible measurement matrix that achieves (29) is obtained by choosing $\Phi = \phi^T$, where $\phi \in \mathbb{R}^{N \times 1}$ is a vector in \mathcal{N}_1 or \mathcal{N}_2 that is not contained in the intersection $\mathcal{N}_1 \cap \mathcal{N}_2$.

Proof: The proof of this proposition follows immediately from the evaluation of the expansion exponent d of the upper bound (16). Namely, when $\Phi = \phi^T$, where $\phi \in \mathbb{R}^{N \times 1}$ is a vector in \mathcal{N}_1 or \mathcal{N}_2 that is not contained in the intersection $\mathcal{N}_1 \cap \mathcal{N}_2$, we obtain $d = (2r_{12} - r_1 - r_2)/4 = 1/4 > 0$, which immediately implies (29). Note also that the existence of the vector ϕ is guaranteed by the fact that, if $r_\Sigma < N$, then $\mathcal{R}_1 \neq \mathcal{R}_2$ and, therefore, $\mathcal{N}_1 \neq \mathcal{N}_2$. ■

Proposition 4: Consider the measurement model in (1) where the assumptions A.1–A.4 in Section II are verified and $L = 2$. Then, if $d_0 < R/4$, an upper bound on the minimum

number of measurements for

$$\lim_{\sigma^2 \rightarrow 0} -\frac{\log P_e}{\log(1/\sigma^2)} > d_0 \quad (31)$$

is

$$M = \lfloor 4d_0 \rfloor + 1, \quad (32)$$

and a measurement matrix Φ that achieves (31) is obtained by choosing arbitrarily $\lfloor 4d_0 \rfloor + 1$ out of the R rows of matrix

$$\Phi_0 = [\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_{n_\Sigma}, \mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_{n_\Sigma}]^T, \quad (33)$$

where the sets $[\mathbf{u}_1, \dots, \mathbf{u}_{n_{12}}]$, $[\mathbf{u}_1, \dots, \mathbf{u}_{n_{12}}, \mathbf{v}_1, \dots, \mathbf{v}_{n_\Sigma}]$, $[\mathbf{u}_1, \dots, \mathbf{u}_{n_{12}}, \mathbf{w}_1, \dots, \mathbf{w}_{n_\Sigma}]$, $\mathbf{u}_i, \mathbf{v}_i, \mathbf{w}_i \in \mathbb{R}^N$, constitute an orthonormal basis of the linear spaces $\mathcal{N}_{12}, \mathcal{N}_1$ and \mathcal{N}_2 , respectively, and $n_{12} = \lfloor N - 2r_\Sigma \rfloor^+$, $n_\Sigma = \min\{N - r_\Sigma, r_\Sigma\} = R/2$.

Proof: See Appendix B. ■

We can observe that a designed kernel can offer marked improvements over a random one in the low-noise regime. Namely, perfect separation of the measured signals can be achieved with a single measurement – with a random measurement kernel we require $M \geq r_\Sigma + 1$ – and the maximum decay exponent d associated with the upper bound (16), i.e., $R/4$, is achieved with $M = R$ – with a random measurement kernel we require $M = \min\{N, 2r_\Sigma\} \geq R$.

We also observe that the kernel design embedded in Proposition 4 relates to previous results in the literature about measurement kernel optimization for the 2-classes classification problem. In particular, for the case of zero-mean classes, it was shown in [37] that the measurement kernel minimizing the Bhattacharyya bound of the misclassification probability for two zero-mean classes is obtained via the eigenvalue decomposition of the matrix $\Sigma_1^{-1}\Sigma_2$, where the covariance matrices Σ_1 and Σ_2 are assumed to be full rank.

A generalization of this construction for the case when Σ_1 and Σ_2 are not invertible is presented in [38]. Such kernel design leverages the generalized singular value decomposition (GSVD) [47] of the pair of matrices (Σ_1, Σ_2) in order to minimize the corresponding Bhattacharyya upper bound. In particular, it is shown that the most discriminant measurements are those corresponding to generalized eigenvectors which lie in the intersections $\mathcal{R}_1 \cap \mathcal{N}_2$ or $\mathcal{R}_2 \cap \mathcal{N}_1$. Then, on recalling that $\mathcal{R}_i = \mathcal{N}_i^\perp$, we can note that the most discriminant measurements are picked from a subspace contained in $\mathcal{N}_1(\mathcal{N}_2)$ that is also orthogonal to (and therefore, not contained in) $\mathcal{N}_2(\mathcal{N}_1)$. In this sense, the construction described by Proposition 4 is similar to this result. However, there are significant differences between our results and the results in [38]. First, our analysis applies to a sensing scenario in *lieu* of feature extraction; so the measurements in (1) are contaminated by noise whereas the measurements in [38] are not. More importantly, the analysis in [38] does not offer an explicit characterization of the number of measurements needed to guarantee a given misclassification probability performance. On the other hand, our analysis offers sufficient conditions for reliable classification in the low-noise regime and a direct connection between the number of measurements taken on the source signal and the low-noise behavior of the

corresponding upper bound to the misclassification probability via the exponent d .

B. Multiple Classes

The following propositions offer an upper bound to the minimum number of measurements required to drive the misclassification probability to zero, and a procedure to determine an upper bound to the minimum number of measurements required to guarantee that the misclassification probability decays to zero with an exponent higher than a given value d_0 .

Proposition 5: Consider the measurement model in (1) and the assumptions A.1–A.4 in Section II. Then, an upper bound on the minimum number of measurements for

$$\lim_{\sigma^2 \rightarrow 0} P_e = 0 \quad (34)$$

is

$$M = \min\{L - 1, r_\Sigma + 1\}. \quad (35)$$

Moreover, a measurement matrix Φ that achieves (34) is obtained as follows: let \mathbf{N}_i be a matrix that contains a basis for the null space \mathcal{N}_i . Then, the $M = \min\{L - 1, r_\Sigma + 1\}$ rows of the matrix Φ are obtained by randomly picking one row from each of the matrices $\mathbf{N}_{\pi(1)}^T, \dots, \mathbf{N}_{\pi(\min\{L-1, r_\Sigma+1\})}^T$, where $\pi(\cdot)$ is any permutation function of the integers $1, \dots, L$.

Proof: See Appendix C. ■

Note that the characterization embodied in Proposition 5 is obtained by taking the measurement matrix to belong to a certain restricted subset of $\mathbb{R}^{M \times N}$ rather than the entire $\mathbb{R}^{M \times N}$.

The choice of such subset of $\mathbb{R}^{M \times N}$ is inspired by our characterization pertaining to the two-class problem embodied in Propositions 3 and 4. Namely, let $\mathbf{N}_i \in \mathbb{R}^{N \times (N-r_\Sigma)}$ be a matrix that contains a basis for the null space \mathcal{N}_i and let $\mathbf{N} = [\mathbf{N}_1, \dots, \mathbf{N}_L]$ be a matrix that contains the concatenation of the bases for all the null spaces $\mathcal{N}_1, \dots, \mathcal{N}_L$. Then, we take the measurement matrix to consist of M rows of \mathbf{N}^T rather than M arbitrary vectors from \mathbb{R}^N .

Note also that the result embodied in Proposition 5 – which is shown to be very sharp both with synthetic data and real data simulations – provides a fundamental insight in the role of measurement design in comparison with random measurement kernels in the discrimination of subspaces. In particular, we can clearly identify two operational regimes that depend on the relationship between two fundamental geometrical parameters describing the source: the number of classes and the dimension of the linear subspaces associated to the different classes.

- When, the number of classes in the source is lower than or equal to the dimension of the spaces spanned by signals in each class, the designed measurement matrix is such that we take one measurement from $L - 1$ out of the L null spaces $\mathcal{N}_i, i = 1, \dots, L$. In this sense, the construction that achieves the upper bound implements a one-vs-all approach, where each measurement is able to perfectly detect the presence of signals coming from a specific class against signals from all the remaining classes. Note that in this regime, proper design of the measurement kernel can provide a dramatic performance advantage

with respect to random measurements, as it can guarantee that the misclassification probability approaches zero, in the low-noise regime, even when random measurements yield an error floor.

- On the other hand, when the number of classes is larger than the dimension spanned by signals in a given class, (more precisely, when $L > r_{\Sigma} + 1$), then $r_{\Sigma} + 1$ measurements are sufficient to drive to zero the misclassification probability in the low-noise regime. In this case, the designed measurement kernel obtains the same performance of random measurements in terms of phase transition of upper bounds to the misclassification probability. However, properly designing the measurement kernel can have an impact on the value of the error floor or the speed of the decay of the misclassification probability with $1/\sigma^2$.

Proposition 6: Consider the measurement model in (1) and the assumptions A.1–A.4 in Section II. Then, if $d_0 < R/4$, an upper bound on the minimum number of measurements for

$$\lim_{\sigma^2 \rightarrow 0} -\frac{\log P_e}{\log(1/\sigma^2)} > d_0 \quad (36)$$

is given by the solution to the integer programming problem

$$\begin{aligned} & \underset{(M_1, \dots, M_L) \in \mathbb{N}^L}{\text{minimize}} && M = \sum_{i=1}^L M_i \\ & \text{subject to:} && M_i \leq N - r_{\Sigma}, \forall i \\ & && f(M, M_i, M_j) - 2(M - 2r_{\Sigma}) > d_0, \forall i \neq j \\ & && f(M, M_i, M_j) - 2(M_i - r_{\Sigma}) > d_0, \forall i \neq j \\ & && f(M, M_i, M_j) - 2(M_j - r_{\Sigma}) > d_0, \forall i \neq j \\ & && f(M, M_i, M_j) > d_0, \quad \forall i \neq j, \end{aligned} \quad (37)$$

where $f(M, M_i, M_j) = \max\{M - r_{\Sigma}, M_i\} + \max\{M - r_{\Sigma}, M_j\}$.

