



HAL
open science

Sketching for Large-Scale Learning of Mixture Models

Nicolas Keriven, Anthony Bourrier, Rémi Gribonval, Patrick Pérez

► **To cite this version:**

Nicolas Keriven, Anthony Bourrier, Rémi Gribonval, Patrick Pérez. Sketching for Large-Scale Learning of Mixture Models. *Information and Inference*, 2018, 7 (3), pp.447-508. 10.1093/imaiai/iax015 . hal-01329195v2

HAL Id: hal-01329195

<https://inria.hal.science/hal-01329195v2>

Submitted on 5 May 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Sketching for Large-Scale Learning of Mixture Models

NICOLAS KERIVEN*,

Université Rennes 1, INRIA Rennes-Bretagne Atlantique
Campus de Beaulieu, Rennes, France

*Corresponding author: Email: nicolas.keriven@inria.fr

ANTHONY BOURRIER

Gipsa-Lab, 11 rue des Mathématiques, St-Martin-d’Hérès, France

Email: anthony.bourrier@gmail.com

RÉMI GRIBONVAL

INRIA Rennes-Bretagne Atlantique, Campus de Beaulieu, Rennes, France

Email: remi.gribonval@inria.fr

AND

PATRICK PÉREZ

Technicolor, 975 Avenue des Champs Blancs, Cesson Sévigné, France

Email: patrick.perez@technicolor.fr

May 5, 2017

Abstract

Learning parameters from voluminous data can be prohibitive in terms of memory and computational requirements. We propose a “compressive learning” framework where we estimate model parameters from a *sketch* of the training data. This sketch is a collection of generalized moments of the underlying probability distribution of the data. It can be computed in a single pass on the training set, and is easily computable on streams or distributed datasets. The proposed framework shares similarities with compressive sensing, which aims at drastically reducing the dimension of high-dimensional signals while preserving the ability to reconstruct them.

To perform the estimation task, we derive an iterative algorithm analogous to sparse reconstruction algorithms in the context of linear inverse problems. We exemplify our framework with the compressive estimation of a Gaussian Mixture Model (GMM), providing heuristics on the choice of the sketching procedure and theoretical guarantees of reconstruction. We experimentally show on synthetic data that the proposed algorithm yields results comparable to the classical Expectation-Maximization (EM) technique while requiring significantly less memory and fewer computations when the number of database elements is large. We further demonstrate the potential of the approach on real large-scale data (over 10^8 training samples) for the task of model-based speaker verification.

Finally, we draw some connections between the proposed framework and approximate Hilbert space embedding of probability distributions using random features. We show that the proposed sketching operator can be seen as an innovative method to design translation-invariant kernels adapted to the analysis of GMMs. We also use this theoretical framework to derive information preservation guarantees, in the spirit of infinite-dimensional compressive sensing.

Compressive Sensing, Compressive Learning, Database Sketching, Gaussian Mixture Model

1 Introduction

An essential challenge in machine learning is to develop scalable techniques able to cope with large-scale training collections comprised of numerous elements of high dimension. To achieve this goal, it is necessary to come up with approximate learning schemes which perform the learning task with a fair precision while reducing the memory and computational requirements compared to standard techniques. A possible way to achieve such savings is to compress data beforehand, as illustrated in Figure 1. Instead of the more classical individual compression of each element of the database with random projections

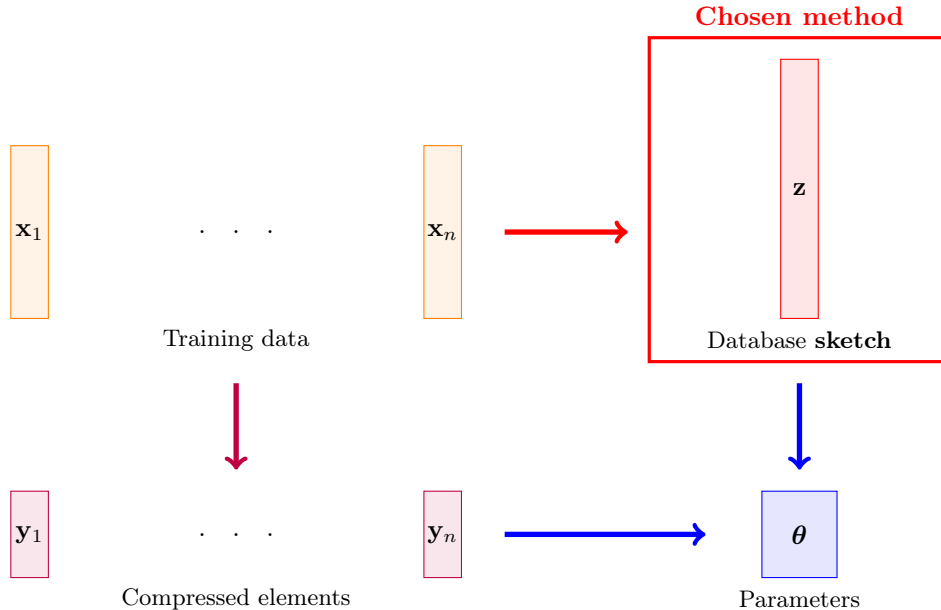


Figure 1: **Paths to compressive learning.** The training data is compressed into a smaller representation, typically through random projections. This can either consist in reducing the dimensions of each individual entry (left bottom) or in computing a more global compressed representation of the data, called *sketch* (top right). Parameters are then inferred from such a compressed representation by a learning algorithm adapted to it. The proposed approach explores the second sketch-based option.

[2, 53, 24, 84, 72] the framework we consider in this paper corresponds to the top right scheme: a large collection of training data is compressed into a fixed-size representation called *sketch*. The dimension of the sketch –and therefore the cost of inferring/learning the parameters of interest from this sketch– does not depend on the number of data items in the initial collection. A complementary path to handling large-scale collections is online learning [30]. Sketching, which leverages ideas originating from streaming algorithms [41], can trivially be turned into an online version and is amenable to distributed computing.

1.1 From compressive sensing to sketched learning.

As we will see, compressing a training collection into a sketch before learning is reminiscent of –and indeed inspired by– compressive sensing (CS) [52] and streaming algorithms [40, 41]. The main goal of CS is to find a dimensionality-reducing linear operator \mathbf{M} such that certain high-dimensional vectors (or signals) can be reconstructed from their observations by \mathbf{M} . Initial work on CS [29, 46] showed that such a reconstruction is possible for k -sparse signals of dimension d ; from only $\mathcal{O}(k)$ linear measurements by (theoretically) solving an intractable NP-complete problem ([52], chap. 2), and in practice from $\mathcal{O}(k \ln(d))$ linear measurements by solving a convex relaxation of this problem. Matrices \mathbf{M} with such reconstruction guarantees can be obtained as typical draws of certain random matrix ensembles. This reconstruction paradigm from few random measurements has subsequently been considered and proven to work for more general signal models [17]. Examples of such models include low-rank matrices [25], cospase vectors [75] and dictionary models [26]. Reconstruction from compressive measurements for these models is made possible, among other properties, by the fact that they correspond to unions of subspaces [12] which have a much lower dimension than the ambient dimension.

Low-dimensional models also intervene in learning procedures, where one aims at fitting a model of moderate “dimension” to some training data $\{\mathbf{x}_1 \dots \mathbf{x}_n\} \subset X$ in order to prevent overfitting and ensure good generalization properties. In this paper, we consider *mixture models* comprising probability distributions on the set X of the form

$$P = \sum_{k=1}^K \alpha_k P_k, \quad (1)$$

where the P_k ’s are probability distributions taken in a certain set and the α_k ’s, $\alpha_k \geq 0$, $\sum_k \alpha_k = 1$, are the weights of the mixture. Such a model can be considered as a generalized sparse model in the linear space E of finite measures over the set X .

Similarly to compressive sensing, one can define a linear compressive operator $\mathcal{A} : E \rightarrow \mathbb{C}^m$ which computes *generalized moments* [60] of a measure ν :

$$\mathcal{A} : \nu \mapsto \mathcal{A}\nu := \frac{1}{\sqrt{m}} \left[\int_X M_1 d\nu, \dots, \int_X M_m d\nu \right], \quad (2)$$

where the M_j 's are well-chosen functions on X and the constant $\frac{1}{\sqrt{m}}$ is used for normalization purposes. In the case where ν is a probability measure P , the integrals are the expectations of $M_j(\mathbf{x})$ with $\mathbf{x} \sim P$.

Given some training data $\mathcal{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ drawn from X , the corresponding empirical distribution is

$$\hat{P} = \frac{1}{n} \sum_{i=1}^n \delta_{\mathbf{x}_i}, \quad (3)$$

where $\delta_{\mathbf{x}_i}$ is a unit mass at \mathbf{x}_i . A practical sketch of the data can then be defined¹ and computed as

$$\hat{\mathbf{z}} = \mathcal{A}\hat{P} = \frac{1}{n\sqrt{m}} \left[\sum_{i=1}^n M_j(\mathbf{x}_i) \right]_{j=1, \dots, m}. \quad (4)$$

Denoting $\Sigma \subset E$ the set of probability distributions satisfying (1), fitting a probability mixture to the training collection \mathcal{X} in a compressive fashion can be expressed as the following optimization problem

$$\operatorname{argmin}_{P \in \Sigma} \|\hat{\mathbf{z}} - \mathcal{A}P\|_2, \quad (5)$$

which corresponds to the search for the probability mixture in the model set Σ whose sketch is closest to the empirical data sketch $\hat{\mathbf{z}}$. By analogy with sparse reconstruction, we propose an iterative greedy reconstruction algorithm to empirically address this problem, and exemplify our framework on the estimation of GMMs.

1.2 Related works

A large set of the existing literature on random projections for dimension reduction in the context of learning focuses on the scheme represented on the bottom left of Figure 1: each item of the training collection is individually compressed with random projections [2, 53] prior to learning for classification [24, 84] or regression [72], or to fitting a GMM [44]. In contrast, we consider here a framework where the whole training collection is compressed to a fixed-size sketch, corresponding to the top right scheme in Figure 1. This framework builds on work initiated in [19][20]. Compared to [19][20], the algorithms we propose here: a) are inspired by Orthogonal Matching Pursuits rather than Iterative Hard Thresholding; b) can handle GMMs with arbitrary diagonal variances; c) yield much better empirical performance (see Section 5 for a thorough experimental comparison).

Our approach bears connections with the Generalized Method of Moments (GeMM) [60], where parameters are estimated by matching empirical generalized moments computed from the data with theoretical generalized moments of the distribution. Typically used in practice when the considered probability models do not yield explicit likelihoods, the GeMM also provides mathematical tools to study the identifiability and learnability of parametric models such as GMMs [10, 9, 63]. Using the empirical characteristic function is a natural way of obtaining moment conditions [49, 50, 101]. Following developments of GeMM with a continuum of moments instead of a finite number of them [31], powerful estimators can be derived when the characteristic function is available at all frequencies simultaneously [32, 107, 33]. Yet, these estimators are rarely implementable in practice.

This is naturally connected with a formulation of mixture model estimation as a linear inverse problem. In [22, 11] for example, this is addressed by considering a finite and incoherent dictionary of densities and the unknown density is reconstructed from its scalar products with every density of the dictionary. These scalar products can be interpreted as generalized moments of the density of interest. Under these assumptions, the authors provide reconstruction guarantees for their cost functions. In our framework, we consider possibly infinite and even uncountable dictionaries, and only collect a limited number of ‘‘measurements’’, much smaller than the number of elements in the dictionary.

Sketching is a classical technique in the database literature [41]. A sketch is a fixed-size data structure which is updated with each element of a data stream, allowing one to perform some tasks without

¹Any other unbiased empirical estimator of the moments, for example using the empirical median, can be envisioned.

keeping the data stored. Applications include the detection of frequent elements, called *heavy hitters* [39] and simple statistical estimations on the data [55]. The sketches used in these works typically involve quantization steps which we do not perform in our work. We also consider the density estimation problem, which is closer to machine learning than the typical application of sketches. Closer to our work, the authors in [100] consider the estimation of 2-dimensional histograms from random linear sketches. Even though this last method is theoretically applicable in higher dimensions, the complexity would grow exponentially with the dimension of the problem. Such a “curse of dimensionality” is also found in [22, 11]: the size of the dictionary grows exponentially with the dimension of the data vectors, and naturally impacts the cost of the estimation procedure. In our work, we rather consider parametric dictionaries that are described by a finite number of parameters. This enables us to empirically leverage the structure of iterative algorithms from sparse reconstruction and compressive sensing to optimize with respect to these parameters, offering better scalability to higher dimensions. This is reminiscent of generalized moments methods for the reconstruction of measures supported on a finite subset of the real line [34], and can be applied to much more general families of probability measures.

The particular sketching operator that we propose to apply on GMMs (see Section 3 and further) is obtained by randomly sampling the (empirical) characteristic function of the distribution of the training data. This can be seen as a combination between two techniques from the Reproducing Kernel Hilbert Space (RKHS) literature, namely embedding of probability distributions in RKHS using a feature map referred to as the “Mean Map” [16, 94, 98], and Random Fourier Features (RFFs) for approximating translation-invariant reproducing kernels [82].

Embedding of probability distributions in RKHS with the Mean Map has been successfully used for a large variety of tasks, such as two-sample test [57], classification [74] or even performing algebraic operations on distributions [92]. In [96], the estimation of a mixture model with respect to the metric of the RKHS is considered with a greedy algorithm. The proposed algorithm is however designed to approximate the target distribution by a large mixture with many components, resulting in an approximation error that decreases as the number of components increases, while our approach considers (1) as a “sparse” combination of a fixed, limited number of components which we aim at identifying. Furthermore, unlike our method that uses RFFs to obtain an efficient algorithm, the algorithm proposed in [96] does not seem to be directly implementable.

Many kernel methods can be performed efficiently using finite-dimensional, nonlinear embeddings that approximate the feature map of the RKHS [82, 103]. A popular method approximates translation-invariant kernels with RFFs [82]. There has been since many variants of RFFs that are faster [69, 108, 35], more precise [106], or designed for a different type of kernel [103]. Similar to our sketching operator, structures combining RFFs with Mean Map embedding of probability distributions have been recently used by the kernel community [15, 99, 77] to accelerate methods such as classification with the so-called Support Measure Machine [74, 99, 77] or two-sample test [110, 36, 65, 78].

Our point of view, *i.e.* that of generalized compressive sensing, is sensibly different: we consider the sketch as a *compressed* representation of the probability distribution, and demonstrate that it contains enough information to robustly recover the distribution from it, resulting in an effective “compressive learning” alternative to usual mixture estimation algorithms.