Proof: Note that $R/4$ is the maximum decay exponent d associated with the upper bound in (16), and note also that, if $d_0 < R/4$, a sufficient condition for (36) is given by $d(i, j) > d_0$, for all $(i, j), i \neq j$. Then, the proposed upper bound follows from taking the measurement matrix to belong to the same restricted subset considered in Proposition 5. In this case, on denoting by M_i the number of measurements in Φ that are also columns of \mathbf{N}_i , so that $M = \sum_{i=1}^L M_i$, we can write⁹

$$\begin{aligned} d(i, j) = f(M, M_i, M_j) \\ - 2 \max\{M - 2r_{\Sigma}, M_i - r_{\Sigma}, M_j - r_{\Sigma}, 0\}, \end{aligned} \quad (38)$$

which leads to the formulation of the problem (37). ■

Note that, although a general closed-form solution to the optimization problem in (37) is difficult to provide, our formulation allows to drastically reduce the number of (integer) optimization variables, which is now equal to the number of classes L . Moreover, the integer programming problem in (37) involves a linear objective function and constraints that are

expressed via linear functions combined via the max function, thus allowing the use of efficient numerical methods for its solution.

It is also important to emphasize the differences between the result in Proposition 5 and other results in the literature. The result in (35) is reminiscent of a result associated to multiclass LDA, that involves the extraction of $L - 1$ linear features from the data using the LDA rule [15]. However, such LDA construction does not provide conditions on the number of measurements needed for reliable classification. Moreover, in contrast with the analysis here proposed, LDA approaches are usually applied to the nonzero-mean classes scenario rather than the zero-mean case considered here. In fact, LDA methods are shown to be ineffective in the case of zero-mean classes, due to the measurement kernel construction approach that is based on the computation of the GSVD of inter-class and intra-class scatter matrices, where the first one is a function of the class means.

A modified version of LDA which can cope also with zero-mean classes has been presented in [38]. Such method is based on recasting a multiclass classification problem into a binary pattern classification problem. However, in this case the measurement kernel Φ is not determined on the basis of the statistical description of the classes, but rather it is derived via a non-parametric approach, which involves the computation of scatter matrices from labeled training samples. In particular, on denoting by Σ_b the between-class scatter matrix and by Σ_w the within-class scatter matrix, measurements are designed in order to maximize the objective function

$$J(\Phi) = \text{tr}((\Phi \Sigma_w \Phi^T)^{-1} (\Phi \Sigma_b \Phi^T)), \quad (39)$$

leading to measurement designs that are associated with the generalized eigenvectors corresponding to the largest generalized eigenvalues of (Σ_b, Σ_w) . In addition, in this case, conditions on number of measurements needed for reliable classification are not available in general.

VI. NUMERICAL RESULTS

We now show how our theory aligns with practice, both for synthetic data and real data associated with a video segmentation application and with a face recognition application. We also show how our upper bound on the minimum number of measurements required for the phase transition compares to those associated with state-of-the-art measurement designs such as information discriminant analysis (IDA) methods [21] and methods based on the maximization of Shannon mutual information and quadratic Rényi entropy [14].

A. Synthetic Data

We first consider experiments with synthetic data by concentrating on two examples that reflect the two regimes embodied in Proposition 5. In the first example, the data is generated by a mixture of $L = 11$ Gaussian distributions with dimension $N = 64$, with probability $p_i = 1/11$, for $i = 1, \dots, 11$. The input covariance matrices have all rank $r_{\Sigma} = 14$, and their images

⁹The details are provided in Appendix C.

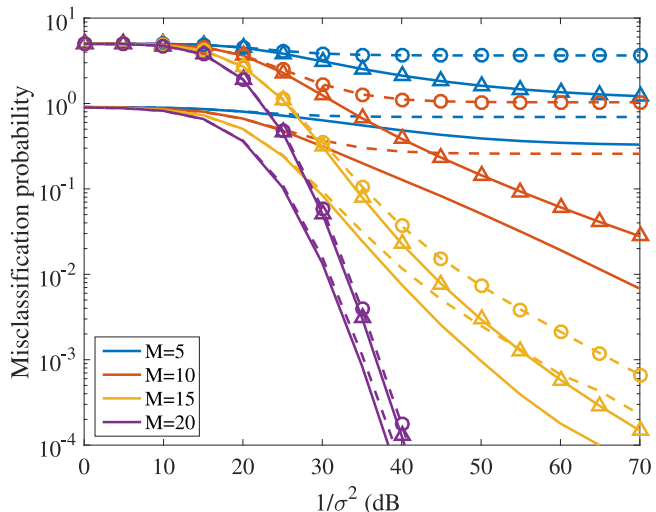


Fig. 1. Upper bound and true misclassification probability vs. $1/\sigma^2$. $N = 64$, $L = 11$, $r_\Sigma = 14$. True misclassification probability with random measurement kernels (dashed lines) and designed kernels (solid lines). Upper bound to the misclassification probability with random measurement kernels (dashed lines with circles) and designed kernels (solid lines with triangles).

are drawn uniformly at random from the Grassmann manifold of 14-dimensional spaces in \mathbb{R}^{64} .

Fig. 1 reports the upper bound to the misclassification probability and the true misclassification probability, respectively, vs $1/\sigma^2$ both for random kernel designs and measurement designs that obey the construction embodied in Proposition 5.¹⁰ The measurement kernels are also normalized such that $\text{tr}(\Phi^T \Phi) \leq M$.

Note that theoretical results are aligned with experimental results in the sense that both theory and practice suggest that the low-noise phase transition occurs with $M \geq L - 1 = 10$ for designed kernels and $M \geq r_\Sigma + 1 = 15$ for random kernels. This is observed from Fig. 1, suggesting that our analysis is sharp.

In the second example, the data is drawn from a mixture of $L = 12$ Gaussian distributions with dimension $N = 64$, with probability $p_i = 1/12$ for $i = 1, \dots, 12$. The input covariance matrices have all rank $r_\Sigma = 9$, and their images are drawn uniformly at random from the Grassmann manifold of 9-dimensional spaces in \mathbb{R}^{64} .

Fig. 2 showcases the upper bound to the misclassification probability and the true misclassification probability, respectively, vs $1/\sigma^2$ both for random kernel designs and measurement designs that obey the construction embodied in Proposition 5. It is evident – as predicted by Proposition 5 – that both random and designed kernels achieve a low-noise phase transition in the upper bound to the misclassification probability with $M \geq r_\Sigma + 1 = 10$. However, designed kernels offer a lower misclassification probability than random kernels for finite noise levels. It is also evident by comparing the true mis-

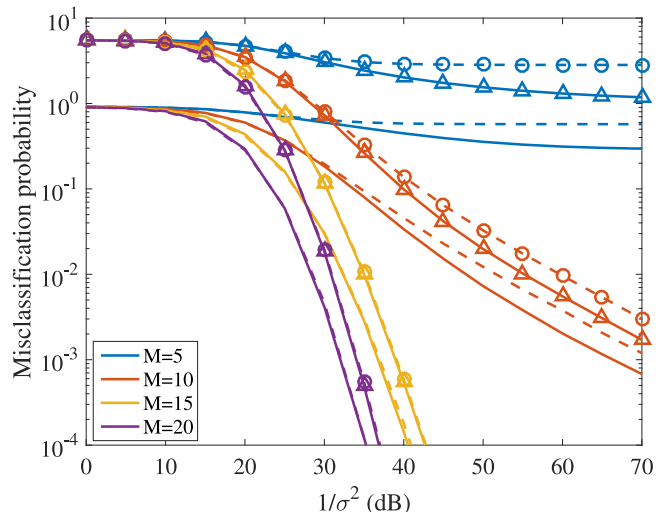


Fig. 2. Upper bound and true misclassification probability vs. $1/\sigma^2$. $N = 64$, $L = 12$, $r_\Sigma = 9$. True misclassification probability with random measurement kernels (dashed lines) and designed kernels (solid lines). Upper bound to the misclassification probability with random measurement kernels (dashed lines with circles) and designed kernels (solid lines with triangles).

TABLE I
MINIMUM NUMBER OF MEASUREMENTS M REQUIRED TO ACHIEVE THE LOW-NOISE PHASE TRANSITION OF THE MISCLASSIFICATION PROBABILITY

	random	design (35)	IDA	MI	Rényi
case i)	15	10	10	10	10
case ii)	10	10	10	10	10

classification probability values and the upper bounds in Fig. 2 that our analysis is sharp.

Our upper bound to the minimum number of measurements required for the phase transition of the misclassification probability relied on a specific construction. It is therefore relevant to examine how such a bound compares to the number of measurements required for the phase transition associated with state-of-the-art kernel designs. To that end, we consider three state-of-the-art measurement kernel designs applied to the two previous examples: these are the IDA method in [21] and methods based on the maximization of Shannon mutual information (MI) and Rényi quadratic entropy [14], respectively.

Table I reports the minimum number of measurements needed by such methods in order to drive to zero the numerically simulated misclassification probability, as well as the theoretical predictions derived in the previous sections for both random and designed kernel. It is interesting to see that the bound embodied in Proposition 5 predicts very well the behavior of state-of-the-art kernel design methods. This means that our bound can be used to gauge a suitable number of measurements to be used in state-of-the-art kernel design approaches.

B. Real Data: Motion Segmentation

We now consider experiments with real data by concentrating on a motion segmentation application, where the goal is to

¹⁰Note that the construction embodied in Proposition 5 is shown to achieve the low-noise phase transition with a number of measurements equal to (35).