1.3 Contributions and outline

Our main contributions can be summarized as follows:

- Inspired by Orthogonal Matching Pursuit (OMP) and its variant OMP with Replacement (OMPR), we design in Section 2 an algorithmic framework for general compressive mixture estimation.
- In the specific context of GMMs, we design in Section 3.2 an algorithm that scales better with the number K of mixture components.
- Inspired by random Fourier sampling for compressive sensing, we consider sketching operators \mathcal{A} defined in terms of random sampling of the characteristic function [19, 20]. However we show that the sampling pattern of [19, 20] is not adapted in high dimension. In Section 3.3, in the specific case of GMMs we propose a new heuristic and devise a practical scheme for randomly drawing the frequencies that define \mathcal{A} . This is empirically demonstrated to yield significantly improved performance in Section 5.

- We establish in Section 4 a connection between the choice of the proposed sampling pattern and the design of a reproducing kernel on probability distributions. Compared to existing literature [58, 77], our method is relatively simpler, faster to perform and fully unsupervised.
- Extensive tests on synthetic data in Section 5 demonstrate that our approach matches the estimation precision of a state-of-the-art C++ implementation of the EM algorithm while enabling significant savings in time and memory.
- In the context of hypothesis testing-based speaker verification, we also report in Section 6 results on real data, where we exploit a corpus of 1000 hours of speech at scales inaccessible to traditional methods, and match using a very limited number of measurements the results obtained with EM.
- We provide preliminary theoretical results (Theorem 3 in Section 7) on the information preservation guarantees of the sketching operator. The proofs of these results (Appendices B and C) introduce a new variant of the Restricted Isometry Property (RIP) [28], connected here to kernel mean embedding and Random Features. Compared to usual guarantees in the GeMM literature, our results have less of a “statistics” flavor and more of a “signal processing” one, such as robustness to modeling error, *i.e.* the true distribution of the data is not exactly a GMM but close to one.

2 A Compressive Mixture Estimation Framework

In classical compressive sensing [52], a signal $\mathbf{x} \in \mathbb{R}^d$ is encoded with a measurement matrix $\mathbf{M} \in \mathbb{R}^{m \times d}$ into a compressed representation $\mathbf{z} \in \mathbb{R}^m$:

$$\mathbf{z} = \mathbf{M}\mathbf{x} \quad (6)$$

and the goal is to recover \mathbf{x} from those linear measurements. Often the system is underdetermined ($m < d$) and recovery can only be done with additional assumptions, typically sparsity. The vector $\mathbf{x} = [x_\ell]_{\ell=1}^d$ is said to be *sparse* if it has only a limited number $k < d$ of non-zero coefficients. Its *support* is the set of indices of non-zero entries: $\Gamma(\mathbf{x}) = \{\ell \mid x_\ell \neq 0\}$. The notation \mathbf{M}_Γ (resp. \mathbf{x}_Γ) denotes the restriction of matrix \mathbf{M} (resp. vector \mathbf{x}) to columns (resp. entries) with indices in Γ .

A sparse vector can be seen as a combination of few basis elements: $\mathbf{x} = \sum_{\ell \in \Gamma} x_\ell \mathbf{e}_\ell$, where $\{\mathbf{e}_\ell\}_{\ell=1, \dots, d}$ is the canonical basis of \mathbb{R}^d . The measurement vector \mathbf{z} is thus expressed as a combination of few *atoms*, corresponding to the columns of the measurement matrix: $\mathbf{z} = \sum_{\ell \in \Gamma} x_\ell \mathbf{M}\mathbf{e}_\ell$. The set of all atoms is referred to as a *dictionary* $\{\mathbf{M}\mathbf{e}_\ell\}_{\ell=1, \dots, d}$.

2.1 Mixture model and generalized compressive sensing

Let $E = E(X)$ be a space of signed finite measures over a measurable space (X, \mathcal{B}) , and \mathcal{P} the set of probability distributions over X , $\mathcal{P} := \{P \in E; P \geq 0, \int_X dP = 1\}$. In our framework, a distribution $P \in \mathcal{P}$ is encoded with a linear measurement operator (or *sketching* operator) $\mathcal{A} : \mathcal{P} \rightarrow \mathbb{C}^m$:

$$\mathbf{z} = \mathcal{A}P. \quad (7)$$

As in classical compressive sensing, we define a “sparse” model in \mathcal{P} . As mentioned in the introduction, it is here assimilated to a mixture model (1), generated by combining elements from some given set $\mathcal{G} = \{P_\theta\}_{\theta \in \mathcal{T}} \subset \mathcal{P}$ of basic distributions indexed by a parameter $\theta \in \mathcal{T}$. For some finite $K \in \mathbb{N}^*$, a distribution is thus said to be *K-sparse* (in \mathcal{G}) if it is a convex combination of K elements from \mathcal{G} :

$$P_{\Theta, \alpha} = \sum_{k=1}^K \alpha_k P_{\theta_k}, \quad (8)$$

with $P_{\theta_k} \in \mathcal{G}$, $\alpha_k \geq 0$ for all k 's, and $\sum_{k=1}^K \alpha_k = 1$. We name *support* of the representation² (Θ, α) of such a sparse distribution the set of parameters $\Theta = \{\theta_1, \dots, \theta_K\}$.

The measurement vector $\mathbf{z} = \mathcal{A}P_{\Theta, \alpha} = \sum_{k=1}^K \alpha_k \mathcal{A}P_{\theta_k}$ of a sparse distribution is a combination of *atoms* selected from the *dictionary* $\{\mathcal{A}P_\theta\}_{\theta \in \mathcal{T}}$ indexed by the parameter θ . Table 1 summarizes the notations used in the context of compressive mixture estimation and their correspondence with more classical notions from compressive sensing.

²Note that this representation might not be unique.

	Usual compressive sensing	Compressive mixture estimation
Signal	$\mathbf{x} \in \mathbb{R}^d$	$P \in \mathcal{P}$
Model	K -sparse vectors	K -sparse mixtures $P_{\Theta, \alpha} = \sum_{k=1}^K \alpha_k P_{\theta_k}$
Measurement operator	$\mathbf{M} \in \mathbb{R}^{m \times d}$	$\mathcal{A} : \mathcal{P} \rightarrow \mathbb{C}^m$
Support	$\Gamma(\mathbf{x}) = \{\ell \mid x_\ell \neq 0\}$	$\Gamma(P_{\Theta, \alpha}) = \Theta = \{\theta_1, \dots, \theta_K\}$
Dictionary of atoms	$\{\mathbf{M}e_\ell\}_{\ell=1, \dots, d}$	$\{\mathcal{A}P_\theta\}_{\theta \in \mathcal{T}}$

Table 1: Correspondance between objects manipulated in usual compressive sensing of finite-dimensional signals and in the proposed compressive mixture estimation framework.

2.2 Principle for reconstruction: moment matching

As mentioned in Section 1, usual reconstruction algorithms (also known as *decoders* [37, 17]) in generalized compressive sensing are designed with the purpose of minimizing the measurement error while enforcing sparsity [13], as formulated in equation (5). Here it also corresponds to traditional parametric optimization in the Generalized Method of Moments (GeMM) [60]:

$$\operatorname{argmin}_{\Theta, \alpha} \|\hat{\mathbf{z}} - \mathcal{A}P_{\Theta, \alpha}\|_2, \quad (9)$$

where $\hat{\mathbf{z}} = \mathcal{A}\hat{P} = \frac{1}{n} \sum_{i=1}^n \mathcal{A}\delta_{\mathbf{x}_i}$ is the empirical sketch. This problem is usually highly non-convex and does not allow for an efficient direct optimization, nevertheless we show in Section 7 that in some cases it yields a decoder robust to modeling errors and empirical estimation errors, with high probability.

Convex relaxations of (9) based on sparsity-promoting penalization terms [22, 11, 34, 80] can be envisioned in certain settings, however their direct adaptation to general uncountable dictionaries of atoms (*e.g.*, with GMMs) seems difficult. The main alternative is greedy algorithms. Using an algorithm inspired by Iterative Hard Thresholding (IHT) [14], Bourrier *et al.* [19] estimate mixtures of isotropic Gaussian distributions with fixed variance, using a sketch formed by sampling the empirical characteristic function. As will be shown in Section 5.6, this IHT-like algorithm often yields an unsatisfying local minimum of (9) when the variance is estimated. Instead, we propose a greedy approach similar to Orthogonal Matching Pursuit (OMP) [73, 79] and its extension OMP with Replacement (OMPR) [64].

Another intuitive solution would be to discretize the space of the parameter θ to obtain a finite dictionary of atoms and apply the classic convex relaxation or greedy methods mentioned above. However, one quickly encounters the well-known curse of dimensionality: for a grid with precision ε and a parameter of dimension p , the size of the dictionary is as $\mathcal{O}(\varepsilon^{-p})$, which is intractable even for moderate p . Initial experiments for learning GMMs in dimension $d = 2$ with diagonal covariance (*i.e.* the dimension of the parameter θ is $p = 4$) show that this approach is extremely long and has a very limited precision. Instead, in the next section we propose an adaptation of OMPR directly in the continuous domain.

2.3 Inspiration: OMPR for classical compressive sensing

Matching Pursuit [73] and Orthonormal Matching Pursuit [79] deal with general sparse approximation problems. They gradually extend the sparse support by selecting atoms most correlated with the residual signal, until the desired sparsity is attained.

An efficient variation of OMP called OMP with Replacement (OMPR) [64] exhibits better reconstruction guarantees. Inspired by IHT[14], and similar to CoSAMP or Subspace Pursuit [52], it increases the number of iterations of OMP and extends the size of the support *further than* the desired sparsity before reducing it with Hard Thresholding to suppress spurious atoms.

2.4 Proposed algorithms: Compressive Learning OMP/OMPR

Adapting OMPR to the considered compressive mixture estimation framework requires several modifications. We detail them below, and summarize them in Algorithm 1.

Several aspects of this framework must be highlighted:

- **Non-negativity.** Unlike classical compressive sensing, the compressive mixture estimation framework imposes a non-negativity constraint on the weights α , that we enforce at each iteration. Thus Step 1 is modified compared to classical OMPR by replacing the modulus of the correlation by its real part, to avoid negative correlation between atom and residual. Similarly, in Step 4 we perform a Non-Negative Least-Squares (NNLS) [68] instead of a classical Least-Squares.

Algorithm 1: Compressive mixture learning *à la* OMP: CL-OMP ($T = K$) and CL-OMPR ($T = 2K$)

Data: Empirical sketch $\hat{\mathbf{z}}$, sketching operator \mathcal{A} , sparsity K , number of iterations $T \geq K$
Result: Support Θ , weights α
 $\hat{\mathbf{r}} \leftarrow \hat{\mathbf{z}}; \Theta \leftarrow \emptyset;$
for $t \leftarrow 1$ **to** T **do**
 Step 1: Find a normalized atom highly correlated with the residual with a gradient descent
 | $\boldsymbol{\theta} \leftarrow \underset{\boldsymbol{\theta}}{\text{maximize}} \left(\text{Re} \left\langle \frac{\mathcal{A}P_{\boldsymbol{\theta}}}{\|\mathcal{A}P_{\boldsymbol{\theta}}\|_2}, \hat{\mathbf{r}} \right\rangle_2, \text{init} = \text{rand} \right);$
 end
 Step 2: Expand support
 | $\Theta \leftarrow \Theta \cup \{\boldsymbol{\theta}\};$
 end
 Step 3: Enforce sparsity by Hard Thresholding if needed
 | **if** $|\Theta| > K$ **then**
 | | $\boldsymbol{\beta} \leftarrow \arg \min_{\boldsymbol{\beta} \geq 0} \left\| \hat{\mathbf{z}} - \sum_{k=1}^{|\Theta|} \beta_k \frac{\mathcal{A}P_{\boldsymbol{\theta}_k}}{\|\mathcal{A}P_{\boldsymbol{\theta}_k}\|_2} \right\|_2;$
 | | Select K largest entries $\beta_{i_1}, \dots, \beta_{i_K};$
 | | Reduce the support $\Theta \leftarrow \{\boldsymbol{\theta}_{i_1}, \dots, \boldsymbol{\theta}_{i_K}\};$
 | **end**
 end
 Step 4: Project to find weights
 | $\boldsymbol{\alpha} \leftarrow \arg \min_{\boldsymbol{\alpha} \geq 0} \left\| \hat{\mathbf{z}} - \sum_{k=1}^{|\Theta|} \alpha_k \mathcal{A}P_{\boldsymbol{\theta}_k} \right\|_2;$
 end
 Step 5: Perform a gradient descent *initialized with current parameters*
 | $\Theta, \boldsymbol{\alpha} \leftarrow \underset{\Theta, \boldsymbol{\alpha}}{\text{minimize}} \left(\left\| \hat{\mathbf{z}} - \sum_{k=1}^{|\Theta|} \alpha_k \mathcal{A}P_{\boldsymbol{\theta}_k} \right\|_2, \text{init} = (\Theta, \boldsymbol{\alpha}), \text{constraint} = \{\boldsymbol{\alpha} \geq 0\} \right);$
 end
 Update residual: $\hat{\mathbf{r}} \leftarrow \hat{\mathbf{z}} - \sum_{k=1}^{|\Theta|} \alpha_k \mathcal{A}P_{\boldsymbol{\theta}_k}$
end
Normalize $\boldsymbol{\alpha}$ such that $\sum_{k=1}^K \alpha_k = 1$

- **Normalization.** Note that we do not enforce normalization $\sum_{k=1}^K \alpha_k = 1$ at each iteration. Instead, a normalization of α is performed at the end of the algorithm to obtain a valid distribution. Enforcing the normalization constraint at each iteration was found on initial experiments to have a negligible effect while increasing computation time.
- **Continuous dictionary.** The set \mathcal{G} of elementary distributions is often continuously indexed (as with GMMs, see Section 3.2) and cannot be exhaustively searched. Instead we propose to replace the maximization in Step 1 of classical OMPR with a randomly initialized gradient descent, denoted by a call to a sub-routine `maximize $_{\theta}$` , leading to a – local – maximum of the correlation between atom and residual. Note that the atoms are normalized during the search, as is often the case with OMP.
- **Global optimization step to handle coherent dictionaries.** Compared to classical OMPR, the proposed algorithm includes a new step at each iteration (Step 5), which further reduces the cost function with a gradient descent initialized with the *current* parameters (Θ, α) . This is denoted by a call to the sub-routine `minimize $_{\Theta, \alpha}$` . The need for this additional step stems from the lack of incoherence between the elements of the uncountable dictionary. For instance, in the case of GMM estimation, a $(K + 1)$ -GMM approximation of a distribution cannot be directly derived from a K -GMM by simply adding a Gaussian. This is reminiscent of a similar problem handled in High Resolution Matching Pursuit (HRMP) [59], which uses a multi-scale decomposition of atoms, while we handle here the more general case of a continuous dictionary using a global gradient descent that adjusts all atoms. Experiments show that this step is the most time-consuming of the algorithm, but that it is necessary.