TABLE II
EIGENVALUES OF THE INPUT COVARIANCE MATRICES OBTAINED FROM
TRAINING SAMPLES FROM THE VIDEO "1RT2RCR" VIA THE ML ESTIMATOR.
LARGEST FIVE EIGENVALUES FOR EACH CLASS

Σ_1	9.6284	2.2694	0.0194	0.0061	0.003
Σ_2	3.1756	0.7410	0.0267	0.0022	0.000
Σ_3	11.2797	5.9315	0.0672	0.0004	0.000

segment a video in multiple rigidly moving objects. Such application involves the extraction of feature points from the video whose position is tracked over different frames. Then, motion segmentation aims at partitioning pixels extracted from different frames into spatiotemporal regions. In particular, feature point are clustered into different groups, each corresponding to a given motion [46]. The data to be processed by the clustering algorithm is obtained by stacking the coordinate values associated to a given feature point corresponding to different frames. For a detailed description of how clustering data are obtained from feature points coordinates, please refer to [46].

We use the Hopkins 155 motion segmentation dataset [48], which consists of video sequences with two or three motions in each video. Each video of two motions consists of 30 frames, whereas each video of three motions consists of 29 frames. In particular the results reported in this section are obtained by considering the video with three motions in the dataset having the largest number of samples for each motion/class,¹¹ namely, 142 samples for class 1, 114 samples for class 2 and 236 samples for class 3.

We consider in particular a supervised learning approach, in which 50% or 30% of the vectors corresponding to features points are manually labeled, whereas the remaining points are classified automatically, starting from the observation of noisy measurements, where the noise variance is set to $\sigma^2 = -60$ dB. The manually labeled points represent labeled training samples from which the input signal parameters p_i , Σ_i , $i = 1, \dots, L$ are inferred using maximum likelihood (ML) estimators.

As described in [46], [49], [50], features points trajectories belonging to a given motion can be shown to lie on approximately three dimensional affine spaces or four dimensional linear spaces. In fact, the covariance matrices obtained from the training samples present only two dominant principal components, as demonstrated by the magnitudes of eigenvalues of the input covariance matrices reported in Table II. Then, based on the results presented in Propositions 1 and 5, we can expect that at least 3 random measurements and 2 designed measurements are needed for reliable classification, respectively.

Fig. 3(a) and (b) report the misclassification probability vs the number of measurements for random kernels, kernels designed via the construction embodied in Proposition 5, and the designs in [14], [21]. In particular, in view of the fact that the analysis is conducted for the scenario where the MAP classifier is provided with the true model parameters, our results consider both the

¹¹Denoted as "1RT2RCR" in the dataset.

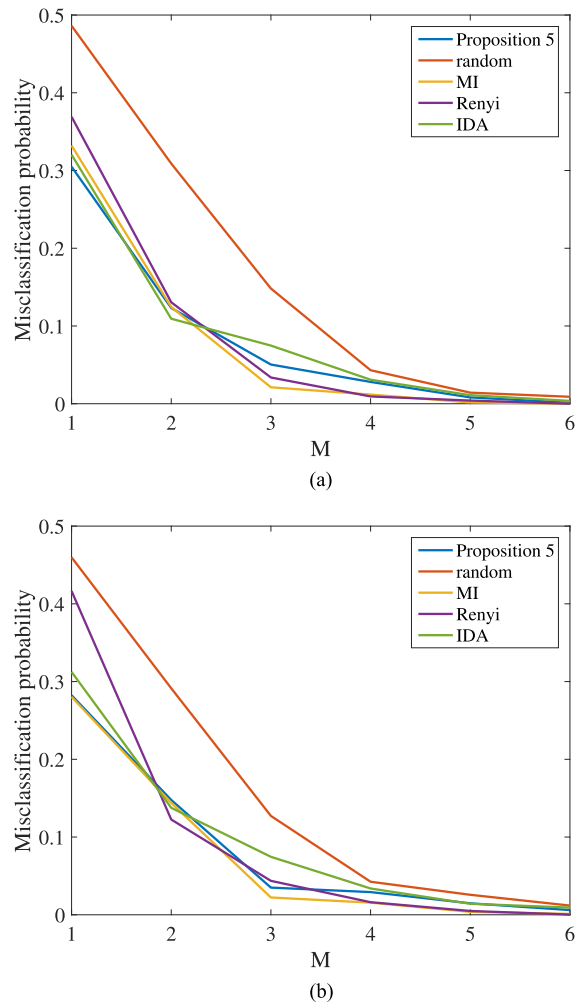


Fig. 3. Misclassification probability vs. M . Hopkins 155, 1RT2RCR dataset. 50% or 30% of the samples are manually labeled and used for training. (a) 50% training samples. (b) 30% training samples.

scenario where a significant number of training samples (50%) is used to learn the underlying models and a scenario where a lower number of training samples (30%) is used to derive the models in order to assess the robustness of the theoretical insights against model mismatch. Note that now the misclassification probability does not exhibit a perfect phase transition in view of the fact that the data covariance matrices are not low-rank anymore but rather approximately low-rank, and due to the mismatch between the model inferred from training data and the actual test data. However, one can still conclude that our theoretical results align with practical ones, since they can unveil the number of measurements required for the misclassification probability to be below a certain low value.

In particular, Table III reports the minimum number of measurements required by the random and designed kernels to achieve a misclassification probability below 15%, 10% and 5%, for both cases when 50% and 30% of the vectors in the dataset are used as training samples. It can be observed that our characterization of the upper bound to the number of measurements required for the phase transition matches

TABLE III
MINIMUM NUMBER OF MEASUREMENTS M REQUIRED TO ACHIEVE A GIVEN VALUE OF THE MISCLASSIFICATION PROBABILITY. HOPKINS 155 DATASET

	random	Prop. 5	IDA	MI	Rényi
50% training samples					
$P_e < 15\%$	3	2	2	2	2
$P_e < 10\%$	4	3	3	3	3
$P_e < 05\%$	4	4	4	3	3
30% training samples					
$P_e < 15\%$	3	2	2	2	2
$P_e < 10\%$	4	3	3	3	3
$P_e < 05\%$	4	3	4	3	3

well the number of measurements required to achieve a low misclassification probability in IDA and methods based on the maximization of Shannon mutual information and Rényi quadratic entropy, in both scenarios where 50% and 30% of the vectors in the dataset are used as training samples.

C. Real Data: Face Recognition

We now consider a different real-world, compressive classification application. In particular, we consider a face recognition problem where the orientation of faces associated to different individuals relative to the camera remains fixed, but the illumination conditions vary. On assuming that faces are approximately convex and that reflect light according to Lambert's law, it is possible to show that the set of images of a same individual under different illuminations lies approximately on a 9-dimensional linear subspace [51]. Therefore, face recognition from linear measurements extracted from such images can be performed via subspace classification.

In this section, we show classification results using cropped images from the Extended Yale Face Database B [52]. In particular, we consider 16×16 images of $L = 5$ different individuals from the 38 available in the dataset. For each individual, 63 images corresponding to 63 different illumination conditions are considered.

As for the video motion segmentation application described in Section VI-B, classification is performed via the MAP classifier (2), where we assume Gaussian distribution for each class and the parameters p_i, Σ_i are obtained via ML estimators by using 50% or 30% of the available images as training samples. Moreover, we set the noise variance to $\sigma^2 = -60$ dB.

In contrast with the case of the Hopkins 155 dataset, samples in the Extended Yale Face Database B are described via an approximately low-rank model which is characterized by a slower decay of the eigenvalues of the corresponding covariance matrices, as reported in Fig. 4. In this sense, experimental results for this dataset represent a way to test the predictions provided by our analysis also for a scenario which departs further from the assumption of signals lying on a union of low-dimensional subspaces.

Fig. 5(a) and (b) report the misclassification probability vs the number of measurements for random kernels, kernels designed via the construction embodied in Proposition 5, and the

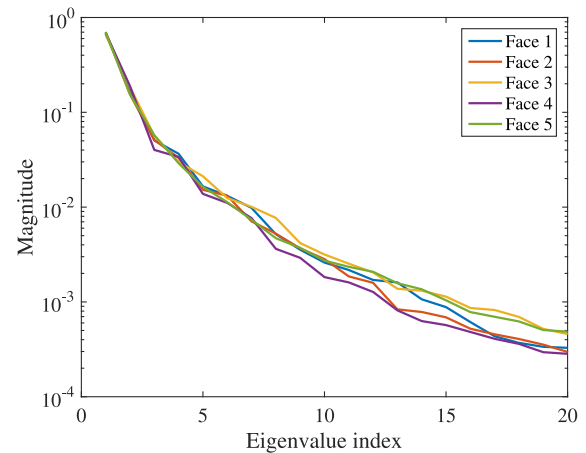


Fig. 4. Largest 20 eigenvalues of the covariance matrices associated to the first $L = 5$ classes in the Extended Yale Face Database B. The covariance matrices are obtained from training samples via the ML estimator.

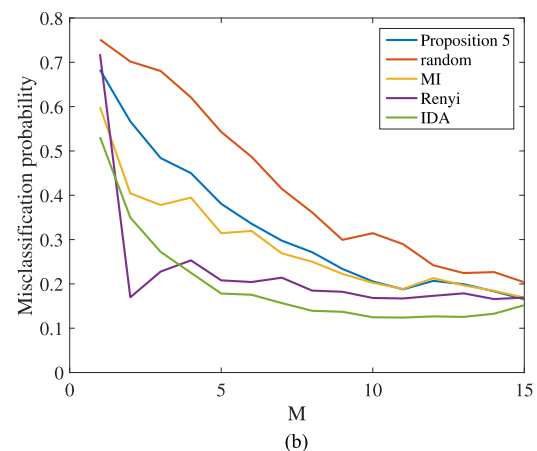
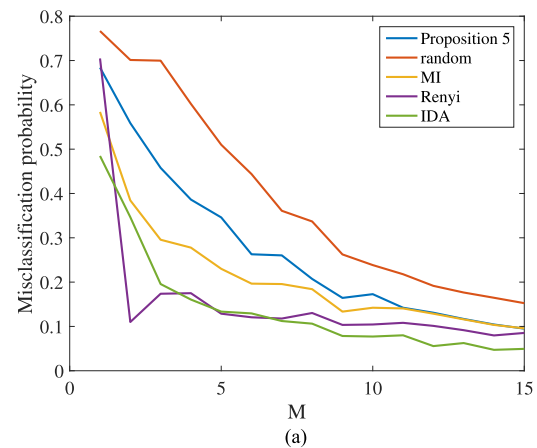


Fig. 5. Misclassification probability vs. M . Extended Yale Face Database B. 50% or 30% of the samples are manually labeled and used for training. (a) 50% training samples. (b) 30% training samples.

designs in [14], [21]. Also in this case, motivated by the fact that the analysis is conducted for the scenario where the MAP classifier is provided with the true model parameters, our results consider both the scenario where a significant number of training samples (50%) is used to learn the underlying models and

TABLE IV
MINIMUM NUMBER OF MEASUREMENTS M REQUIRED TO ACHIEVE
 $P_e < 25\%$. EXTENDED YALE FACE DATABASE B, $L = 5$

training	random	Prop. 5	IDA	MI	Rényi
50%	10	8	3	5	2
30%	12	9	4	8	2

a scenario where a lower number of training samples (30%) is used to derive the models in order to assess the robustness of the theoretical insights against model mismatch. We note that in this case, due to the slow eigenvalue decay reported in Fig. 4, the measurement design described in Section V-B does not provide state-of-the-art classification results, as classification based on measurements extracted via the methods in [14], [21] guarantee lower misclassification probabilities.

On the other hand, it is possible to observe that the theoretical results in Proposition 1 and Proposition 5 indeed capture the actual behavior of classification with state-of-the-art measurement design. In fact, the upper bounds (25) (35) applied to the face recognition scenario under exam predicts that $M = r_{\Sigma} + 1 = 10$ random measurements or $M = L - 1 = 4$ designed measurements are required for reliable classification. Then, based on numerical simulations of classification with non-compressive measurements, we set the baseline misclassification probability for reliable classification at 25%. We observe that the predictions offered by Proposition 1 and Proposition 5 are in line with the trends shown in Table IV, which reports the minimum number of measurements required by random and designed kernels to achieve a misclassification probability below 25% for both cases when 50% and 30% of the vectors in the dataset are used as training samples.

VII. DISCUSSION: IMPACT OF MODEL MISMATCH

It is also instructive to discuss the impact of model mismatch on the classification performance of the MAP classifier (2) in practical application scenarios.

In fact, the analysis carried out in the previous sections assumed that the MAP classifier is given the true model parameters. On the other hand, in practical applications, the conditional pdfs $p(\mathbf{y}|C = i)$ and the prior probabilities p_i are usually learnt from training data, thus implying the introduction of mismatch between the model adopted by the classifier and the actual statistical description of test data.

A proper derivation of the number of measurements required for reliable classification in practical application scenarios would therefore require a more in-depth analysis that takes into account the model mismatch induced by the learning process. In particular, it would require: i) expressions that articulate about the behaviour of the misclassification probability as a function of the true underlying model and the learnt model; ii) a further analysis that determines how compressive random or designed measurements influence the phase transition associated with the misclassification probability.

A preliminary analysis of the impact of model mismatch in classification problems has been conducted in [53], [54]. These

works consider the classification of signals drawn from Gaussian distributions with mismatched classifiers. In particular, they provide sufficient conditions on the relationship between the true model parameters and the learnt model parameters that guarantee reliable classification in the low-noise regime. However, the results in [53], [54] are derived for a non-compressive classification scenario, therefore they cannot explain how compressive random or designed measurements influence the misclassification probability.

A generalization of our analysis on the minimum number of measurements sufficient for reliable classification to capture the impact of model mismatch does not seem immediate. However, our simulation results associated with real-data subspace classification problems in Sections VI-B and VI-C suggest that our theory can still provide meaningful insights both in the situation where we use a significant number of training samples (as expected because we can learn an accurate data model) and in the situation where we use a lower number of training samples. This is despite the fact that the learning process produces distributions that do not correspond exactly to the true ones and also the modelling process assumes a Gaussian distribution that does not necessarily correspond to the true ones pertaining to the motion segmentation or face classification examples.

We conjecture that the reasons for this phenomenon are related to the fact that reliable classification is achieved when compressive measurements are able to discriminate among linear subspaces spanned by signals in the different classes, irrespectively to the particular shape of the marginal distributions that are supported on such subspaces.

In this sense, motion segmentation is more immune to model mismatch than face recognition because, as it is implied by the quick decay of the eigenvalues of the covariance matrices, the majority of the energy of the samples in the motion segmentation dataset is concentrated in linear subspaces of dimension 2 or 3. Then, even a reduced number of training samples is sufficient to identify the dominating principal components for each class. On the other hand, when considering face recognition, the energy associated to samples drawn from a given class is only approximately concentrated on a low-dimensional subspace. In this case, training sets with increased cardinality can guarantee a refined estimation of the principal components associated to each class.

VIII. CONCLUSION

In this paper we have offered a characterization of the number of measurements required to reliably classify linear subspaces modeled via low-rank, zero-mean Gaussian distributions. In particular, we have provided upper bounds to the number of measurements required to drive the misclassification probability to zero both for random measurements as well as designed measurements for two-class classification problems and more challenging multi-class problems. Our characterization suggests that the minimum number of measurements required for phase transition may be achieved by either a one-vs-all approach, or by randomly spreading measurements over the Grassmann manifold, depending on the relationship between the number of

classes and the dimension of the spaces spanned by signals in each class.

One of the hallmarks of our characterizations relates to its ability to predict the minimum number of measurements required to achieve a low-misclassification probability in state-of-the-art measurement design methods. Therefore, it offers engineers a concrete tool to gauge the number of measurements for reliable classification, thereby bypassing the need for time-consuming simulations.

APPENDIX A PROOF OF LEMMA 1

Consider the eigenvalue decomposition of the following matrices $\mathbf{S}_i = \Phi \Sigma_i \Phi^T$, $\mathbf{S}_j = \Phi \Sigma_j \Phi^T$ and $\mathbf{S}_{ij} = \Phi (\Sigma_i + \Sigma_j) \Phi^T$, which yields

$$\mathbf{S}_i = \mathbf{U}_i \Lambda_i \mathbf{U}_i^T \quad (40)$$

$$\mathbf{S}_j = \mathbf{U}_j \Lambda_j \mathbf{U}_j^T \quad (41)$$

$$\mathbf{S}_{ij} = \mathbf{U}_{ij} \Lambda_{ij} \mathbf{U}_{ij}^T. \quad (42)$$

where $\mathbf{U}_i, \mathbf{U}_j, \mathbf{U}_{ij} \in \mathbb{R}^{M \times M}$ are orthogonal matrices; the diagonal matrices $\Lambda_i = \text{diag}(\lambda_{i_1}, \dots, \lambda_{i_{r_i}}, 0, \dots, 0)$, $\Lambda_j = \text{diag}(\lambda_{j_1}, \dots, \lambda_{j_{r_j}}, 0, \dots, 0)$ and $\Lambda_{ij} = \text{diag}(\lambda_{ij_1}, \dots, \lambda_{ij_{r_{ij}}}, 0, \dots, 0)$ contain the eigenvalues of $\mathbf{S}_i, \mathbf{S}_j$ and \mathbf{S}_{ij} , respectively. Note that the number of strictly positive eigenvalues of $\mathbf{S}_i, \mathbf{S}_j$ and \mathbf{S}_{ij} , i.e., the number of strictly positive diagonal entries in Λ_i, Λ_j and Λ_{ij} , is equal to $r_i = \text{rank}(\mathbf{S}_i) = \text{rank}(\Phi \Sigma_i \Phi^T)$, $r_j = \text{rank}(\mathbf{S}_j) = \text{rank}(\Phi \Sigma_j \Phi^T)$ and $r_{ij} = \text{rank}(\mathbf{S}_{ij}) = \text{rank}(\Phi (\Sigma_i + \Sigma_j) \Phi^T)$, respectively.

Then, we recall the expression of the upper bound to the misclassification probability

$$\bar{P}_e = \sum_{i=1}^L \sum_{\substack{j=1 \\ j \neq i}}^L \sqrt{p_i p_j} e^{-K_{ij}}, \quad (43)$$

and we can re-express K_{ij} as

$$\begin{aligned} K_{ij} &= \frac{1}{4} \log \frac{\left(\det \left(\frac{\Phi (\Sigma_i + \Sigma_j) \Phi^T + 2\sigma^2 \mathbf{I}}{2} \right) \right)^2}{\det(\Phi \Sigma_i \Phi^T + \sigma^2 \mathbf{I}) \det(\Phi \Sigma_j \Phi^T + \sigma^2 \mathbf{I})} \\ &= \frac{1}{4} \log \frac{\left(\det \left(\frac{\mathbf{S}_{ij} + 2\sigma^2 \mathbf{I}}{2} \right) \right)^2}{\det(\mathbf{S}_i + \sigma^2 \mathbf{I}) \det(\mathbf{S}_j + \sigma^2 \mathbf{I})} \\ &= \frac{1}{4} \log \left[2^{-2r_{ij}} (\sigma^2)^{r_i + r_j - 2r_{ij}} \right. \\ &\quad \left. \cdot \frac{\prod_{k=1}^{r_{ij}} (\lambda_{ij_k} + 2\sigma^2)^2}{\prod_{k=1}^{r_i} (\lambda_{i_k} + \sigma^2) \prod_{k=1}^{r_j} (\lambda_{j_k} + \sigma^2)} \right], \quad (44) \end{aligned}$$

thus leading to

$$\begin{aligned} \bar{P}_e &= \sum_{i=1}^L \sum_{\substack{j=1 \\ j \neq i}}^L \sqrt{p_i p_j} 2^{r_{ij}/2} (\sigma^2)^{(2r_{ij} - r_i - r_j)/4} \\ &\quad \cdot \left[\frac{\sqrt{\prod_{k=1}^{r_i} (\lambda_{i_k} + \sigma^2) \prod_{k=1}^{r_j} (\lambda_{j_k} + \sigma^2)}}{\prod_{k=1}^{r_{ij}} (\lambda_{ij_k} + 2\sigma^2)^2} \right]^{1/2}. \quad (45) \end{aligned}$$

Then, on letting $\sigma^2 \rightarrow 0$, we note that the term in square brackets converges to the positive constant $[\frac{\sqrt{v_i v_j}}{v_{ij}}]^{1/2}$, moreover, the decay of \bar{P}_e as a function of σ^2 is dominated by the terms in the sum corresponding to the minimum value of the exponent $d(i, j) = (2r_{ij} - r_i - r_j)/4$, thus leading to the result in (18)–(20).