Similar to classical OMPR, Algorithm 1 yields two algorithms depending on the number of iterations:

1. **Compressive Learning OMP** (CL-OMP) if run *without* Hard Thresholding (*i.e.* with $T = K$);
2. **CL-OMPR** (with Replacement) if run with $T = 2K$ iterations.

Learning the number of components ? In the proposed framework, the number of components K is known in advance and provided by the user. However, it is known that greedy approaches such as OMP are convenient to derive stopping conditions, that could be readily applied to CL-OMP: when the residual falls below a fixed (or adaptive) threshold, stop the algorithm (additional strategies would be required for CL-OMPR however). In this paper however, we only compare the proposed method with classical approaches such as EM for Gaussian Mixture Models, that consider the number of components K *known in advance*. We leave for future work the implementation of a stopping condition for CL-OMP(R) and comparison with existing methods for model selection.

2.5 Complexity of CL-OMP(R).

Just as OMP, which complexity scales quadratically with the sparsity parameter K , proposed greedy approaches CL-OMP or CL-OMPR have a computational cost of the order of $\mathcal{O}(mdKT)$, where $T \geq K$ is the number of iterations, resulting in a quadratic cost with respect to the number of components K .

This is potentially a limiting factor for the estimation of mixtures of many basic distributions (large K). In classical sparse approximation, approximate least squares approaches such as Gradient Pursuit [13] or LocOMP [71] have been developed to overcome this computational bottleneck. One could probably get inspiration from these approaches to further scale up compressive mixture estimation, however in the context of GMMs we propose in Section 3.2 to rather leverage ideas from existing fast Expectation-Maximization (EM) algorithms that are *specific to GMMs*.

2.6 Sketching by randomly sampling the characteristic function

Let us now assume $X = \mathbb{R}^d$. *The reader will notice that in classical compressive sensing, the compressed object is a vector $\mathbf{x} \in \mathbb{R}^d$, while in this context, a training collection of vectors $\{\mathbf{x}_1, \dots, \mathbf{x}_n\} \subset \mathbb{R}^d$ is considered as an empirical version of some probability distribution $P \in \mathcal{E}(\mathbb{R}^d)$ which is the compressed object.*

The proposed algorithms CL-OMP(R) are suitable for any sketching operator \mathcal{A} and any mixture model of parametric densities P_{θ} , *as long as the optimization schemes in Steps 1 and 5 can be performed.*

In the case of a continuous dictionary the optimization steps can be performed with simple gradient descents implicitly represented by calls to the sub-routines `maximize θ` and `minimize θ, α` , provided $\mathcal{A}P_\theta$ and its gradient with respect to θ have a closed-form expression.

In many important applications such as medical imaging (MRI and tomography), astrophysics or geophysical exploration, one wishes to reconstruct a signal from incomplete samples of its discrete Fourier transform. Random Fourier sampling was therefore one of the first problems to give rise to the classical notions of compressive sensing [28, 27, 29]. Indeed, a random uniform selection of rows of the full Fourier matrix, *i.e.* a random selection of frequencies, forms a partial Fourier matrix that satisfies a certain Restricted Isometry Property (RIP) with overwhelming probability [29], and is therefore appropriate for the encoding and recovery of sparse signals. For more details, we refer the reader to [52, 29, 28, 27, 6] and references therein. Inspired by this methodology, we form the sketch by sampling the characteristic function (*i.e.* the Fourier transform) of the probability distribution P .

The characteristic function ψ_ν of a finite measure $\nu \in E(\mathbb{R}^d)$ is defined as:

$$\psi_\nu(\boldsymbol{\omega}) = \int \left(e^{-i\boldsymbol{\omega}^T \mathbf{x}} \right) d\nu(\mathbf{x}) \quad \forall \boldsymbol{\omega} \in \mathbb{R}^d. \quad (10)$$

For a sparse distribution $P_{\Theta, \alpha}$ (in some given set of basic distributions $\mathcal{G} \subset \mathcal{P}$), we also denote $\psi_{\Theta, \alpha} = \psi_{P_{\Theta, \alpha}}$ for simplicity.

The characteristic function ψ_P of a probability distribution P is a well-known object with many desirable properties. It completely defines any distribution with no ambiguity and often has a closed-form expression (even for distributions which may not have a probability density function with closed-form expression, *e.g.*, for α -stable distributions [91]), which makes it a suitable choice to build a sketch used with CL-OMP(R). It has been used as an estimation tool at an early stage [50] as well as in more recent developments on GeMM [32].

The proposed sketching operator is defined as a sampling of the characteristic function. Given frequencies $\Omega = \{\boldsymbol{\omega}_1, \dots, \boldsymbol{\omega}_m\}$ in \mathbb{R}^d , we define generalized moment functions:

$$M_j(\mathbf{x}) = \exp(-i\boldsymbol{\omega}_j^T \mathbf{x}), \quad j = 1 \dots m, \quad (11)$$

and the sketching operator (2) is therefore expressed as

$$\mathcal{A}\nu = \frac{1}{\sqrt{m}} [\psi_\nu(\boldsymbol{\omega}_1), \dots, \psi_\nu(\boldsymbol{\omega}_m)]^T. \quad (12)$$

Given a training collection $\mathcal{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ in \mathbb{R}^d , we denote $\hat{\psi}(\boldsymbol{\omega}) = \frac{1}{n} \sum_{i=1}^n e^{-i\boldsymbol{\omega}^T \mathbf{x}_i}$ the empirical characteristic function³. The empirical sketch $\hat{\mathbf{z}} = \mathcal{A}\hat{P}$ is

$$\hat{\mathbf{z}} = \frac{1}{\sqrt{m}} [\hat{\psi}(\boldsymbol{\omega}_1), \dots, \hat{\psi}(\boldsymbol{\omega}_m)]^T. \quad (13)$$

To fully specify the sketching operator (12), one needs to choose the frequencies $\boldsymbol{\omega}_j$ that define it. In the spirit of Random Fourier Sampling, we propose to define a probability distribution $\Lambda \in \mathcal{P}$ to draw $(\boldsymbol{\omega}_1, \dots, \boldsymbol{\omega}_m) \stackrel{i.i.d.}{\sim} \Lambda$. Choosing this distribution is a problem of its own that will be discussed in details in Section 3.3.

Connections with Random Neural Networks. It is possible to draw connections between the proposed sketching operation and neural networks. Denote $\mathbf{W} := [\boldsymbol{\omega}_1, \dots, \boldsymbol{\omega}_m] \in \mathbb{R}^{d \times m}$ and $\mathbf{X} := [\mathbf{x}_1, \dots, \mathbf{x}_n] \in \mathbb{R}^{d \times n}$. To derive the sketch, one needs to compute the matrix $\mathbf{U} := \mathbf{W}^T \mathbf{X} \in \mathbb{R}^{m \times n}$, take the complex exponential of each individual entry $\mathbf{V} := \bar{f}(\mathbf{U})$ where \bar{f} is the pointwise application of the function $f(x) = e^{-ix}/\sqrt{m}$ and finally pool the columns of \mathbf{V} to obtain the empirical sketch $\hat{\mathbf{z}} = (\sum_{i=1}^n \mathbf{v}_i)/n$. This procedure indeed shares similarities with a 1-layer neural network: the output $\mathbf{y} \in \mathbb{R}^m$ of such a network can be expressed as $\mathbf{y} = f(\mathbf{W}^T \mathbf{x})$, where $\mathbf{x} \in \mathbb{R}^d$ is the input signal, f is a pointwise non-linear function and $\mathbf{W} \in \mathbb{R}^{d \times m}$ is some weighting matrix. Therefore, in the proposed framework, such a 1-layer network is applied to many inputs \mathbf{x}_i , then the empirical average of the outputs \mathbf{z}_i is taken to obtain a representation of the underlying distribution of the \mathbf{x}_i .

Neural networks with weights \mathbf{W} chosen at random rather than learned on training data – as is done with the frequencies $\boldsymbol{\omega}_j$ – have been studied in the so-called *Random Kitchen Sinks* [83] or in the context

³Other more robust estimators can be envisioned such as the empirical median. The empirical average allows more easy streaming or distributed computing.

of Deep Neural Networks (DNN) with Gaussian weights [56]. In the latter, they have been shown to perform a stable embedding of the input \mathbf{x} when it lives on a low-dimensional manifold. Similar to [56], we show in Section 7 that with high probability the sketching procedure is a stable embedding of the probability distribution of \mathbf{x} when this distribution belongs to, *e.g.*, a compact manifold.

2.7 Complexity of sketching

The main computational load of the sketching operation is the computation of the matrix $\mathbf{U} = \mathbf{W}^T \mathbf{X}$, which theoretically scales in $\mathcal{O}(dmn)$. Large multiplications by random matrices are indeed a well-known computational bottleneck in Compressive Sensing, and some methods circumvent this issue by using approximated fast transforms [45, 70]. Closer to our work (see Section 4), fast transforms have also been used in the context of Random Fourier Features and kernel methods [69, 108]. We leave the study of a possible adaptation of these acceleration methods for future work, and focus on simple practical remarks about the computation of the sketch.

GPU computing. Matrix multiplication is one of the most studied problem in the context of large-scale computing. A classical way to drastically reduce its cost consists in using GPU computing [104]. Recent architectures can even leverage giant matrix multiplication using multiple GPUs in parallel [109].

Distributed/online computing. The computation of the sketch can also be performed in a distributed manner. One can divide the database \mathcal{X} in T subsets \mathcal{X}_t containing n_t items respectively, after which individual computing units can compute the sketches $\hat{\mathbf{z}}_t$ of each subset \mathcal{X}_t in parallel, using the same frequencies. Those sketches are then easily merged⁴ as $\hat{\mathbf{z}} = \sum_{t=1}^T n_t \hat{\mathbf{z}}_t / \sum_{t=1}^T n_t$. The cost of computing the sketch is thus divided by the number of units T . Similarly, this simple observation allows the sketch to be computed in an online fashion.

3 Sketching Gaussian Mixture Models

In this section, we instantiate the proposed framework in the context of Gaussian Mixture Models (GMMs). A more scalable algorithm specific to GMMs is first introduced as a possible alternative to CL-OMP(R). We then focus on the design of the sketching operator \mathcal{A} , *i.e.* on the design of the frequency distribution Λ (see Section 2.6).

3.1 Gaussian Mixture Models

In the GMM framework, the basic distributions $P_{\theta} \in \mathcal{G}$ are Gaussian distributions with density functions denoted p_{θ} :

$$p_{\theta}(\mathbf{x}) = \mathcal{N}(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma}) = \frac{1}{(2\pi)^{d/2} |\boldsymbol{\Sigma}|^{1/2}} \exp\left(-\frac{1}{2}(\mathbf{x} - \boldsymbol{\mu})^T \boldsymbol{\Sigma}^{-1}(\mathbf{x} - \boldsymbol{\mu})\right), \quad (14)$$

where $\theta = (\boldsymbol{\mu}, \boldsymbol{\Sigma})$ represents the parameters of the Gaussian with mean $\boldsymbol{\mu} \in \mathbb{R}^d$ and covariance $\boldsymbol{\Sigma} \in \mathbb{R}^{d \times d}$. A Gaussian is said to be isotropic, or spherical, if its covariance matrix is proportional to the identity: $\boldsymbol{\Sigma} = \sigma^2 \mathbf{I}$.

Densities of GMMs are denoted $p_{\Theta, \boldsymbol{\alpha}} = \sum_{k=1}^K \alpha_k p_{\theta_k}$. A K -GMM is then naturally parametrized by weight vector $\boldsymbol{\alpha} \in \mathbb{R}^K$ and parameters $\theta_k = (\boldsymbol{\mu}_k, \boldsymbol{\Sigma}_k)$, $k = 1 \dots K$.

Compressive density estimation with mixtures of isotropic Gaussians with fixed, known variance was considered in [19]. In this work, we consider mixtures of Gaussians with diagonal covariances, which is known to be sufficient for many applications [86, 111] and is the default framework in well-known toolboxes such as VLFeat [102]. We denote $(\sigma_{k,1}^2, \dots, \sigma_{k,d}^2) = \text{diag}(\boldsymbol{\Sigma}_k)$. Depending on the context, we equivalently denote $\theta_k = [\boldsymbol{\mu}_k; \boldsymbol{\sigma}_k] \in \mathbb{R}^{2d}$ where $\boldsymbol{\sigma}_k = [\sigma_{k,\ell}^2]_{\ell=1 \dots d}$ is the diagonal of the covariance of the k -th Gaussian.

The characteristic function of a Gaussian P_{θ} has a closed-form expression:

$$\psi_{\theta}(\boldsymbol{\omega}) = \exp(-i\boldsymbol{\omega}^T \boldsymbol{\mu}) \exp\left(-\frac{1}{2}\boldsymbol{\omega}^T \boldsymbol{\Sigma} \boldsymbol{\omega}\right), \quad (15)$$

⁴A similar strategy can also be used on a single machine when the matrix \mathbf{U} is too large to be stored

from which we can easily derive the expression of the gradients necessary to perform the optimization schemes maximize_{θ} , $\text{minimize}_{\theta, \alpha}$ in CL-OMP(R), with the sketching operator introduced in Section 2.6.

3.2 A faster algorithm specific to GMM estimation

To handle mixtures of many Gaussians (large K), the fast hierarchical EM used in [90] alternates between binary splits of each Gaussian k along its first principal component (in the case of Gaussians with diagonal covariance, this is the dimension ℓ with the highest variance $\sigma_{k,\ell}^2$) and a few EM iterations.

Our compressive adaptation is summarized in Algorithm 2. The binary split is performed by calling the function **Split**, and the EM steps are replaced by **Step 5** of Algorithm 1 to adjust all Gaussians. In the case where the targeted sparsity level K is not a power of 2, we split the GMM until the support reaches a size $2^{\lceil \log_2 K \rceil} > K$, then reduce it with a Hard Thresholding (Step 3 of Algorithm 1), similar to CL-OMPR.

Since the number of iterations in Algorithm 2 is $T = \lceil \log_2 K \rceil$, the computational cost of this algorithm scales in $\mathcal{O}(mdK \log K)$, which is much faster for large K than the quadratic cost of CL-OMPR.

In practice, Algorithm 2 is seen to sometimes yield worse results than CL-OMPR (see Section 5.6), but be well-adapted to other tasks such as, *e.g.*, GMM estimation for large-scale speaker verification (see Section 6).