APPENDIX B PROOF OF PROPOSITION 4

The derivation of the upper bound on the number of measurements needed to verify (31) is based on the analysis of the upper bound \bar{P}_e in (16).

Recall that the low-noise expansion exponent d of the upper bound to the misclassification probability for the classification problem of two, zero-mean classes is given by $d = (2r_{12} - r_1 - r_2)/4$.

We first show that, for all possible choices of Φ , it holds $d \leq R/4$ so that there is not any M such that

$$\lim_{\sigma^2 \rightarrow 0} -\frac{\log \bar{P}_e}{\log(1/\sigma^2)} > d_0 \quad (46)$$

for $d_0 \geq R/4$.

Then, we consider the case $d_0 < R/4$ and we derive the minimum number of measurements M needed to verify (46), which represents an upper bound on the minimum number of measurements needed to verify (31).

A. Case Where $d_0 \geq R/4$

Let $r_{\Sigma_{12}} = \text{rank}(\Sigma_1 + \Sigma_2)$. In the following, we show that, for all possible choices of Φ , it holds

$$d = (2r_{12} - r_1 - r_2)/4 \leq (2r_{\Sigma_{12}} - 2r_{\Sigma})/4 = R/4, \quad (47)$$

or, instead,

$$r_{\Sigma_{12}} - r_{12} \geq r_{\Sigma} - r_1 \wedge r_{\Sigma_{12}} - r_{12} \geq r_{\Sigma} - r_2, \quad (48)$$

since (48) implies (47). Consider the generalized eigenvalue decomposition of the positive semidefinite matrices Σ_1 and Σ_2 given by [47, Theorem 8.7.1], namely,

$$\Sigma_1 = \mathbf{X}^{-T} \mathbf{D}_1 \mathbf{X}^{-1} = \mathbf{X}^{-T} \text{diag}(d_{1_1}, \dots, d_{1_N}) \mathbf{X}^{-1}, \quad (49)$$

with $d_{1_i} \geq 0, i = 1, \dots, N$ and

$$\Sigma_2 = \mathbf{X}^{-T} \mathbf{D}_2 \mathbf{X}^{-1} = \mathbf{X}^{-T} \text{diag}(d_{2_1}, \dots, d_{2_N}) \mathbf{X}^{-1}, \quad (50)$$

with $d_{2_i} \geq 0, i = 1, \dots, N$, where \mathbf{X} is a non-singular matrix.

Note that we have

$$r_{12} = \text{rank}(\Phi \mathbf{X}^{-T} (\mathbf{D}_1 + \mathbf{D}_2) \mathbf{X}^{-1} \Phi^{-T}) \quad (51)$$

$$= \text{rank}(\tilde{\Phi} (\mathbf{D}_1 + \mathbf{D}_2)^{\frac{1}{2}}) \quad (52)$$

and likewise,

$$r_1 = \text{rank}(\Phi \mathbf{X}^{-T} \mathbf{D}_1 \mathbf{X}^{-1} \Phi^{-T}) = \text{rank}(\tilde{\Phi} \mathbf{D}_1^{\frac{1}{2}}) \quad (53)$$

and

$$r_2 = \text{rank}(\Phi \mathbf{X}^{-T} \mathbf{D}_2 \mathbf{X}^{-1} \Phi^{-T}) = \text{rank}(\tilde{\Phi} \mathbf{D}_2^{\frac{1}{2}}), \quad (54)$$

where $\tilde{\Phi} = \Phi \mathbf{X}^{-T}$.

On the other hand, the ranks of the input covariance matrices can be expressed as

$$r_{\Sigma_{12}} = \text{rank}(\mathbf{X}^{-T} (\mathbf{D}_1 + \mathbf{D}_2) \mathbf{X}^{-1}) \quad (55)$$

$$= \text{rank}(\mathbf{X}^{-T} (\mathbf{D}_1 + \mathbf{D}_2)^{\frac{1}{2}}) \quad (56)$$

$$= \text{rank}((\mathbf{D}_1 + \mathbf{D}_2)^{\frac{1}{2}}) \quad (57)$$

and

$$r_{\Sigma} = \text{rank}((\mathbf{D}_1)^{\frac{1}{2}}) = \text{rank}((\mathbf{D}_2)^{\frac{1}{2}}). \quad (58)$$

Let us now define the cardinalities of the following sets:

$$k_c = |\{i : d_{1i} > 0 \wedge d_{2i} > 0\}| \quad (59)$$

$$k_1 = |\{i : d_{1i} > 0\}| \quad (60)$$

$$k_2 = |\{i : d_{2i} > 0\}|. \quad (61)$$

Then, it becomes evident that, $r_{\Sigma_{12}} - r_{\Sigma} = k_1 + k_2 - k_c - k_1 = k_2 - k_c = k_1 - k_c$, and, in view of the possible dependence between columns of $\tilde{\Phi}$, $r_{12} - r_1 \leq k_2 - k_c$ and $r_{12} - r_2 \leq k_1 - k_c$, thus concluding the proof of (47).

B. Case Where $d_0 < R/4$

We start by describing an explicit measurement matrix construction that achieves an expansion exponent of the upper bound to the misclassification probability strictly greater than d_0 with $M = \lfloor 4d_0 \rfloor + 1$ measurements. After that, we prove that $M \leq \lfloor 4d_0 \rfloor$ implies $d \leq d_0$ for all possible choices of Φ .

1) *Achievability*: Consider the matrix

$$\Phi_0 = [\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_{n_{\Sigma}}, \mathbf{w}_1, \mathbf{w}_2, \dots, \mathbf{w}_{n_{\Sigma}}]^T, \quad (62)$$

where the sets $[\mathbf{u}_1, \dots, \mathbf{u}_{n_{12}}]$, $[\mathbf{u}_1, \dots, \mathbf{u}_{n_{12}}, \mathbf{v}_1, \dots, \mathbf{v}_{n_{\Sigma}}]$, $[\mathbf{u}_1, \dots, \mathbf{u}_{n_{12}}, \mathbf{w}_1, \dots, \mathbf{w}_{n_{\Sigma}}]$, $\mathbf{u}_i, \mathbf{v}_i, \mathbf{w}_i \in \mathbb{R}^N$, constitute an orthonormal basis of the linear spaces $\mathcal{N}_{12} = \text{Null}(\Sigma_1) \cap \text{Null}(\Sigma_2)$, $\mathcal{N}_1 = \text{Null}(\Sigma_1)$ and $\mathcal{N}_2 = \text{Null}(\Sigma_2)$, respectively, and $n_{12} = [N - 2r_{\Sigma}]^+$, $n_{\Sigma} = \min\{N - r_{\Sigma}, r_{\Sigma}\} = R/2$.

Then, we can write

$$\Phi_0 \Sigma_1 \Phi_0^T = \begin{bmatrix} \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{Q} \end{bmatrix} \quad (63)$$

where

$$\mathbf{Q} = [\mathbf{w}_1, \dots, \mathbf{w}_{n_{\Sigma}}]^T \Sigma_1 [\mathbf{w}_1, \dots, \mathbf{w}_{n_{\Sigma}}], \quad (64)$$

so that we also have $r_1 = \text{rank}(\Phi_0 \Sigma_1 \Phi_0^T) = \text{rank}(\mathbf{Q})$. Now, note that the matrix \mathbf{Q} is the Gram matrix of the set of vectors $\mathbf{q}_i = \Sigma_1^{\frac{1}{2}} \mathbf{w}_i$, $i = 1, \dots, n_{\Sigma}$, and, therefore, $r_1 = \text{rank}(\mathbf{Q}) = n_{\Sigma}$ if and only if the vectors \mathbf{q}_i , $i = 1, \dots, n_{\Sigma}$, are linearly independent.

Assume by contradiction that the vectors \mathbf{q}_i are linearly dependent. Then, there exists a set of n_{Σ} scalars α_i (with $\alpha_i \neq 0$ for at least one index i) such that $\Sigma_1^{\frac{1}{2}} \sum_i \alpha_i \mathbf{w}_i = \mathbf{0}$. It is known that $\sum_i \alpha_i \mathbf{w}_i \neq \mathbf{0}$ because \mathbf{w}_i are linearly independent by construction. Therefore, the linearly dependence among the vectors \mathbf{q}_i implies that $\sum_i \alpha_i \mathbf{w}_i \in \mathcal{N}_1$, which is false since, by construction, $\sum_i \alpha_i \mathbf{w}_i \in \mathcal{N}_2$ and $\sum_i \alpha_i \mathbf{w}_i \notin \mathcal{N}_{12}$. Therefore, we can establish that $r_1 = \text{rank}(\Phi_0 \Sigma_1 \Phi_0^T) = \text{rank}(\mathbf{Q}) = n_{\Sigma}$, and, we can similarly establish that $r_2 = \text{rank}(\Phi_0 \Sigma_2 \Phi_0^T) = n_{\Sigma}$ and $r_{12} = \text{rank}(\Phi_0 (\Sigma_1 + \Sigma_2) \Phi_0^T) = 2n_{\Sigma} = R$.

Finally, we generate Φ by picking arbitrarily only $M = \lfloor 4d_0 \rfloor + 1$ among the R row vectors of the matrix Φ_0 in (62). In particular, we take M_1 rows from the set $[\mathbf{v}_1, \dots, \mathbf{v}_{n_{\Sigma}}]$ and M_2 rows from the set $[\mathbf{w}_1, \dots, \mathbf{w}_{n_{\Sigma}}]$, where $M_1 + M_2 = \lfloor 4d_0 \rfloor + 1$, which is always possible as $\lfloor 4d_0 \rfloor + 1 \leq R$. Then, by following steps similar to the previous ones, it is possible to show that $r_1 = \text{rank}(\Phi \Sigma_1 \Phi^T) = M_2$, $r_2 = \text{rank}(\Phi \Sigma_2 \Phi^T) = M_1$ and $r_{12} = \text{rank}(\Phi (\Sigma_1 + \Sigma_2) \Phi^T) = M_1 + M_2$, thus implying $d = (\lfloor 4d_0 \rfloor + 1)/4 > d_0$.

2) *Converse*: Assume now $M \leq \lfloor 4d_0 \rfloor$. In this case, we can show that, for all possible choices of Φ , it holds

$$d \leq M/4 \leq d_0. \quad (65)$$

This upper bound follows from the solution to the following integer-valued optimization problem¹²:

$$\max_{r_1, r_2, r_{12}} (2r_{12} - r_1 - r_2)/4 \quad (66)$$

subject to: $r_1 + r_2 \geq r_{12}$, $r_1 \leq M$, $r_2 \leq M$, $r_{12} \leq M$ and $r_1, r_2, r_{12} \in \mathbb{Z}_0^+$.

The solution, which can be obtained by considering a linear programming relaxation along with a Branch and Bound approach [55], is given by¹³:

$$r_1 + r_2 = r_{12}, \quad r_{12} = M, \quad (67)$$

$$d = (2r_{12} - r_1 - r_2)/4 = M/4. \quad (68)$$

APPENDIX C

PROOF OF PROPOSITION 5

Let $\mathbf{N}_i \in \mathbb{R}^{N \times (N - r_{\Sigma})}$ be a matrix that contains a basis for the null space \mathcal{N}_i and let $\mathbf{N} = [\mathbf{N}_1, \dots, \mathbf{N}_L]$ be a matrix that contains the concatenation of the bases for all the null spaces $\mathcal{N}_1, \dots, \mathcal{N}_L$. Then, consider the measurement matrices $\Phi \in \mathbb{R}^{M \times N}$ that consist of M rows of \mathbf{N}^T . More precisely, such matrices Φ are obtained by picking M_i rows from \mathbf{N}_i^T , so that $\sum_{i=1}^L M_i = M$.¹⁴

A sufficient condition for (34) is represented by $d > 0$, where d is the decay exponent associated to the misclassification

¹²Note that this problem represents a relaxation of the problem which aims at maximizing d , as it incorporates only some of the constraints dictated by the geometrical description of the scenario. For example, it does not take into account the actual value of some parameters of the input description as r_{Σ} and $r_{\Sigma_{12}}$.

¹³The solution of the optimization problem is not unique. Nevertheless, the maximum value achieved by the objective function is indeed unique.

¹⁴Throughout the proof, we assume $M \leq N$, since the decay exponent d associated to any matrix Φ is always smaller than or equal to the decay exponent associated to the identity matrix \mathbf{I}_N , as it was shown in Appendix B-A.

probability upper bound (16). Moreover, $d > 0$ if and only if $d(i, j) > 0$ for all the pairs (i, j) with $i \neq j$.

We can now express the conditions $d(i, j) > 0$ in terms of the values M_i as follows. On recalling Sylvester's rank theorem [56], which states

$$\text{rank}(\mathbf{AB}) = \text{rank}(\mathbf{B}) - \dim(\text{Im}(\mathbf{B}) \cap \text{Null}(\mathbf{A})), \quad (69)$$

we can write each term $d(i, j)$ as

$$\begin{aligned} d(i, j) &= (2r_{ij} - r_i - r_j) / 4 \\ &= [\dim(\text{Im}(\Phi^T) \cap \mathcal{N}_i) + \dim(\text{Im}(\Phi^T) \cap \mathcal{N}_j) \\ &\quad - 2 \dim(\text{Im}(\Phi^T) \cap \mathcal{N}_{ij})] / 4. \end{aligned} \quad (70)$$

We first show that

$$\dim(\text{Im}(\Phi^T) \cap \mathcal{N}_i) = \max\{M - r_\Sigma, M_i\}. \quad (72)$$

Notice that, since the images \mathcal{R}_i are independently drawn from a continuous pdf over the Grassmann manifold, any $\min\{N, L(N - r_\Sigma)\}$ columns of \mathbf{N} are linearly independent with probability 1. Then, by leveraging the expression of the dimension of the intersection of two linear spaces, we can write

$$\begin{aligned} \dim(\text{Im}(\Phi^T) \cap \mathcal{N}_i) &= \text{rank}(\Phi) + \text{rank}(\mathbf{N}_i) - \text{rank}[\Phi^T \mathbf{N}_i] \\ &= M + (N - r_\Sigma) - \text{rank}[\Phi^T \mathbf{N}_i]. \end{aligned} \quad (73)$$

Moreover,

$$\text{rank}[\Phi^T \mathbf{N}_i] = \text{rank}[\bar{\Phi}^T \mathbf{N}_i], \quad (74)$$

where $\bar{\Phi}^T$ is obtained from Φ^T by deleting the M_i columns corresponding to vectors taken from the basis of the null space \mathcal{N}_i . Then, given that the columns of $\bar{\Phi}^T$ are picked from spaces drawn at random from the Grassmann manifold, we can conclude that

$$\text{rank}[\bar{\Phi}^T \mathbf{N}_i] = \min\{N, M - M_i + N - r_\Sigma\}, \quad (75)$$

and on replacing (75) into (73) we immediately obtain (72).

Consider now the last term in (71) and recall that, since the linear spaces \mathcal{N}_i are drawn independently at random from a continuous pdf, then

$$\dim(\mathcal{N}_{ij}) = \dim(\mathcal{N}_i \cap \mathcal{N}_j) = \max\{N - 2r_\Sigma, 0\}, \quad (76)$$

thus implying immediately that $\dim(\text{Im}(\Phi^T) \cap \mathcal{N}_{ij}) = 0$ if $N \leq 2r_\Sigma$. Therefore, we assume $N > 2r_\Sigma$ and we show that

$$\dim(\text{Im}(\Phi^T) \cap \mathcal{N}_{ij}) = \max\{M - 2r_\Sigma, M_i - r_\Sigma, M_j - r_\Sigma, 0\}. \quad (77)$$

In order to do that, we first note that we can leverage the expression of the dimension of the intersection of two linear subspaces to write

$$\dim(\text{Im}(\Phi^T) \cap \mathcal{N}_{ij}) = M + (N - 2r_\Sigma) - \text{rank}[\Phi^T \mathbf{N}_{ij}], \quad (78)$$

where the columns of \mathbf{N}_{ij} form a basis of the linear space \mathcal{N}_{ij} . Let us also write Φ as

$$\Phi = [\tilde{\Phi}^T \Phi_i^T \Phi_j^T]^T, \quad (79)$$

where the M_i columns of Φ_i^T are vectors picked from a basis of \mathcal{N}_i , the M_j columns of Φ_j^T are vectors picked from a basis of \mathcal{N}_j and the $M - M_i - M_j$ columns of $\tilde{\Phi}^T$ are vectors picked from the bases of the remaining null spaces. Then, on leveraging again the assumption that the linear spaces associated to the different classes are picked independently at random from a continuous distribution, we can write

$$\begin{aligned} \text{rank}[\Phi^T \mathbf{N}_{ij}] &= \min\{N, M - M_i - M_j \\ &\quad + \text{rank}[\Phi_i^T \Phi_j^T \mathbf{N}_{ij}]\}. \end{aligned} \quad (80)$$

On the other hand, on introducing the notation $r_{\Phi_{ij}\mathbf{N}_{ij}} = \text{rank}[\Phi_i^T \Phi_j^T \mathbf{N}_{ij}]$ we also have

$$\begin{aligned} r_{\Phi_{ij}\mathbf{N}_{ij}} &= \text{rank}[\Phi_i^T \mathbf{N}_{ij} \Phi_j^T \mathbf{N}_{ij}] \\ &= \dim(\text{Im}[\Phi_i^T \mathbf{N}_{ij}]) + \dim(\text{Im}[\Phi_j^T \mathbf{N}_{ij}]) \\ &\quad - \dim(\text{Im}[\Phi_i^T \mathbf{N}_{ij}] \cap \text{Im}[\Phi_j^T \mathbf{N}_{ij}]) \\ &= \min\{N - r_\Sigma, M_i + N - 2r_\Sigma\} \\ &\quad + \min\{N - r_\Sigma, M_j + N - 2r_\Sigma\} - (N - 2r_\Sigma). \end{aligned}$$

In fact, $\dim(\text{Im}[\Phi_i^T \mathbf{N}_{ij}]) = \min\{N - r_\Sigma, M_i + N - 2r_\Sigma\}$ derives from the fact that the columns of Φ_i^T and \mathbf{N}_{ij} are all picked at random from the space \mathcal{N}_i – which has dimension $N - r_\Sigma$. Moreover, we have used the fact

$$\dim(\text{Im}[\Phi_i^T \mathbf{N}_{ij}] \cap \text{Im}[\Phi_j^T \mathbf{N}_{ij}]) = N - 2r_\Sigma, \quad (81)$$

which follows from

$$\text{Im}(\mathbf{N}_{ij}) \subseteq \text{Im}[\Phi_i^T \mathbf{N}_{ij}] \cap \text{Im}[\Phi_j^T \mathbf{N}_{ij}] \subseteq \mathcal{N}_i \cap \mathcal{N}_j. \quad (82)$$

Then, on using the symbol $r_{\Phi\mathbf{N}_{ij}} = \text{rank}[\Phi^T \mathbf{N}_{ij}]$, we have

$$\begin{aligned} r_{\Phi\mathbf{N}_{ij}} &= \min\{N, M - M_i - M_j \\ &\quad + \min\{N - r_\Sigma, M_i + N - 2r_\Sigma\} \\ &\quad + \min\{N - r_\Sigma, M_j + N - 2r_\Sigma\} - (N - 2r_\Sigma)\} \end{aligned}$$

and, therefore,

$$\begin{aligned} \dim(\text{Im}(\Phi^T) \cap \mathcal{N}_{ij}) &= \max\{M - 2r_\Sigma, M_i + M_j \\ &\quad + \min\{N - r_\Sigma, M_i + N - 2r_\Sigma\} \\ &\quad + \min\{N - r_\Sigma, M_j + N - 2r_\Sigma\} \\ &\quad - (N - 2r_\Sigma)\}. \end{aligned} \quad (83)$$

Finally, it is possible to show that (83) is equivalent to (77) by considering separately the cases for which $M_i \leq r_\Sigma$ and $M_j \leq r_\Sigma$.