<p>Function Split(Θ): split each Gaussian in the support along its dimension of highest variance</p> <p>Data: Support $\Theta = \{\theta_1, \dots, \theta_K\}$ where $\theta_k = [\mu_k; \sigma_k]$</p> <p>Result: New support Θ^{new} of size $\Theta^{new} = 2K$</p> <p>$\Theta^{new} \leftarrow \emptyset$;</p> <p>for $k \leftarrow 1$ to K do</p> <p> $\ell \leftarrow \arg \max_{j \in [1,d]} \sigma_{k,j}^2$;</p> <p> $\Theta^{new} \leftarrow \Theta^{new} \cup \{(\mu_k - \sigma_{k,\ell} \mathbf{e}_\ell, \Sigma_k), (\mu_k + \sigma_{k,\ell} \mathbf{e}_\ell, \Sigma_k)\}$;</p> <p>end</p>

<p>Algorithm 2: An algorithm with complexity $\mathcal{O}(K \log K)$ for compressive GMM estimation</p> <p>Data: Sketch $\hat{\mathbf{z}}$, sketching operator \mathcal{A}, sparsity K</p> <p>Result: Support Θ, weights α</p> <p>$\Theta \leftarrow \emptyset$;</p> <p>begin Initialize with <i>one</i> atom highly correlated with the sketch</p> <p> $\Theta \leftarrow \left\{ \text{maximize}_{\theta} \left(\text{Re} \left\langle \frac{\mathcal{A}P_{\theta}}{\ \mathcal{A}P_{\theta}\ _2}, \hat{\mathbf{z}} \right\rangle_2, \text{init} = \text{rand} \right) \right\}$;</p> <p>end</p> <p>for $t \leftarrow 1$ to $\lceil \log_2 K \rceil$ do</p> <p> begin Split each Gaussian in the support along its dimension of highest variance</p> <p> $\Theta \leftarrow \text{Split}(\Theta)$;</p> <p> end</p> <p> Perform Step 3, Step 4 and Step 5 of Algorithm 1 ;</p> <p>end</p> <p>Normalize α such that $\sum_{k=1}^K \alpha_k = 1$</p>

3.3 Designing the frequency sampling pattern

A key ingredient in designing the sketching operator \mathcal{A} is the choice of a probability distribution Λ to draw the frequency sampling pattern $\{\omega_1, \dots, \omega_m\}$ as defined in Section 2.6. We show in Section 4 that this corresponds to the design of a translation-invariant kernel in the data domain. Interestingly, working in the Fourier domain seems to make the heuristic design strategy more direct. Literature on designing kernels in this context usually focus on maximizing the distance between the sketch of two distributions [97, 77], which cannot be readily applied in our context since we sketch only a single

distribution. However, as we will see, the proposed approach follows the general idea of maximizing the capacity of the sketch to distinguish this distribution from others, which amounts to maximizing the variations of the sketch with respect to the parameters of the GMM at the selected frequencies.

3.3.1 Oracle frequency sampling pattern for a single known Gaussian

We start by designing a heuristic for choosing frequencies adapted to the estimation of a single Gaussian P_θ , assuming the parameters $\theta = (\boldsymbol{\mu}, \boldsymbol{\Sigma})$ are available – which is obviously not the case in practice. We will deal in due time with mixtures, and with unknown parameters.

Gaussian frequency distribution. Recall the expression (15) of the characteristic function $\psi_\theta(\boldsymbol{\omega}) = e^{-i\boldsymbol{\omega}^T \boldsymbol{\mu} - \frac{1}{2}\boldsymbol{\omega}^T \boldsymbol{\Sigma} \boldsymbol{\omega}}$ of the Gaussian P_θ . It is an oscillating function with Gaussian amplitude of inverted variance with respect to the original Gaussian. Given that $|\psi_\theta| \propto \mathcal{N}(0, \boldsymbol{\Sigma}^{-1})$, choosing a Gaussian frequency distribution $\Lambda_\Sigma^{(G)} = \mathcal{N}(0, \boldsymbol{\Sigma}^{-1})$ is a possible intuitive choice [19, 20] to sample the characteristic function ψ_θ . It concentrates frequencies in the regions where the sampled characteristic function has high amplitude.

However, points drawn from a high-dimensional Gaussian concentrate on an ellipsoid which moves away from the origin as the dimension d increases. Such a Gaussian sampling therefore “undersamples” low or even middle frequencies. This phenomenon has long been one of the reasons for using dimensionality reduction for GMM estimation [44]. Hence, in high dimension the amplitude of the characteristic function becomes negligible (with high probability) at all selected frequencies.

Folded Gaussian radial frequency distribution. In light of the problem observed with the Gaussian frequency distribution, we propose to draw frequencies from a distribution allowing for an accurate control of the quantity $\boldsymbol{\omega}^T \boldsymbol{\Sigma} \boldsymbol{\omega}$, and thus of the amplitude $e^{-\frac{1}{2}\boldsymbol{\omega}^T \boldsymbol{\Sigma} \boldsymbol{\omega}}$ of the characteristic function. This is achieved by drawing

$$\boldsymbol{\omega} = R\boldsymbol{\Sigma}^{-\frac{1}{2}}\boldsymbol{\varphi}, \quad (16)$$

where $\boldsymbol{\varphi} \in \mathbb{R}^d$ is uniformly distributed on the ℓ_2 unit sphere \mathcal{S}_{d-1} , and $R \in \mathbb{R}_+$ is a radius chosen independently from $\boldsymbol{\varphi}$ with a distribution p_R we will now specify.

With the decomposition (16), the characteristic function ψ_θ is now expressed as

$$\psi_\theta \left(R\boldsymbol{\Sigma}^{-\frac{1}{2}}\boldsymbol{\varphi} \right) = \exp \left(-iR\boldsymbol{\varphi}^T \boldsymbol{\Sigma}^{-\frac{1}{2}}\boldsymbol{\mu} \right) \exp \left(-\frac{1}{2}R^2 \right) = \psi_\theta(R),$$

where ψ_θ is the characteristic function of a *one-dimensional* Gaussian with mean $\boldsymbol{\mu} = \boldsymbol{\varphi}^T \boldsymbol{\Sigma}^{-\frac{1}{2}}\boldsymbol{\mu}$ and variance $\sigma^2 = 1$. We thus consider the estimation of a one-dimensional Gaussian $P_{\theta_0} = \mathcal{N}(\mu_0, \sigma_0^2)$, with $\sigma_0 = 1$, as our baseline to design a radius distribution p_R .

In this setting, we no longer suffer from unwanted concentration phenomena and can resort to the intuitive Gaussian radius distribution to sample ψ_{θ_0} . It corresponds to a radius density function $p_R = \mathcal{N}^+(0, \frac{1}{\sigma_0^2}) = \mathcal{N}^+(0, 1)$ (*i.e.* Gaussian with absolute value, referred to as *folded* Gaussian). Using this radius distribution with the decomposition (16) yields a frequency distribution $\Lambda_\Sigma^{(FGr)}$ referred to as *Folded Gaussian radius* frequency distribution. Note that, similar to the Gaussian frequency distribution, the Folded Gaussian radius distribution only depends on the (oracle) covariance $\boldsymbol{\Sigma}$ of the sketched distribution P_θ .

Adapted radius distribution Though we will see it yields decent results in Section 5, the Folded Gaussian radius frequency distribution somehow produces too many frequencies with a low radius R . These carry a limited quantity of information about the original distribution, since all characteristic functions equal 1 at the origin⁵. We now present a heuristics that may avoid this “waste” of frequencies.

Intuitively, the chosen frequencies should properly discriminate Gaussians with different parameters, at least in the neighborhood of the true parameter $\theta_0 = (\mu_0, \sigma_0) = (\mu_0, 1)$. This corresponds to promoting frequencies $\boldsymbol{\omega}$ leading to a large difference $|\psi_\theta(\boldsymbol{\omega}) - \psi_{\theta_0}(\boldsymbol{\omega})|$ for parameters θ close to θ_0 . A way to achieve this is to promote frequencies where the norm of the gradient $\|\nabla_\theta \psi_\theta(\boldsymbol{\omega})\|_2$ is large at $\theta = \theta_0$.

⁵In a way, numerous measures of the characteristic function near the origin essentially measure its derivatives at various orders, which are associated to classical polynomial moments.

Recall that for a one-dimensional Gaussian $\psi_\theta(\omega) = e^{-i\mu\omega} e^{-\frac{1}{2}\sigma^2\omega^2}$. The norm of the gradient is expressed as:

$$\|\nabla_\theta \psi_\theta(\omega)\|_2^2 = |\nabla_\mu \psi_\theta(\omega)|^2 + |\nabla_{\sigma^2} \psi_\theta(\omega)|^2 = |-i\omega \psi_\theta(\omega)|^2 + \left| -\frac{1}{2}\omega^2 \psi_\theta(\omega) \right|^2 = \left(R^2 + \frac{R^4}{4} \right) e^{-\sigma^2 R^2}$$

and therefore $\|\nabla_\theta \psi_{\theta_0}(\omega)\|_2 = \left(R^2 + \frac{R^4}{4} \right)^{\frac{1}{2}} e^{-\frac{1}{2}R^2}$ since $\sigma_0^2 = 1$. This expression still has a Gaussian decrease (up to polynomial factors), and indeed avoids very low frequencies. It can be normalized to a density function:

$$p_R = C \left(R^2 + \frac{R^4}{4} \right)^{\frac{1}{2}} e^{-\frac{1}{2}R^2}, \quad (17)$$

with C some normalization constant. Using this radius distribution with the decomposition (16) yields a distribution $\Lambda_\Sigma^{(Ar)}$ referred to as *Adapted radius* frequency distribution. Once again, this distribution only depends on the covariance Σ .

3.3.2 Oracle frequency sampling pattern for a known mixture of Gaussians

Any frequency distribution $\Lambda_\Sigma^{(\cdot)}$ selected for sampling the characteristic function of a single known Gaussian P_θ can be immediately and naturally extended to a frequency distribution $\Lambda_{\Theta, \alpha}^{(\cdot)}$ to sample the characteristic function of a known GMM $P_{\Theta, \alpha}$, by mixing the frequency distributions corresponding to each Gaussian:

$$\Lambda_{\Theta, \alpha}^{(\cdot)} = \sum_{k=1}^K \alpha_k \Lambda_{\Sigma_k}^{(\cdot)}. \quad (18)$$

Each component $\Lambda_{\Sigma_k}^{(\cdot)}$ has the same weight than its corresponding Gaussian P_{θ_k} . Indeed, a Gaussian with a high weight must be precisely estimated, as its influence on the overall reconstruction error (*e.g.* in terms of Kullback-Leibler divergence) is more important than the components with low weights. Thus more frequencies adapted to this Gaussian are selected.

The draw of frequencies with an oracle distribution $\Lambda_{\Theta, \alpha}^{(\cdot)}$ is summarized in Function `DrawFreq`.

3.3.3 Choosing the frequency sampling pattern in practice

In practice the parameters (Θ, α) of the GMM to be estimated are obviously unknown beforehand, so the oracle frequency distributions $\Lambda_{\Theta, \alpha}^{(\cdot)}$ introduced in the previous section cannot be computed. We propose a simple method to obtain an approximate distribution that yields good results in practice. The reader must also keep in mind that it is very easy to integrate some prior knowledge in this design, especially since the proposed frequency distributions only take into account the variances of the GMM components, not their means.

The idea is to estimate the average variance $\bar{\sigma}^2 = \frac{1}{Kd} \sum_{k=1}^K \sum_{\ell=1}^d \sigma_{k, \ell}^2$ of the components in the GMM – note that this parameter may be significantly different from the global variance of the data, for instance in the case of well-separated components with small variances. This estimation is performed using a light sketch \mathbf{z}_0 with m_0 frequencies, computed on a small subset of n_0 items from the database, then a frequency distribution corresponding to a single isotropic Gaussian $\Lambda_{\bar{\sigma}^2 \mathbf{I}}^{(\cdot)}$ is selected.

Indeed, if the variances $\sigma_{k, \ell}^2$'s are not too different from each other, the amplitude of the empirical characteristic function $|\hat{\psi}(\boldsymbol{\omega})|$ *approximately* follows $e^{-\frac{1}{2}\|\boldsymbol{\omega}\|_2^2 \bar{\sigma}^2}$ with high oscillations, allowing for a very simple amplitude estimation process: assuming the m_0 frequencies used to compute the sketch \mathbf{z}_0 are ordered by increasing Euclidean radius, the sketch \mathbf{z}_0 is divided into consecutive blocks, maximal peaks of its modulus are identified within each block forming a curve that approximately follows $e^{-\frac{1}{2}R^2 \bar{\sigma}^2}$, then a simple regression is used to estimate $\bar{\sigma}^2$. This process is illustrated in Figure 2.

To cope with the fact that the "range" of frequencies that must be considered to compute \mathbf{z}_0 is also not known beforehand, we initialize $\bar{\sigma}^2 = 1$ and reiterate this procedure several times, each time drawing m_0 frequencies adapted to the current estimation of $\bar{\sigma}^2$, *i.e.* with some choice of frequency distribution $\Lambda_{\bar{\sigma}^2 \mathbf{I}}^{(\cdot)}$, and update $\bar{\sigma}^2$ at each iteration. In practice, the procedure quickly converges in three iterations.

The entire process is summarized in detail in Function `EstimMeanSigma` and Algorithm 3.

Function DrawFreq($\{\Sigma_k\}_{k=1,\dots,K}, \alpha, m, f$): drawing frequencies for a GMM with known variances and weights, choosing one of the three distributions described in Section 3.3

Data: Set of variances and weights of a GMM $\{\Sigma_k\}_{k=1,\dots,K}, \alpha$, number of frequencies m , type of frequency distribution $f \in \{(G), (FGr), (Ar)\}$

Result: Set of frequencies $\Omega = \{\omega_1, \dots, \omega_m\}$

for $j \leftarrow 1$ **to** m **do**

 Draw a label according to the weights of the GMM $k_j \sim \sum_{k=1}^K \alpha_k \delta_k$;

if $f = (G)$ **then**

$\omega_j \sim \mathcal{N}(0, \Sigma_{k_j}^{-1})$; // Gaussian

end

else

 Draw a direction $\varphi \sim \mathcal{U}(\mathcal{S}_{d-1})$;

if $f = (FGr)$ **then**

$R \sim \mathcal{N}^+(0, 1)$; // Folded Gaussian radius

end

else if $f = (Ar)$ **then**

$R \sim p_R$ with p_R defined by (17); // Adapted radius

end

$\omega_j \leftarrow R \Sigma_{k_j}^{-\frac{1}{2}} \varphi$;

end

end

Function EstimMeanSigma($\mathcal{X}, n_0, m_0, c, T$): Estimation of the mean variance $\bar{\sigma}^2$

Data: Dataset $\mathcal{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$, small number of items $n_0 \leq n$, small number of frequencies m_0 , number of blocks $c \in \mathbb{N}_+^*$, number of iterations T

Result: Estimated mean variance $\bar{\sigma}^2$

begin Initialize

$\bar{\sigma}^2 \leftarrow 1$;

end

for $t \leftarrow 1$ **to** T **do**

begin Draw some frequencies adapted to the current $\bar{\sigma}^2$

$\{\omega_1, \dots, \omega_{m_0}\} \leftarrow \text{DrawFreq}(\bar{\sigma}^2 \mathbf{I}, 1, m_0, (Ar))$;

 Sort the frequencies $\{\omega_1, \dots, \omega_{m_0}\}$ by increasing radius $\|\omega_j\|_2$;

end

begin Compute small empirical sketch, without $\frac{1}{\sqrt{m_0}}$ normalization (Figure 2, left)

$\hat{\mathbf{z}}_0 \leftarrow \left[\frac{1}{n_0} \sum_{i=1}^{n_0} e^{-i\omega_j^T \mathbf{x}_i} \right]_{j=1 \dots m_0}$;

end

begin Divide sketch into blocks, find maximum peak in each block (Figure 2, center)

$s \leftarrow \lfloor m_0/c \rfloor$;

for $q \leftarrow 1$ **to** c **do**

$j_q = \arg \max_{j \in [(q-1)s+1; qs]} |\hat{z}_{0,j}|$;

end

end

begin Update $\bar{\sigma}^2$ (Figure 2, right)

$\hat{\mathbf{e}} \leftarrow [\hat{z}_{0,j_q}]_{q=1 \dots c}$;

$\bar{\sigma}^2 = \arg \min_{\sigma^2 > 0} \left\| \hat{\mathbf{e}} - \left[e^{-\frac{1}{2} R_{j_q}^2 \sigma^2} \right]_{q=1 \dots c} \right\|_2$;

end

end




Figure 2: Estimation of $\bar{\sigma}^2$ (Function `EstimMeanSigma`), for $d = 10$, $K = 5$, $m_0 = 500$ and $n_0 = 5000$. Green: modulus of the sketch with respect to the norm of frequencies (ordered by increasing radius). Blue: visualization of the peaks in each block of 20 consecutive values. Red: fitted curve $e^{-\frac{1}{2}R^2\bar{\sigma}^2}$ for the estimated $\bar{\sigma}^2$.

Algorithm 3: Draw the frequencies in practice.

Data: Dataset $\mathcal{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$, number of frequencies m , type of frequency distribution $f \in \{(G), (FG\text{r}), (Ar)\}$, parameters for the estimation of the mean variance (n_0, m_0, c, T)
Result: Set of frequencies $\Omega = \{\omega_1, \dots, \omega_m\}$
begin Estimate the mean variance
 | $\bar{\sigma}^2 \leftarrow \text{EstimMeanSigma}(\mathcal{X}, n_0, m_0, c, T);$
end
begin Draw frequencies with the distribution $\Lambda_{\bar{\sigma}^2 \mathbf{I}}^{(\cdot)}$
 | $\Omega \leftarrow \text{DrawFreq}(\bar{\sigma}^2 \mathbf{I}, 1, m, f);$
end

3.4 Summary

At this point, all procedures necessary for compressive GMM estimation have been defined. Given a database $\mathcal{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$, a number of measurements m and a number of components K , the entire process is as follow:

- In the absence of prior knowledge, draw m frequencies using Algorithm 3 on (a fraction of) the dataset. The proposed Adapted radius frequency distribution is recommended, other parameters have default values (see Section 5.3), the effect of modifying them has been found to be negligible.
- Compute the empirical sketch $\hat{\mathbf{z}} = \frac{1}{\sqrt{m}} \left[\frac{1}{n} \sum_{i=1}^n e^{-i\omega_j^T \mathbf{x}_i} \right]_{j=1 \dots m}$. One may use GPU and/or distributed/online computing.
- Estimate a K -GMM using CL-OMP(R) or the more scalable but less precise Algorithm 2.

Connections with Distilled sensing. The reader may note that designing a measurement operator *adapted* to some particular data does not fit the classical paradigm of compressive sensing.

The two-stage approaches used to choose the frequency distribution presented above can be related to a line of work referred to as adaptive (or *distilled*) sensing [61], in which a portion of the computational budget is used to crudely design the measurement operator while the rest is used to actually measure the signal. Most often these methods are extended to multi-stage approaches, where the measurement operator is refined at each iteration, and have been used in machine learning [38] or signal processing [7]. Allocating the resources and choosing between *exploration* (designing the measurement operator) and *precision* (actually measuring the signal) is a classic trade-off in areas such as reinforcement learning or game theory.

4 Kernel design and sketching

It turns out that sketching operators as (12) are intimately related to RKHS embedding of probability distributions [94, 98] and RFFs [82]. The proposed, carefully-designed choice of frequency distribution appears as an innovative method to design a reproducing kernel, which is faster and provides better results than traditional choices in the kernel literature [99], as experimentally shown in Section 5.4. Additionally, approximate RKHS embedding of distributions turns out to be an appropriate framework to derive the information preservation guarantees of Section 7.

4.1 Reproducing Kernel Hilbert Spaces (RKHS)

We refer the reader to Appendix A for definitions related to positive definite (p.d.) kernels and measures.

Let $\kappa : X \times X \rightarrow \mathbb{C}$ be a p.d. kernel. By Moore-Aronszajn Theorem [4], to this kernel is associated a unique Hilbert space $\mathcal{H} \subset \mathbb{C}^X$ that satisfies the following properties: for any $\mathbf{x} \in X$ the function $\kappa(\mathbf{x}, \cdot)$ belongs to \mathcal{H} , and the kernel satisfies the reproducing property $\forall f \in \mathcal{H}, \forall \mathbf{x} \in X, \langle \kappa(\mathbf{x}, \cdot), f \rangle_{\mathcal{H}} = f(\mathbf{x})$. The space \mathcal{H} is referred to as the Reproducing Kernel Hilbert Space (RKHS) associated with the kernel κ . We denote $\langle \cdot, \cdot \rangle_{\mathcal{H}}$ the scalar product of the RKHS \mathcal{H} . We refer the reader to [62] and references therein for a review of RKHS and kernel methods.

We focus here on the space $X = \mathbb{R}^d$ and on *translation-invariant* kernels of the form

$$\kappa(\mathbf{x}, \mathbf{y}) = \mathbf{K}(\mathbf{x} - \mathbf{y}), \quad (19)$$

where $\mathbf{K} : \mathbb{R}^d \rightarrow \mathbb{C}$ is a positive definite function. Translation-invariant positive definite kernels are characterized by the Bochner theorem:

Theorem 1 (Bochner [88], Thm. 1.4.3). *A continuous function $\mathbf{K} : \mathbb{R}^d \rightarrow \mathbb{C}$ is positive definite if and only if it is the Fourier transform of a finite (i.e., $\Lambda(\mathbb{R}^d) = \int_{\mathbb{R}^d} d\Lambda(\boldsymbol{\omega}) < \infty$) nonnegative measure Λ on \mathbb{R}^d , that is:*

$$\mathbf{K}(\mathbf{x}) = \int_{\mathbb{R}^n} e^{-i\boldsymbol{\omega}^T \mathbf{x}} d\Lambda(\boldsymbol{\omega}). \quad (20)$$

This expression implies the normalization $|\kappa(\mathbf{x}, \mathbf{y})| \leq |\kappa(\mathbf{x}, \mathbf{x})| = \Lambda(\mathbb{R}^d)$. Hence, without loss of generality, up to a scaling factor, we suppose $\Lambda(\mathbb{R}^d) = 1$. This means in particular that Λ is a probability distribution on \mathbb{R}^d .

4.2 Hilbert Space Embedding of probability distributions

Embeddings of probability distributions [94, 98] are obtained by defining the Mean Map $\varphi : \mathcal{P} \rightarrow \mathcal{H}$:

$$\varphi(P) = \mathbb{E}_{\mathbf{x} \sim P}(\kappa(\mathbf{x}, \cdot)). \quad (21)$$

Using the Mean Map we can define the following p.d. kernel $\mathcal{K}(P, Q) = \langle \varphi(P), \varphi(Q) \rangle_{\mathcal{H}}$ over the set of probability distributions \mathcal{P} . Note that this expression is also equivalent to the following definition [98, 74]: $\mathcal{K}(P, Q) = \mathbb{E}_{\mathbf{x} \sim P, \mathbf{y} \sim Q}(\kappa(\mathbf{x}, \mathbf{y}))$.

We denote $\gamma_{\kappa}(P, Q) = \|\varphi(P) - \varphi(Q)\|_{\mathcal{H}}$ the pseudometric induced by the Mean Map on the set of probability distributions, often referred to as Maximum Mean Discrepancy (MMD) [99]. Fukumizu *et al.* introduced in [54] the concept of *characteristic kernel*, that is, a kernel κ for which the map φ is injective and the pseudometric γ_{κ} is a true metric. Translation-invariant characteristic kernels are characterized by the following Theorem [98].

Theorem 2 (Sriperumbudur *et al.* [98]). *Assume that $\kappa(\mathbf{x}, \mathbf{y}) = \mathbf{K}(\mathbf{x} - \mathbf{y})$ where \mathbf{K} is a positive definite function on \mathbb{R}^d . Then κ is characteristic if and only if $\text{supp}(\Lambda) = \mathbb{R}^d$, where Λ is defined as (20).*

Many classical translation-invariant kernels (Gaussian, Laplacian....) indeed exhibit a Fourier transform with a support that is equal to \mathbb{R}^d , and are therefore characteristic. This is also the case for all kernels corresponding to the proposed frequency distributions $\Lambda_{\Sigma}^{(\cdot)}$ that we defined in Section 3.3.

4.3 From Kernel design to the design of a frequency sampling pattern

In the case of a translation-invariant kernel, the metric γ_κ can be expressed as $\gamma_\kappa^2(P, Q) = \gamma_\Lambda^2(P, Q)$ where by abuse of notation, for a given frequency distribution Λ , we introduce

$$\gamma_\Lambda^2(P, Q) := \int_{\mathbb{R}^d} |\psi_P(\boldsymbol{\omega}) - \psi_Q(\boldsymbol{\omega})|^2 d\Lambda(\boldsymbol{\omega}), \quad (22)$$

where we recall that $\psi_P(\boldsymbol{\omega})$ is the characteristic function of P . The proof of Theorem 2 is actually based on this reformulation [98].

Hence, given m frequencies $(\boldsymbol{\omega}_1, \dots, \boldsymbol{\omega}_m) \stackrel{i.i.d.}{\sim} \Lambda$ and the corresponding sketching operator \mathcal{A} (12), we can expect that for large enough m , with high probability,

$$\gamma_\Lambda^2(P, Q) = \mathbb{E}_{\boldsymbol{\omega} \sim \Lambda} |\psi_P(\boldsymbol{\omega}) - \psi_Q(\boldsymbol{\omega})|^2 \approx \frac{1}{m} \sum_{j=1}^m |\psi_P(\boldsymbol{\omega}_j) - \psi_Q(\boldsymbol{\omega}_j)|^2 = \|\mathcal{A}P - \mathcal{A}Q\|_2^2. \quad (23)$$

Building on this relation between the metric $\gamma_\Lambda^2(\cdot, \cdot)$ and the sketching operators considered in this paper, we next derive some connections with kernel design. We further exploit these connections to draw a theoretical analysis of the information preservation guarantees of the sketching operator \mathcal{A} in Section 7.

Traditional kernel design. Given some learning task, the selection of an appropriate kernel –known as kernel design– is a difficult, open problem, usually done by cross-validation. Even when using RFFs to leverage the computational gain of explicit embeddings (compared to potentially costly implicit kernels), the considered frequency distribution Λ is usually *derived from* a translation-invariant kernel κ chosen in a parametric family endowed with closed-form expressions for both κ and Λ [82, 99]. A typical example is the Gaussian kernel (chosen for the simplicity of its Fourier transform), which bandwidth is often selected by cross-validation [97, 99].

Spectral kernel design: designing the frequency sampling pattern. Learning an appropriate kernel by directly deriving a spectral distribution of frequencies for Random Fourier Features has been an increasingly popular idea in the last few years. In the field of usual reproducing kernels on finite-dimensional objects, researchers have explored the possibility of modifying the matrix of frequencies to obtain a better approximation quality [106] or to accelerate the computation of the kernel [69]. Both ideas have been exploited for learning an appropriate frequency distribution, often modeled as a mixture of Gaussians [105, 108, 76] or by optimizing weights over a combination of many distributions [93].

In the context of using the Mean Map with Random Features, learning a kernel has been often explored for the two-sample test problem, mainly based on the idea of maximizing the discriminative power of the MMD [97]. Similar to our approach, such methods often divide the database in two parts to learn the kernel then perform the hypothesis test [58, 65], or are done in a streaming context [78]. Variants of the Random Fourier Features have also been used [36].

Compared to these methods, the approach proposed in Section 3.3 is relatively simple, fast to perform, and based on an approximate theoretical expression of the MMD for clustered data instead of a thorough statistical analysis of the problem. In the spirit of our sketching framework, it only requires a tiny portion of the database and is performed in a *completely unsupervised* manner, which is quite different from most existing literature. Furthermore, in this paper we only consider traditional Random Fourier Features for translation-invariant kernels.

Using more exotic kernels and/or adapting more advanced learning methods to our framework is left for future work. Still, in Section 5.4, we empirically show that the estimation of $\bar{\sigma}^2$ (Function `EstimMeanSigma`) is much faster than estimation by cross-validation, and that the Adapted radius distribution $\Lambda_{\bar{\sigma}^2 \mathbf{1}}^{(Ar)}$ performs better than a traditional Gaussian kernel with optimized bandwidth.

5 Experiments with synthetic data

To validate the proposed framework, the algorithms CL-OMP(R) (Algorithm 1) are first extensively tested against synthetic problems for which the true parameters of the GMM are known. Experiments on real data will be conducted in Section 6, with the additional analysis of Algorithm 2. The full Matlab code is available at [67].

5.1 Generating the data

When dealing with synthetic data for various settings, it is particularly difficult to generate problems that are “equally difficult” when varying the dimension d and the level of sparsity K . We opted for a simple heuristic to draw K Gaussians neither too close nor too far from one another, given their variances.

Since an d -dimensional ball with radius r has volume $D_d r^d$ (with D_d a constant depending on d), we consider that any Gaussian which is *approximately isotropic* with variance σ^2 (i.e., $\mathcal{N}(0, \Sigma)$ with $\Sigma \approx \sigma^2 \mathbf{I}$) “occupies” a volume $V_{\sigma^2} = \sigma^d V_1$, where V_1 is a reference volume for $\sigma = 1$.