Therefore, by using (71), (72) and (77), we can write the condition $d(i, j) > 0$ as the set of equivalent conditions

$$f(M, M_i, M_j) - 2(M - 2r_\Sigma) > 0 \quad (84)$$

$$f(M, M_i, M_j) - 2(M_i - r_\Sigma) > 0 \quad (85)$$

$$f(M, M_i, M_j) - 2(M_j - r_\Sigma) > 0 \quad (86)$$

$$f(M, M_i, M_j) > 0, \quad (87)$$

where $f(M, M_i, M_j) = \max\{M - r_\Sigma, M_i\} + \max\{M - r_\Sigma, M_j\}$. Then, the upper bound in (35) is obtained as the solution of the integer optimization problem that aims at minimizing $M = \sum_{i=1}^L M_i$ subject to the constraints (84)-(87) and $0 \leq M_i \leq N - r_\Sigma$.

In the remainder of this appendix, we will show that the solution of such minimization problem is given by $M = \min\{L - 1, r_\Sigma + 1\}$, by considering separately two cases. In particular, when $L - 1 \leq r_\Sigma$, we can show that the optimal solution is given by $M = \sum_{i=1}^L M_i = L - 1$. We first observe that such value represents a feasible solution: in fact, by picking only 1 measurement from $L - 1$ out of L null spaces, e.g., by choosing $M_1 = \dots = M_{L-1} = 1$ and $M_L = 0$, we can immediately prove that all the constraints are verified. Then, we also observe that any solutions for which $M < L - 1$ is not feasible: in fact, if $M < L - 1$ there exist at least two indexes k and ℓ such that $M_k = M_\ell = 0$, and therefore at least one of the constraints (87) is not verified.

Consider now the case $L - 1 > r_\Sigma$. In this case the optimal solution of the minimization problem yields $M = r_\Sigma + 1$. In a similar way to the previous case, we start by observing that $M = r_\Sigma + 1$ is a feasible solution, which can be achieved by picking 1 measurement from $r_\Sigma + 1$ different null space, e.g., by picking $M_1 = \dots = M_{r_\Sigma+1} = 1$ and $M_{r_\Sigma+2} = \dots = M_L = 0$. Also in this case it is straightforward to prove that all the constraints are verified. Moreover, it is possible to observe that there is not any feasible solution such that $M < r_\Sigma + 1$, as $r_\Sigma < L - 1$ implies that there exist at least two indexes k and ℓ such that $M_k = M_\ell = 0$, and, therefore, at least one of the constraints (87) is not verified.

REFERENCES

- [1] H. Reboredo, F. Renna, R. Calderbank, and M. R. D. Rodrigues, "Compressive classification," in *Proc. IEEE Int. Symp. Inf. Theory*, Jul. 2013, pp. 674–678.
- [2] H. Reboredo, F. Renna, R. Calderbank, and M. R. D. Rodrigues, "Projections designs for compressive classification," in *Proc. IEEE Global Conf. Signal Inf. Process.*, Dec. 2013, pp. 1029–1032.
- [3] E. Candès, J. K. Romberg, and T. Tao, "Stable signal recovery for incomplete and inaccurate measurements," *Commun. Pure Appl. Math.*, vol. 59, no. 8, pp. 1207–1223, Aug. 2006.
- [4] E. Candès, J. Romberg, and T. Tao, "Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information," *IEEE Trans. Inf. Theory*, vol. 52, no. 2, pp. 489–509, Feb. 2006.
- [5] D. Donoho, "Compressed sensing," *IEEE Trans. Inf. Theory*, vol. 52, no. 4, pp. 1289–1306, Apr. 2006.
- [6] R. Baraniuk, M. Davenport, R. DeVore, and M. Wakin, "A simple proof of the restricted isometry property for random matrices," *Construct. Approx.*, vol. 28, no. 3, pp. 253–263, Dec. 2008.
- [7] J. Tropp and S. Wright, "Computational methods for sparse solution of linear inverse problems," *Proc. IEEE*, vol. 98, no. 6, pp. 948–958, Jun. 2010.
- [8] A. Ashok, P. Baheti, and M. Neifeld, "Compressive imaging system design using task-specific information," *Appl. Opt.*, vol. 47, no. 25, pp. 4457–4471, Sep. 2008.
- [9] P. K. Baheti and M. A. Neifeld, "Adaptive feature-specific imaging: A face recognition example," *Appl. Opt.*, vol. 47, no. 10, pp. B21–B31, Apr. 2008.
- [10] P. Baheti and M. Neifeld, "Recognition using information-optimal adaptive feature-specific imaging," *J. Opt. Soc. Am. A*, vol. 26, no. 4, pp. 1055–1070, 2009.
- [11] J. Mairal, F. Bach, J. Ponce, G. Sapiro, and A. Zisserman, "Discriminative learned dictionaries for local image analysis," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, Jun. 2008, pp. 1–8.
- [12] Y. Hitomi, J. Gu, M. Gupta, T. Mitsunaga, and S. Nayar, "Video from a single coded exposure photograph using a learned over-complete dictionary," in *Proc. IEEE Int. Conf. Comput. Vis.*, 2011, pp. 287–294.
- [13] J. Wright, A. Y. Yang, A. Ganesh, S. Shankar Sastry, and Y. Ma, "Robust face recognition via sparse representation," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 31, no. 2, pp. 210–227, Feb. 2009.
- [14] M. Chen, W. Carson, M. R. D. Rodrigues, R. Calderbank, and L. Carin, "Communication inspired linear discriminant analysis," in *Proc. Int. Conf. Mach. Learn.*, Jun. 2012, pp. 919–926.
- [15] R. O. Duda, P. E. Hart, and D. G. Stork, *Pattern Classification*, 2nd ed. New York, NY, USA: Wiley-Interscience, 2000.
- [16] R. Fisher, "The use of multiple measurements in taxonomic problems," *Ann. Hum. Genet.*, vol. 7, no. 2, pp. 179–188, 1936.
- [17] D. Erdogmus and J. C. Principe, "Lower and upper bounds for misclassification probability based on Renyi's information," *J. VLSI Signal Process. Syst. Signal Image Video Technol.*, vol. 37, nos. 2/3, pp. 305–317, 2004.
- [18] K. Hild, D. Erdogmus, K. Torkkola, and J. Principe, "Feature extraction using information-theoretic learning," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 28, no. 9, pp. 1385–1392, Sep. 2006.
- [19] S. Kaski and J. Peltonen, "Informative discriminant analysis," in *Proc. Int. Conf. Mach. Learn.*, 2003, pp. 329–336.
- [20] L. Liu and P. Fieguth, "Texture classification from random features," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 34, no. 3, pp. 574–586, Mar. 2012.
- [21] Z. Nenadic, "Information discriminant analysis: Feature extraction with an information-theoretic objective," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 29, no. 8, pp. 1394–1407, Aug. 2007.
- [22] D. Tao, X. Li, X. Wu, and S. Maybank, "Geometric mean for subspace selection," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 31, no. 2, pp. 260–274, Feb. 2009.
- [23] K. Torkkola, "Learning discriminative feature transforms to low dimensions in low dimensions," in *Proc. Adv. Neural Inf. Process. Syst.*, 2001, pp. 3–8.
- [24] K. Torkkola, "Feature extraction by non-parametric mutual information maximization," *J. Mach. Learn. Res.*, vol. 3, pp. 1415–1438, Mar. 2003.
- [25] Q. Qiu and G. Sapiro, "Learning transformations for clustering and classification," *J. Mach. Learn. Res.*, vol. 16, pp. 187–225, 2015.
- [26] R. Durrant and A. Kabán, "Random projections as regularizers: Learning a linear discriminant from fewer observations than dimensions," *Mach. Learn.*, vol. 99, no. 2, pp. 257–286, 2015.
- [27] M. Davenport, P. Boufounos, M. Wakin, and R. Baraniuk, "Signal processing with compressive measurements," *IEEE J. Sel. Topics Signal Process.*, vol. 4, no. 2, pp. 445–460, Apr. 2010.
- [28] M. Duarte, M. Davenport, M. Wakin, and R. Baraniuk, "Sparse signal detection from incoherent projections," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2006, vol. 3, pp. 305–308.
- [29] J. Haupt, R. Castro, R. Nowak, G. Fudge, and A. Yeh, "Compressive sampling for signal classification," in *Proc. Asilomar Conf. Signals Syst. Comput.*, 2006, pp. 1430–1434.
- [30] J. Haupt and R. Nowak, "Compressive sampling for signal detection," in *Proc. IEEE Int. Conf. Acoust., Speech, Signal Process.*, 2007, vol. 3, pp. 1509–1512.
- [31] T. Wimalajeewa, H. Chen, and P. Varshney, "Performance limits of compressive sensing-based signal classification," *IEEE Trans. Signal Process.*, vol. 60, no. 6, pp. 2758–2770, Jun. 2012.
- [32] K. Krishnamurthy, M. Raginsky, and R. Willett, "Hyperspectral target detection from incoherent projections," in *Proc. IEEE Int. Conf. Acoust. Speech Signal Process.*, 2010, pp. 3550–3553.
- [33] K. Krishnamurthy, M. Raginsky, and R. Willett, "Hyperspectral target detection from incoherent projections: Nonequiprobable targets and inhomogeneous SNR," in *Proc. IEEE Int. Conf. Image Process.*, 2010, pp. 1357–1360.
- [34] M. A. Neifeld, A. Ashok, and P. K. Baheti, "Task-specific information for imaging system analysis," *J. Opt. Soc. Am. A*, vol. 24, no. 12, pp. B25–B41, Dec. 2007.
- [35] J. Ke, A. Ashok, and M. A. Neifeld, "Object reconstruction from adaptive compressive measurements in feature-specific imaging," *Appl. Opt.*, vol. 49, no. 34, pp. H27–H39, Dec. 2010.
- [36] J. M. Duarte-Carvajalino, G. Yu, L. Carin, and G. Sapiro, "Task-driven adaptive statistical compressive sensing of Gaussian mixture models," *IEEE Trans. Signal Process.*, vol. 61, no. 3, pp. 585–600, Feb. 2013.
- [37] K. Fukunaga, *Introduction to Statistical Pattern Recognition*, vol. 2. New York, NY, USA: Academic, 1990.
- [38] S. Zhang and T. Sim, "Discriminant subspace analysis: A Fukunaga-Koontz approach," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 29, no. 10, pp. 1732–1745, Oct. 2007.

- [39] T. Blumensath and M. Davies, "Sampling theorems for signals from the union of finite-dimensional linear subspaces," *IEEE Trans. Inf. Theory*, vol. 55, no. 4, pp. 1872–1882, Apr. 2009.
- [40] Y. Eldar and M. Mishali, "Robust recovery of signals from a structured union of subspaces," *IEEE Trans. Inf. Theory*, vol. 55, no. 11, pp. 5302–5316, Nov. 2009.
- [41] R. G. Baraniuk, V. Cevher, M. F. Duarte, and C. Hegde, "Model-based compressive sensing," *IEEE Trans. Inf. Theory*, vol. 56, no. 4, pp. 1982–2001, Apr. 2010.
- [42] R. G. Baraniuk and M. Wakin, "Random projections of smooth manifolds," *Found. Comput. Math.*, vol. 9, no. 1, pp. 51–77, 2009.
- [43] M. Chen, J. Silva, J. Paisley, C. Wang, D. Dunson, and L. Carin, "Compressive sensing on manifolds using a nonparametric mixture of factor analyzers: Algorithm and performance bounds," *IEEE Trans. Signal Process.*, vol. 58, no. 12, pp. 6140–6155, Dec. 2010.
- [44] F. Renna, R. Calderbank, L. Carin, and M. R. D. Rodrigues, "Reconstruction of Gaussian mixture models from compressive measurements: A phase transition view," in *Proc. IEEE Global Conf. Signal Inf. Process.*, 2013, pp. 628–628.
- [45] F. Renna, R. Calderbank, L. Carin, and M. R. D. Rodrigues, "Reconstruction of signals drawn from a Gaussian mixture via noisy compressive measurements," *IEEE Trans. Signal Process.*, vol. 62, no. 9, pp. 2265–2277, May 2014.
- [46] E. Elhamifar and R. Vidal, "Sparse subspace clustering: Algorithm, theory, and applications," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 35, no. 11, pp. 2765–2781, Nov. 2013.
- [47] G. H. Golub and C. F. Van Loan, *Matrix Computations*, 3rd ed. Baltimore, MD, USA: The Johns Hopkins Univ. Press, 1996.
- [48] R. Tron and R. Vidal, "A benchmark for the comparison of 3-D motion segmentation algorithms," in *Proc. IEEE Conf. Comput. Vis. Pattern Recog.*, Jun. 2007, pp. 1–8.
- [49] C. Tomasi and T. Kanade, "Shape and motion from image streams under orthography: A factorization method," *Int. J. Comput. Vis.*, vol. 9, no. 2, pp. 137–154, Nov. 1992.
- [50] T. E. Boult and L. G. Brown, "Factorization-based segmentation of motions," in *Proc. IEEE Workshop Vis. Motion*, Oct. 1991, pp. 179–186.
- [51] R. Basri and D. W. Jacobs, "Lambertian reflectance and linear subspaces," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 25, no. 2, pp. 218–233, Feb. 2003.
- [52] A. S. Georghiades, P. N. Belhumeur, and D. J. Kriegman, "From few to many: Illumination cone models for face recognition under variable lighting and pose," *IEEE Trans. Pattern Anal. Mach. Intell.*, vol. 23, no. 6, pp. 643–660, Jun. 2001.
- [53] J. Sokolić, F. Renna, R. Calderbank, and M. R. D. Rodrigues, "Mismatch in the classification of linear subspaces: Upper bound to the probability of error," in *Proc. IEEE Int. Symp. Inform. Theory*, Jun. 2015, pp. 2201–2205.
- [54] J. Sokolić, F. Renna, R. Calderbank, and M. R. D. Rodrigues, "Mismatch in the classification of linear subspaces: Sufficient conditions for reliable classification," *IEEE Trans. Signal Process.*, vol. 64, no. 12, pp. 3035–3050, Jun. 2016.
- [55] A. Schrijver, *Theory of Linear and Integer Programming*. West Sussex, U.K.: Wiley, 1998.
- [56] C. D. Meyer, *Matrix Analysis and Applied Linear Algebra*, vol. 2. Philadelphia, PA, USA: SIAM, 2000.



Francesco Renna (S'09–M'11) received the Laurea Specialistica degree in telecommunication engineering and the Ph.D. degree in information engineering, both from the University of Padova, Padova, Italy, in 2006 and 2011, respectively. Between 2007 and 2015, he held visiting researcher and postdoctoral appointments at Infineon Technology AG, Princeton University, Georgia Institute of Technology (Lorraine Campus), Supélec, University of Porto, Duke University, and University College London. Since 2016, he has been a Marie Skłodowska Curie Fellow at the University of Cambridge. His research interests focus on high-dimensional information processing, but also include physical layer security for multicarrier and multiantenna systems.



Robert Calderbank (M'89–SM'97–F'98) received the B.Sc. degree in 1975 from Warwick University, Coventry, U.K., the M.Sc. degree in 1976 from Oxford University, Oxford, U.K., and the Ph.D. degree in 1980 from California Institute of Technology, Pasadena, CA, USA, all in mathematics. He is a Professor of electrical engineering at Duke University where he now directs the Information Initiative at Duke after serving as the Dean of Natural Sciences during 2010–2013. He was previously a Professor of electrical engineering and mathematics at Princeton University, where he directed the Program in Applied and Computational Mathematics. Prior to joining Princeton in 2004, he was the Vice President for Research at AT&T, responsible for directing the first industrial research lab in the world where the primary focus is data at scale. At the start of his career at Bell Labs, innovations by him were incorporated in a progression of voiceband modem standards that moved communications practice close to the Shannon limit. Together with P. Shor and colleagues at AT&T Labs, he showed that good quantum error correcting codes exist and developed the group theoretic framework for quantum error correction. He is a coinventor of space-time codes for wireless communication, where correlation of signals across different transmit antennas is the key to reliable transmission.

He served as the Editor-in-Chief of the IEEE TRANSACTIONS ON INFORMATION THEORY from 1995 to 1998, and an Associate Editor for Coding Techniques from 1986 to 1989. He was a Member of the Board of Governors of the IEEE Information Theory Society from 1991 to 1996 and from 2006 to 2008. He received the IEEE Information Theory Prize Paper Award in 1995 for his work on the Z4 linearity of Kerdock and Preparata Codes (joint with A. R. Hammons, Jr., P. V. Kumar, N. J. A. Sloane, and P. Solé), and again in 1999 for the invention of space-time codes (joint with V. Tarokh and N. Seshadri). He also received the 2006 IEEE Donald G. Fink Prize Paper Award, the IEEE Millennium Medal, the 2013 IEEE Richard W. Hamming Medal, the 2015 Shannon Award, and was elected to the U.S. National Academy of Engineering in 2005.



Miguel R. D. Rodrigues (S'98–M'02–SM'15) received the Licenciatura degree in electrical engineering from the Faculty of Engineering, University of Porto, Porto, Portugal, in 1998, and the Ph.D. degree in electronic and electrical engineering from University College London, London, U.K., in 2002.

He is currently a Reader in the Department of Electronic and Electrical Engineering, University College London. He was previously in the Department of Computer Science, University of Porto, Portugal, where he also led the Information Theory and Communications Research Group at Instituto de Telecomunicações Porto. He has also held visiting research appointments at Princeton University, USA, Duke University, USA, Cambridge University, U.K., and University College London, in the period 2007 to 2013. His research interests include the general areas of information theory, communications theory, and signal processing.

He received the IEEE Communications and Information Theory Societies Joint Paper Award in 2011 for the work on Wireless Information-Theoretic Security (with M. Bloch, J. Barros, and S. W. McLaughlin). He also received the Prize Engenheiro António de Almeida, the Prize Engenheiro Cristiano Spratley, and the Merit Scholarship from the University of Porto, and the Best Student Poster Prize at the 2nd IMA Conference on Mathematics in Communications. He also received doctoral and postdoctoral research fellowships from the Portuguese Foundation for Science and Technology, and research fellowships from Foundation Calouste Gulbenkian. He is an Associate Editor of the IEEE COMMUNICATIONS LETTERS.



Hugo Reboredo (S'10) received the integrated M.S. degree in electrical and computers engineering, major in telecommunications, from Faculdade de Engenharia da Universidade do Porto, Porto, Portugal, in 2009, and the Ph.D. degree in telecommunications from a joint program between Universities of Minho, Aveiro and Porto, Portugal, in 2014. He is currently working in planning and optimization of mobile networks. His research interests include physical layer security, signal processing, and compressed sensing.