Variances $\sigma_{k,\ell}^2$ are randomly drawn uniformly from 0.25 to 1.75, so that $\mathbb{E}(\bar{\sigma}^2) = 1$. The means $\boldsymbol{\mu}$ are chosen so that they lie in a ball sufficiently large to accommodate for a volume $K \times V_{\bar{\sigma}^2}$. We therefore choose $\boldsymbol{\mu} \sim \mathcal{N}(0, \sigma_{\boldsymbol{\mu}}^2 \mathbf{I})$ with $\sigma_{\boldsymbol{\mu}}$ that verifies $V_{\sigma_{\boldsymbol{\mu}}^2} = K V_{\bar{\sigma}^2}$, which yields $\sigma_{\boldsymbol{\mu}} = K^{\frac{1}{d}} \bar{\sigma}$, i.e., $\sigma_{\boldsymbol{\mu}} = K^{\frac{1}{d}}$ by considering the expected value of $\bar{\sigma}$. In practice this choice seems to offer problems that are neither too elementary nor too difficult.

5.2 Evaluation measure

To evaluate reconstruction performance when the true distribution p of the data is known, one can resort to the classic Kullback-Leibler divergence $D(P||Q)$ [42]. A symmetric version of the KL-divergence is more comfortable to work with: $d(P, \bar{P}) = D(P||\bar{P}) + D(\bar{P}||P)$, where P is the true distribution and \bar{P} is the estimated one (we still refer to this measure as “KL-divergence”). In our framework, P and \bar{P} are GMMs with density functions denoted p and \bar{p} , hence as in [19] to estimate the KL-divergence in practice we draw $(\mathbf{y}_1, \dots, \mathbf{y}_{n'}) \stackrel{i.i.d.}{\sim} P$ with $n' = 5 \times 10^5$ and compute

$$d(P, \bar{P}) \approx \frac{1}{n'} \sum_{i=1}^{n'} \left[\ln \left(\frac{p(\mathbf{y}_i)}{\bar{p}(\mathbf{y}_i)} \right) + \frac{\bar{p}(\mathbf{y}_i)}{p(\mathbf{y}_i)} \ln \left(\frac{\bar{p}(\mathbf{y}_i)}{p(\mathbf{y}_i)} \right) \right]. \quad (24)$$

5.3 Basic Setup

The basic setup is the same for all following experiments, given data $(\mathbf{x}_1, \dots, \mathbf{x}_n) \in \mathbb{R}^d$ and parameters m, K . We suppose the data to be approximately centered (see Section 5.1).

First, unless specified otherwise (e.g., in Section 5.4), we draw frequencies according to an Adapted radius distribution $\Lambda_{\bar{\sigma}^2 \mathbf{I}}^{(Ar)}$, using Algorithm 3 with parameters $m_0 = 500$, $n_0 = \min(n, 5000)$, $c = 30$, $T = 5$. The empirical sketch $\hat{\mathbf{z}}$ is then computed.

The compressive algorithms are performed with their respective number of iterations. For Step 1 of CL-OMP(R), the gradient descent is initialized with a centered isotropic Gaussian with random variance⁶ $\sigma^2 \sim \mathcal{U}([0.5\bar{\sigma}^2; 1.5\bar{\sigma}^2])$. Furthermore, during all optimization steps in CL-OMP(R) or Algorithm 2, we constrain all variances to be larger than a small 10^{-15} for numerical stability. All continuous optimization schemes in CL-OMP(R) are performed with Stephen Becker’s adaptation of the L-BFGS-B algorithm [23] in C, with Matlab wrappers [8].

We compare our results with an adaptation of the previous IHT algorithm [19, 20] for isotropic Gaussians with fixed variances, in which all optimization steps have been straightforwardly adapted for our non-isotropic framework with relaxed variances. The IHT is performed with 10 iterations, which is the default value in the original paper.

We use the VLFeat toolbox [102] to perform the EM algorithm with diagonal variances. The algorithm is repeated 10 times with random initializations and the result yielding the best log-likelihood is kept. Each run is limited to 100 iterations.

All following reconstruction results are computed by taking the geometric mean over 50 experiments (i.e. the regular mean in logarithmic scale).

5.4 Results: role of the choice of frequency distribution

In this section we compare different choices of frequency distributions. We draw $N = 300000$ items in two settings, respectively low and high dimensional: $d = 2$, $K = 3$ and $d = 20$, $K = 5$. In each setting we construct the sketch with $m = 10K(2d+1)$ frequencies (see Section 5.5). Reconstruction is performed with CL-OMPR.

⁶After trying many variants of initialization strategy, we came to the conclusion that the algorithm is very robust to initialization. This is one of the possible choices.

		(G)	(FGr)	(Ar)
$d = 2, K = 3$	$\Lambda_{\Theta_0, \alpha_0}^{(\cdot)}$	-9.30	-8.90	-9.38
	$\Lambda_{\bar{\sigma}^2 \mathbf{I}}^{(\cdot)}$	-8.93	-8.67	-9.20
$d = 20, K = 5$	$\Lambda_{\Theta_0, \alpha_0}^{(\cdot)}$	3.15	-6.17	-6.48
	$\Lambda_{\bar{\sigma}^2 \mathbf{I}}^{(\cdot)}$	3.41	-5.80	-6.32

Table 2: Log-KL-divergence on synthetic data using the CL-OMPR algorithm, for $m = 10(2d + 1)K$ frequencies and $N = 300\,000$ items. We compare the three proposed frequency distributions: Gaussian [19] (G), Folded Gaussian radius (FGr) or Adapted radius (Ar), using either the oracle distribution defined in Section 3.3.2 or the approximate distribution used in practice, learned with `EstimMeanSigma`.

		Reconstruction result	Est. time
$d = 2, K = 3$	Adapted radius (proposed)	-9.38	$(\bar{\sigma}^2)$ 1.02s
	Gaussian kernel [82, 99]	-9.14	(σ_{best}^2) 23.10s
$d = 20, K = 5$	Adapted radius (proposed)	-6.32	$(\bar{\sigma}^2)$ 0.98s
	Gaussian kernel [82, 99]	-5.18	(σ_{best}^2) 74.10s

Table 3: Log-KL-divergence results on synthetic data using the CL-OMPR algorithm, for $m = 10(2d + 1)K$ frequencies and $n = 300\,000$ items, comparing the proposed Adapted radius distribution $\Lambda_{\bar{\sigma}^2 \mathbf{I}}^{(Ar)}$ and a frequency distribution $\Lambda_{\sigma_{best}^2}^{(Gk)}$ corresponding to a Gaussian kernel (Gk) [82, 99], with a bandwidth σ_{best}^2 selected among 15 values exponentially spread from 10^{-1} to 10^2 yielding the best result in each case. Estimation times for parameters $\bar{\sigma}^2$ and σ_{best}^2 are also given.

In Table 2, we compare the three frequency distributions introduced in Section 3, both with the oracle frequency distribution $\Lambda_{\Theta_0, \alpha_0}^{(\cdot)}$ (*i.e.* using Function `DrawFreq` with the true parameters of the GMM) – we remind the reader that this setting unrealistically assumes that the variances and weights of the GMM are known beforehand –, and with the approximate one $\Lambda_{\bar{\sigma}^2 \mathbf{I}}^{(\cdot)}$ (Algorithm 3). The results show that the Gaussian frequency distribution indeed yields poor reconstruction results in high dimension ($d = 20$), while the Adapted radius frequency distribution outperforms the two others. The use of the approximate $\Lambda_{\bar{\sigma}^2 \mathbf{I}}^{(\cdot)}$ instead of the oracle $\Lambda_{\Theta_0, \alpha_0}^{(\cdot)}$ is shown to have little effect.

In Table 3 we compare the proposed Adapted radius frequency distribution with a frequency distribution $\Lambda = \mathcal{N}(0, \mathbf{I}\sigma_{best}^{-2})$ corresponding to a Gaussian kernel $\kappa(\mathbf{x}, \mathbf{y}) = \exp\left(-\frac{\|\mathbf{x}-\mathbf{y}\|_2^2}{2\sigma_{best}^2}\right)$ [82, 99], where the bandwidth σ_{best}^2 is selected among 15 values exponentially spread from 10^{-1} to 10^2 to yield the best reconstruction results in each setting. It is seen that the Adapted radius distribution $\Lambda_{\bar{\sigma}^2}^{(Ar)}$ outperforms the Gaussian kernel in each case, and the estimation of the parameter $\bar{\sigma}^2$ is significantly faster than the tedious learning of the bandwidth σ_{best}^2 .

We conduct an experiment to determine how robust is the isotropic distribution of frequencies $\Lambda_{\bar{\sigma}^2 \mathbf{I}}^{(Ar)}$ when treating strongly non-isotropic data. We first generate a GMM in dimension $d = 10$ with $K = 10$ components, with the process described in Section 5.3 but with identity covariances. Then the first five dimensions of the parameters are divided by a factor $A > 0$, meaning that: for $\ell = 1, \dots, 5$ and $k = 1, \dots, K$, we do $\mu_{k,\ell} \leftarrow \mu_{k,\ell}/A$ and $\sigma_{k,\ell}^2 \leftarrow \sigma_{k,\ell}^2/A^2$. It simulates data which do not have the same range in each dimension (Fig. 3, top). In Fig. 3 (bottom), for increasing values of A we compare CL-OMPR with the learned frequency distribution $\Lambda_{\bar{\sigma}^2 \mathbf{I}}^{(Ar)}$, CL-OMPR using the ideal frequency distribution $\Lambda_{\Theta_0, \alpha_0}^{(Ar)}$, and EM. It is indeed seen that, in terms of KL-divergence (left), the results produced by CL-OMPR with the learned frequency distribution deteriorates as the non-isotropy of the GMM increases. On the contrary, CL-OMPR with the oracle frequency distribution and EM do not seem to be affected. Hence in the first case the problem lies with the learned frequency distribution and not the CL-OMPR algorithm itself. To further confirm this fact, we examine the MMD γ_Λ with the corresponding frequency distribution Λ in each case. It is seen that CL-OMPR with the learned frequency distribution performs as well as the oracle one, meaning that the MMD defined with an isotropic choice of frequencies $\gamma_{\Lambda_{\bar{\sigma}^2 \mathbf{I}}^{(Ar)}}$ is indeed not adapted for strongly non-isotropic problems, and that despite the ability of CL-OMPR to approximately minimize the cost function (9) (which approximates the MMD, see (23)), it does not produce good results.

In that particular case, the problem could be potentially resolved by a whitening of the data, *e.g.* by

computing the empirical covariance of the fraction of the database used for the frequency distribution design phase, and multiplying each datapoint by its inverse during the sketching phase. However there are of course cases where the proposed methods would be further challenged, for instance if components are flat in a dimension but far apart: the global covariance of the data along that dimension would be large, even if the variance of each Gaussian components is small (these cases are however rarely encountered in practice). As mentioned before (see Sec. 4.3), another solution would be to use more advanced methods for designing the MMD γ_Λ than the proposed simple one. Overall, we leave treatment of strongly non-isotropic data for future work.

Nevertheless, from now on, all experiments are performed with an approximate Adapted radius distribution $\Lambda_{\sigma^2 \mathbf{I}}^{(Ar)}$.

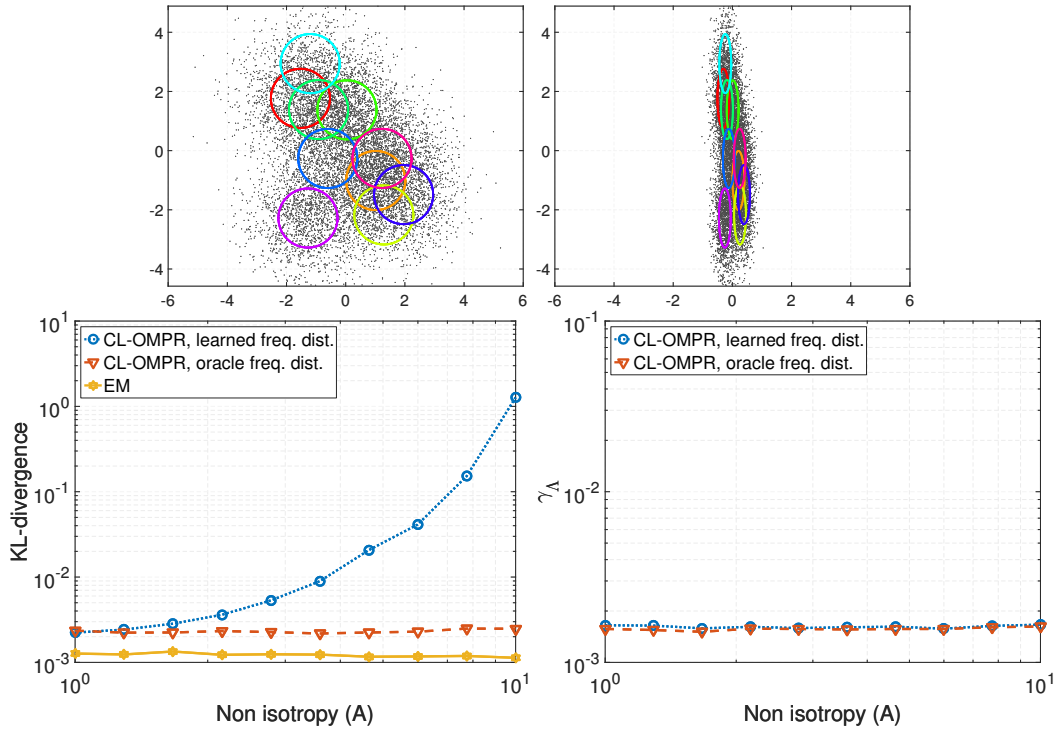


Figure 3: Top: isotropic GMM in dimension $n = 10$ with $K = 10$ components displayed along the first two dimensions (left), the same GMM with one dimension divided by $A = 5$ (right). Bottom: reconstruction results for the KL-divergence (left) or MMD (right) with $n = 2.10^5$ items, using $m = 2000$ frequencies for CL-OMPR, with respect to the coefficient A .

5.5 Results: number of measurements

We now evaluate the quality of the reconstruction with respect to the number m of frequencies, for varying dimension d and number of components K in the GMM (Figure 4).

It is seen that the KL-divergence exhibits a sharp phase-transition with respect to m , whose occurrence seems to be proportional to the “dimension” of the problem, *i.e.* the number of free parameters $(2d + 1)K$. This phenomenon is akin to usual compressive sensing. In light of this observation, the number of frequencies in all following experiments is chosen as $m = 5(2d + 1)K$, beyond the empirically obtained phase transition.

5.6 Results: comparison of the algorithms

We compare the compressive algorithms and EM in terms of reconstruction, computation time and memory usage, with respect to the number of samples N .

Precision of the reconstruction In Figure 5, KL-divergence reconstruction results are shown for EM and all compressed algorithms: IHT [19], CL-OMP (Algorithm 1 with $T = K$), CL-OMPR (Algorithm 1 with $T = 2K$) and Algorithm 2. We consider two settings, one with low $K = 5$ (left) and one with high




Figure 4: Log KL-divergence reconstruction results for CL-OMPR with respect to the normalized number of frequencies $m/((2d+1)K)$ and the dimension d (left) or number of components K (right), using $n = 300\,000$ items. On the left $K = 10$, and on the right $d = 10$.

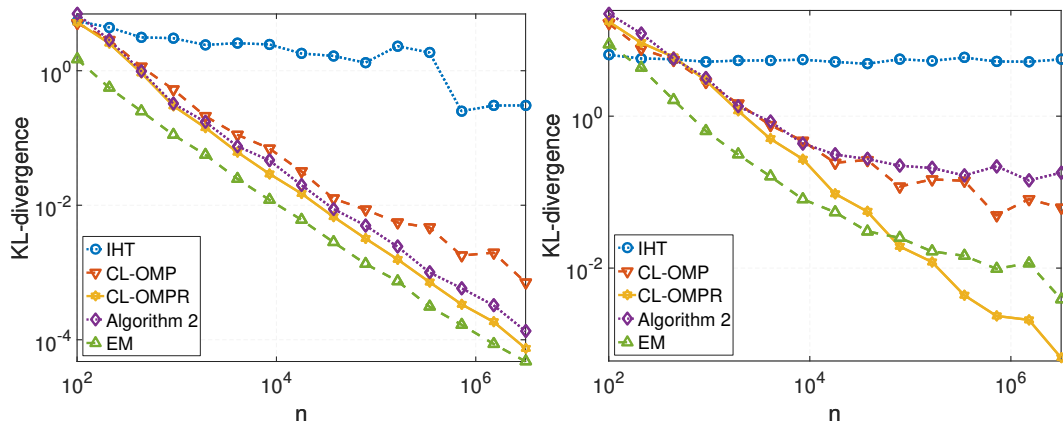


Figure 5: Reconstruction results on synthetic data in dimension $d = 10$, with $K = 5$ component (left) or $K = 20$ (right), and number of frequencies $m = 5(2d+1)K$, with respect to the number of items in the database n .

$K = 20$ (right), in dimension $n = 10$. The number of measurements is set at $m = 5(2d+1)K$ in each case.

With few Gaussians ($K = 5$), all compressive algorithms yield similar results, close to those achieved by EM. The precision of the reconstruction is seen to improve steadily with the size n of the database. With more Gaussians ($K = 20$), CL-OMPR clearly outperforms the other compressive algorithms, and even outperforms EM for very large n .

The IHT algorithm [19] adapted to non-isotropic variances often fails to recover a satisfying GMM. Indeed, IHT fills the support of the GMM at the very first iteration, and is seen to converge toward a local minimum of (9), in which all Gaussians in the GMM are equal to the same large Gaussian that encompasses all data. Note that this situation could not happen in [19], where all Gaussians have fixed, known variance.

Computation time and memory In Figure 6, computation time and memory usage of the compressive methods and EM are presented with respect to the database size n , using an *Intel Core i7-4600U 2.1 GHz* CPU with 8 GB of RAM. In terms of time complexity (resp. memory usage), the EM algorithm scales in $\mathcal{O}(dnKT_{EM})$ for a fixed number of iterations $T_{EM} = 100$ (resp. $\mathcal{O}(nd)$). The CL-OMP or CL-OMPR algorithms scale in $\mathcal{O}(mdK^2)$ (resp. $\mathcal{O}(md)$), while Algorithm 2 scales in $\mathcal{O}(mdK \log K)$ (resp. the same $\mathcal{O}(md)$).

At large n the EM algorithm indeed becomes substantially slower than all compressive methods (Fig 6, left). We also keep in mind that we compare a MATLAB implementation of the compressive methods with a state-of-the-art C++ implementation of EM[102]. Similarly, at large n the compressive algorithms outperform EM by several orders of magnitude in terms of memory usage (Fig 6, right).




Figure 6: Memory (left) and time (right) usage of all algorithms on synthetic data with dimension $n = 10$, number of components $K = 5$, and number of frequencies $m = 5(2d + 1)K$, with respect to the number of items in the database n . On the left, “Sketching” refers to the time of computing the sketch with a naive, direct computation, which must be added to the computation time of the recovery algorithm (that does not vary with n) to obtain the total computation time of the proposed method. However the reader must keep in mind that the sketching step can be massively parallelized, is adapted to streams of data, and so on.

We discussed the computational cost of the sketching operation in Section 2.7, and the possible use of parallelization and distributed computing. In our settings, even when it is done linearly this operation is still faster than EM (Figure 6, left).

6 Large-scale proof of concept: speaker verification

Gaussians Mixture Models are popular for their capacity to smoothly approximate *any* distribution [87] by a large number of Gaussians. This is often the case with real data, and the problem of *fitting* a large GMM to data drawn from some distribution is somewhat different from that of *clustering* data and identifying reasonably well separated components, as presented in the previous section. In order to try out compressive methods on this challenging task, we test them on a speaker verification problem, with a classical approach requiring GMM referred to as Universal Background Model (GMM-UBM) [86].

6.1 Overview of Speaker Verification

Given a fragment of speech and a candidate speaker, the goal is to assess if the fragment was indeed spoken by that person.

We quickly describe GMM-UBM in this section. For more details we refer the reader to the original paper [86]. Similar to many speech processing tasks, this approach uses Mel Frequency Cepstrum Coefficients (MFCC) and their derivatives (Δ -MFCC) as features \mathbf{x}_i . Those features have been often modeled with GMMs or more advanced Markov models. However, in our framework we do *not* use Δ -MFCC; indeed those coefficients typically have a negligible range in dynamic compared to the MFCC, which results in a difficult and unstable choice of frequencies. As mentioned before, this problem may be potentially solved by a pre-whitening of the data, which we leave for future work. In this configuration, the speaker verification results will indeed not be state-of-the-art, but our goal is mainly to test our compressive approach on a different type of problem than that of clustering synthetic data, for which we have already observed excellent results.

In the GMM-UBM model, each speaker S is represented by one GMM (Θ_S, α_S) . The key point is the introduction of a model $(\Theta_{UBM}, \alpha_{UBM})$ that represents a “generic” speaker, referred to as Universal Background Model (UBM). Given speech data \mathcal{X} and a candidate speaker S , the statistic used for hypothesis testing is a likelihood ratio between the speaker and the generic model:

$$T(\mathcal{X}) = \frac{p_{\Theta_S, \alpha_S}(\mathcal{X})}{p_{\Theta_{UBM}, \alpha_{UBM}}(\mathcal{X})}. \quad (25)$$

If $T(\mathcal{X})$ exceeds a threshold τ , the data \mathcal{X} are considered as being uttered by the speaker S .

The GMMs corresponding to each speaker must somehow be “comparable” to each other and to the UBM. Therefore, the UBM is learned *prior* to individual speaker models, using a large database of speech data uttered by many speakers. Then, given training data \mathcal{X}_S specific to one speaker, one M-step from the EM algorithm *initialized with the UBM* is used to adapt the UBM and derive the model (Θ_S, α_S) . We refer the reader to [86] for more details on this procedure.

In our framework, the EM or compressive estimation algorithms are used to learn the UBM.

We note that this type of signal processing task may fully benefit from the advantages of the sketch structure described in Section 2.7. For instance, in practice one can imagine collecting bit by bit the data to train the UBM in a real-life environment, in which case the sketch and the UBM may be progressively updated without having to keep the spoken fragments, possibly of sensitive nature.

6.2 Setup

The experiments were performed on the classical NIST05 speaker verification database. Both training and testing fragments are 5-minutes conversations between two speakers. The database contains approximately 650 speakers, and 30 000 trials.

The MFCCs are computed using the Voicebox toolbox [21]. After filtering the audio data by a speech activity detector, the MFCCs are computed on 23ms frames with a 50% overlap. The first coefficient is removed and we obtain 12-dimensional features ($d = 12$).

Results are presented by choosing the threshold τ that yields the same rates of false alarm and missed detection, referred to as Equal Error Rate (EER). Each result is obtained as the mean of five experiments.

In all experiments, except when indicated otherwise, the compressive methods are performed using a sketch obtained by compressing the entire database of $n = 2.10^8$ MFCC vectors after voice activity detection. The compression is performed taking advantage of distributed computing, by dividing the database into 200 parts that are then compressed simultaneously on a computer cluster. Hence, even for a high number of frequencies $m = 10^5$ the compression of the $n = 2.10^8$ items takes less than an hour.

6.3 Results

	EER (%)		Time (s)	
	CL-OMPR	Algorithm 2	CL-OMPR	Algorithm 2
$m = 10^3$	40.3	32.5	7.10^2	5.10
$m = 10^4$	29.4	29.0	7.10^3	5.10^2
$m = 10^5$	28.8	28.6	7.10^4	5.10^3

Table 4: Comparison between CL-OMPR and Algorithm 2 for speaker verification, with $K = 64$.

Splitting algorithm In the previous section, Algorithm 2 was observed to be less accurate than CL-OMPR (Figure 5). However, as mentioned before the estimation problem considered here is somehow not to identify well-separated components, but rather to fit a GMM with a large number of components to a smooth probability density. In the first case, on synthetic data, Algorithm 2 is indeed expected to sometimes yield poor results: unlike a Matching Pursuit-based approach such as CL-OMPR, at each iteration it locally divides the current Gaussians rather than “exploring” elsewhere. In the second case however, Algorithm 2 may yield a correct approximation of the smooth density, by successively approaching it with GMMs at increasingly finer scales.

In Table 4, we compare the results obtained with CL-OMPR and Algorithm 2 on the speaker verification task using $K = 64$ Gaussians in the UBM. Results are indeed similar when the number of frequencies m is large, and even surprisingly better with Algorithm 2 for a low number of frequencies $m = 1000$. Naturally, Algorithm 2 is much faster than CL-OMPR, with more than a 10 times speedup.

Sketching a large database In Table 5, we compare EER results when using either $n_1 = 3.10^5$ items uniformly selected in the database to cover all speakers, or all $n_2 = 2.10^8$ items in the database. The compressive Algorithm 2 is performed at both scales, while EM is only performed with n_1 items, since the whole database is too large to be handled by the VLFeat toolbox on a machine with 8 GB of RAM. For the compressive approach, the use of the entire database indeed improves the results when compared

		$K = 8$		$K = 64$		$K = 512$	
		n_1	n_2	n_1	n_2	n_1	n_2
EM		31.4	n/a	29.5	n/a	27.5	n/a
Alg. 2	$m = 10^3$	32.5	31.2	31.1	32.5	31.2	29.4
	$m = 10^4$	32.1	30.7	30.2	29.0	30.3	29.1
	$m = 10^5$	32.5	30.7	29.8	28.6	29.4	29.2

Table 5: Comparison EM and Algorithm 2 for speaker verification, in terms of EER. For Algorithm 2, results that outperform these of EM are outlined.




Figure 7: Effect of the number of frequencies m on speaker verification results. Algorithm 2 is performed on the whole database with $n_2 = 2 \cdot 10^8$ items, while EM can only be performed on $n_1 = 3 \cdot 10^5$ items.

to using only n_1 items to compute the sketch. At low $K = 8$ or $K = 64$ and high number of frequencies m , the compressive approach using n_2 items outperforms EM using only n_1 items.

Limitations due to coherence. While increasing the number of components K seems to consistently improve the results of EM, it is not the case with the compressive method for a fixed sketch size m . A possible intuitive explanation could be that, by increasing the number of components we also increase the coherence between them – *i.e.* the Gaussians in the GMM are increasingly overlapping each other – which makes it more and more difficult to handle for any sparsity-based approach. In practice, it results in many components in the GMM having weights $\alpha \approx 0$. In other words, the algorithm outputs a K' -GMM with $K' < K$: there seems to be a “limit” number of components above which additional Gaussians are useless. It may be possible to deal with a higher level of sparsity by drastically increasing the number of frequencies m , at the cost of higher compression and estimation times.

Number of components K and compression. In Figure 7 we study the effect of m for various number of components $K = 8, 64$ and 512 . In each case we observe a sharp phase transition going from an EER of 50%, which corresponds to random guessing, to the results observed in Table 5. Somehow surprisingly, this phase transition does not seem to depend on K , unlike the one observed on synthetic data (Figure 4). As mentioned before it could be interesting to drastically increase m to see if the gap between results obtained EM and those obtained with Algorithm 2 can be bridged in the $K = 512$ case, however the phase transition pattern does not support this idea but rather a limitation of the method itself, maybe in the algorithmic approach.

Overall, results on synthetic and real data show that the *fitting* problem is, as expected, more challenging than the *clustering* problem for the proposed sparsity-based approach. Indeed, while the clustering problem (synthetic data) is that of identifying well-separated components of a sparse distribution, the fitting problem is similar to a *sparse approximation* task, which is known to be challenging when the “signal” (*i.e.* the true distribution of the data) is *not* sparse. Nevertheless, let us point out that in Figure 7, results approaching those of EM are obtained for $m = 3000$ frequencies only, which corresponds to a whopping 33000-fold compression of the database.

7 Information preservation guarantees ?

In this section we derive a number of information preservation guarantees of the proposed sketching operator. Let us come back to the “generic” compressive learning inverse problem introduced in Section 1:

$$\operatorname{argmin}_{P \in \Sigma} \|\hat{\mathbf{z}} - \mathcal{A}P\|_2, \quad (26)$$

where Σ is some “low-dimensional” model. Compressive estimation algorithms such as CL-OMP(R) (Algorithm 1) or the more scalable Algorithm 2 (specifically designed for GMM) seek an approximate solution to this problem in the case $\Sigma = \mathcal{G}_K$, with \mathcal{G}_K the set of K -sparse distributions in \mathcal{G} .

Precise recovery guarantees for CL-OMP(R) or Algorithm 2 are beyond the scope of this paper due to the random nature of several steps in these algorithms and the many non-convex optimization schemes that they contain. Instead, we rather demonstrated empirically in Sections 5 and 6 that these algorithms perform well on a large range of GMM estimation problems, with synthetic and real data.

In parallel, a fundamental question consists in asking if the problem is well-posed, *i.e.* if a potential solution of (26) is somehow guaranteed to be a “good” estimate of the distribution of the data. Namely, we ask the following questions:

- For a distribution $P_\Sigma \in \Sigma$, does the sketch $\mathbf{z} = \mathcal{A}P_\Sigma$ contain “enough” information to retrieve the distribution P_Σ ?
- Is this retrieval robust to using the empirical sketch $\hat{\mathbf{z}}$ instead of the true sketch \mathbf{z} ?
- Is this retrieval robust if $\mathbf{z} = \mathcal{A}P$ where the encoded distribution P is not *exactly* in the model Σ but only *well-approximated* by a distribution in the model (*i.e.* the distance $d(P, \Sigma) = \inf_{P_\Sigma \in \Sigma} d(P, P_\Sigma)$ is small for some metric d) ?

The answers to these questions are linked to the existence of a so-called *instance-optimal* decoder [37, 17], *i.e.* a (not necessarily tractable) reconstruction paradigm that is robust to noise and modeling error. In this paper, we show that with high probability the decoder induced by solving (26) is instance optimal with respect to the Maximum Mean Discrepancy (MMD) metric [99] (Eq. (22)), provided the model and frequency distribution satisfy some assumptions. We then prove this result for GMMs with bounded parameters, starting with the toy example of single Gaussians ($K = 1$).

The proof of our main Theorem (Theorem 3), given in Appendix B, introduces variants of usual tools in compressive sensing. The idea is to use the fact that a Lower Restricted Isometry Property (LRIP) induces the existence of a robust instance optimal decoder, as shown by Bourrier *et al.* [17]. To prove the LRIP for the measurement operator \mathcal{A} , we use the fact that the empirical mean $\|\mathcal{A}P - \mathcal{A}Q\|_2^2$ concentrates around its expectation $\gamma_\Lambda^2(P, Q)$, and use ϵ -coverings to extend this concentration result uniformly over the whole model – this method is similar to the “simple proof” of the RIP developed by Baraniuk *et al.* [5].

As we will see, the results are still in a preliminary state and suboptimal in some cases. However, the use of compressive sensing tools in the context of kernel mean embedding and Random Features is an original and promising lead for future work, as is the introduction of guarantees such as robustness to modeling error for a Generalized Method of Moments problem. They can be seen as our main theoretical contribution.

7.1 Existence of instance optimal decoders for sketched distributions

In this section we formulate a general result that guarantees robust decoding of any model Σ (not necessarily restricted to mixture models) under some hypotheses. The reader should refer to Appendix A for definitions related to ϵ -coverings.

Assumption $\mathbf{A}_1(\Sigma)$ (Compactness of the model). *The model Σ is compact with respect to the total variation norm $\|\cdot\|_{TV}$. In particular, it implies that it has finite covering numbers $N_{\Sigma, \|\cdot\|_{TV}}(\epsilon) < \infty$.*

In the context of mixture models, the following Lemma shows that compactness of any basic set of distributions \mathcal{G} extends to its set of K -sparse distributions \mathcal{G}_K , and that their covering numbers are related. Its proof is given in Appendix C.1.

Lemma 1. *Suppose the set of basic distributions \mathcal{G} is compact with respect to some norm $\|\cdot\|$, denote $C = \max_{P \in \mathcal{G}} \|P\|$. Then for all K the set of K -sparse distributions \mathcal{G}_K is also compact and satisfies, for all $\epsilon > 0$ and $0 < \tau < 1$,*

$$N_{\mathcal{G}_K, \|\cdot\|}(\epsilon) \leq \left(\frac{8C \cdot N_{\mathcal{G}, \|\cdot\|}(\tau\epsilon)}{(1-\tau)\epsilon} \right)^K. \quad (27)$$

Note that, in the case where $\|\cdot\| = \|\cdot\|_{TV}$ as in Assumption $\mathbf{A}_1(\Sigma)$, we have $C = 1$.

The second assumption involves the model Σ , the frequency distribution Λ , a small nonnegative constant $\eta \geq 0$ and a constant $A > 0$.

Assumption $\mathbf{A}_2(\eta, \Sigma, \Lambda, A)$ (Domination of the total variation norm). *For all $P_1, P_2 \in \Sigma$, we have*

$$\left[\gamma_\Lambda(P_1, P_2) \geq \eta \right] \Rightarrow \left[\|P_1 - P_2\|_{TV} \leq A\gamma_\Lambda(P_1, P_2) \right], \quad (28)$$

where γ_Λ is the MMD (22).

Note that, since $\|P_1 - P_2\|_{TV} \leq 2$ for all measures, if $\eta > 0$ Assumption $\mathbf{A}_2(\eta, \Sigma, \Lambda, A)$ is always verified with $A = 2/\eta$.

“Ideal” decoder. Given any sketch $\mathbf{z} \in \mathbb{C}^m$ and measurement operator \mathcal{A} , for all $P_1, P_2 \in \mathcal{P}$ we have

$$\left| \|\mathbf{z} - \mathcal{A}P_1\|_2 - \|\mathbf{z} - \mathcal{A}P_2\|_2 \right| \leq \|\mathcal{A}(P_1 - P_2)\|_2 \leq \|P_1 - P_2\|_{TV}, \quad (29)$$

by Lemma A.2 in Appendix A.2. Hence the function $P \in \Sigma \mapsto \|\mathbf{z} - \mathcal{A}P\|_2$ is continuous with respect to the total variation norm. If Assumption $\mathbf{A}_1(\Sigma)$ is satisfied and the model Σ is compact, this function reaches its minimum on it, which implies that the problem (26) has at least one solution.

In light of this observation, under Assumption $\mathbf{A}_1(\Sigma)$, we analyze below the information-theoretic estimation guarantees of the idealized decoder Δ that minimizes (26), *i.e.* return a distribution verifying:

$$\Delta(\mathbf{z}, \mathcal{A}) \in \underset{P \in \Sigma}{\operatorname{argmin}} \|\mathbf{z} - \mathcal{A}P\|_2. \quad (30)$$

We now turn to the main result of this section.

Theorem 3. *Consider a model Σ , a frequency distribution Λ , a small positive constant $1 \geq \eta > 0$ and a constant $A \geq 1$ such that Assumptions $\mathbf{A}_1(\Sigma)$ and $\mathbf{A}_2(\eta, \Sigma, \Lambda, A)$ hold.*

*Let $\mathbf{x}_i \in \mathbb{R}^d$, $i = 1 \dots n$ be n points drawn *i.i.d.* from an arbitrary distribution $P^* \in \mathcal{P}$, and $\omega_j \in \mathbb{R}^d$, $j = 1 \dots m$ be m frequencies drawn *i.i.d.* from Λ . Denote $\bar{P} = \Delta(\hat{\mathbf{z}}, \mathcal{A})$ the distribution reconstructed from the empirical sketch $\hat{\mathbf{z}}$.*

Let $\rho > 0$. Suppose that

$$m \geq 12A^2 \log \left(\frac{2}{\rho} \cdot N_{\Sigma, \|\cdot\|_{TV}} \left(\frac{\eta^2}{24} \right) \right). \quad (31)$$

Then, with probability at least $1 - \rho$ on the drawing of the items \mathbf{x}_i and sampling frequencies ω_j , we have

$$\gamma_\Lambda(P^*, \bar{P}) \leq 5 d_{TV}(P^*, \Sigma) + \frac{4(1+\sqrt{2\log(2/\rho)})}{\sqrt{n}} + \eta, \quad (32)$$

where γ_Λ is the MMD (22) and $d_{TV}(P^*, \Sigma) = \inf_{P \in \Sigma} \|P^* - P\|_{TV}$ is the distance from P^* to the model.

Hence the MMD between the distribution in the model recovered from the empirical sketch and the original distribution P^* is controlled by the distance (measured by the total variation norm) between P^* and the model, and a small additive error. This proves that the decoding is robust both to the use of the empirical sketch instead of the true sketch, and to the fact that P^* may not be exactly in the model. The choice of the kernel for the MMD is crucial to obtain meaningful guarantees. It somehow justifies the general strategy of choosing a kernel that maximizes the discriminative power of the MMD [97], which as mentioned earlier is the idea behind the proposed Adapted Radius heuristic. Further work will examine relationship between the MMD and other classic metrics such as likelihood or KL-divergence [85].

Limitations and future work. Theorem 3 does not hold for an additive error $\eta = 0$, which would be more akin to usual compressive sensing. In Appendix B.2, we give a version of the Theorem under a different set of Hypotheses \mathbf{H}_i that include the $\eta = 0$ case. However, exhibiting a model Σ and frequency distribution Λ that satisfy those hypotheses in the $\eta = 0$ case is yet to be done.

The control of the MMD $\gamma_\Lambda(P^*, \bar{P})$ with the distance $d_{TV}(P^*, \Sigma)$ is not optimal – ideally we would like to have the same metric on both sides of the inequality. In Appendix B, we formulate all results for a general metric d under some assumptions, then specialize in the case $d = \|\cdot\|_{TV}$ allowing for possible future generalizations.

We now apply Theorem 3 to GMMs, starting with the case $K = 1$ as a toy example.

7.2 Applications to single Gaussians (toy example)

We consider the case $K = 1$ where the model Σ is formed by single Gaussians as a toy example⁷, *i.e.* $\Sigma = \mathcal{G}$. We show that when the set \mathcal{T} of parameters of the Gaussians is compact, Assumption $\mathbf{A}_1(\Sigma)$ is verified with an explicit bound on the covering numbers of the model. Additionally, when the frequency distribution is Gaussian, Assumption $\mathbf{A}_2(\eta, \Sigma, \Lambda, A)$ is verified with a bound A that does not depend on the additive error η . Using the Gaussian kernel has the advantage of yielding closed-form expressions for the mean kernel, which at this point does not seem to be feasible with the proposed Adapted Radius distribution.

Recall the parametrization $\theta = [\mu, \sigma]$ with $\mu \in \mathbb{R}^d$ and $\sigma \in (\mathbb{R}_+ \setminus \{0\})^d$ of the set $\mathcal{G} = \{P_\theta\}_{\theta \in \mathcal{T}}$ of Gaussians with diagonal covariance. We suppose here that the set of parameters $\mathcal{T} \subset \mathbb{R}^{2d}$ is compact, *i.e.* closed and bounded. In particular, the variances of the considered Gaussians are bounded, and we denote $\sigma_{\min}^2 := \min_{[\mu, \sigma] \in \mathcal{T}} \min_\ell \sigma_\ell^2 > 0$ (resp. $\sigma_{\max}^2 := \max_{[\mu, \sigma] \in \mathcal{T}} \max_\ell \sigma_\ell^2 < \infty$) their minimum (resp. maximum) value. We also define $M := \max_{[\mu, \sigma] \in \mathcal{T}} \|\mu\|_2$, and finally we denote $\text{radius}(\mathcal{T}) := \min\{r > 0; \exists \mathbf{x} \in \mathbb{R}^{2d}, \mathcal{T} \subset B(\mathbf{x}, r)\}$ the *Chebyshev radius* [1] of \mathcal{T} , *i.e.* the minimal radius r such that \mathcal{T} is contained in a ball of radius r for the Euclidean norm.

Note that our framework is significantly different from many other works [94, 57, 66] which provide guarantees when the *support* of the distributions is compact, while here the Gaussian densities have *infinite support* but their *parameters* belong to a compact set, which is a far more realistic setting.

Theorem 4. *Suppose the set of parameters $\mathcal{T} \subset \mathbb{R}^{2d}$ is compact. Then the set of Gaussians $\mathcal{G} = \{P_\theta\}_{\theta \in \mathcal{T}}$ is compact. Furthermore, for all $\epsilon > 0$, we have*

$$N_{\mathcal{G}, \|\cdot\|_{TV}}(\epsilon) \leq \left(\frac{B}{\epsilon}\right)^{2d}, \quad (33)$$

where $B := 8 \max(\sigma_{\min}^{-1}, \sigma_{\min}^{-2}/\sqrt{2}) \text{radius}(\mathcal{T})$.

Thus Assumption $\mathbf{A}_1(\Sigma)$ is verified for the model $\Sigma = \mathcal{G}$. Assumption $\mathbf{A}_2(\eta, \Sigma, \Lambda, A)$ is also verified, and for a Gaussian frequency distribution we have the following:

Theorem 5. *Suppose the set of parameters $\mathcal{T} \subset \mathbb{R}^{2d}$ is compact, and the frequency distribution is an isotropic Gaussian $\Lambda = \mathcal{N}(0, \frac{a}{d}\mathbf{I})$ for some $a > 0$.*

Then, for all $P, Q \in \mathcal{G}$, we have

$$\|P - Q\|_{TV} \leq D\gamma_\Lambda(P, Q), \quad (34)$$

where

$$D = \max(\sigma_{\min}^{-1}, \sigma_{\min}^{-2}/\sqrt{2}) \sqrt{\frac{2dD_1 \cdot e^{3a\sigma_{\max}^2}}{a(1 - e^{-D_1})}} \quad \text{with } D_1 = \sigma_{\max}^2 a \left(1 + \frac{2M^2}{d}\right). \quad (35)$$

The proof is given in Appendix C.2.

As a consequence, Assumption $\mathbf{A}_2(\eta, \Sigma, \Lambda, A)$ is verified for $\Sigma = \mathcal{G}$, $\Lambda = \mathcal{N}(0, \sigma_\Lambda^2 \mathbf{I})$, any $\eta \geq 0$ and $A = \min(D, 2/\eta)$. Theorems 4 and 5 lead to the following immediate corollary of Theorem 3.

Corollary 1. *In the case $\Sigma = \mathcal{G}$ of single Gaussians with a compact set of parameters, for a Gaussian frequency distribution $\Lambda = \mathcal{N}(1, \frac{a}{d}\mathbf{I})$ and any constant $0 < \eta \leq 1$, Theorem 3 is verified with (31) replaced by*

$$m \geq 12A^2 \left(4d \log\left(\frac{C}{\eta}\right) + \log\left(\frac{2}{\rho}\right)\right), \quad (36)$$

where $A = \min(D, 2/\eta)$, $C = \sqrt{24B}$, B is defined as in Theorem 4 and D is defined as in Theorem 5.

⁷Obviously, fitting a single Gaussian to a dataset can easily be done by direct empirical estimators.

Hence, for a fixed model $\Sigma = \mathcal{G}$, the additive error η decreases exponentially with the number of measurements m . It is also interesting to note that conversely, for a small additive error $\eta \leq 2/D = \mathcal{O}(1/\sqrt{d})$, assuming that all parameters of the model appearing in the expression of D are constant, we have $A = \mathcal{O}(\sqrt{d})$ and thus the number of measurements m must grow as $\mathcal{O}(d^2)$ with the dimension. It is sub-optimal in the sense that the ideal estimators of the mean and diagonal covariance of a single Gaussian, *i.e.* the empirical mean and covariance, have size $\mathcal{O}(d)$. This number of measurements may scale with the size of the empirical estimators for Gaussian with *full* covariance, although the results presented in this paper do not directly apply to this case.

7.3 Application to GMMs

Theorem 4 and Lemma 1 allow for an immediate extension of Assumption $\mathbf{A}_1(\Sigma)$ from the set of basic distributions to the corresponding mixture models. In the case of GMMs, we have the following corollary, whose proof is given in Appendix C.2.

Corollary 2. *Suppose the set of parameters $\mathcal{T} \subset \mathbb{R}^{2d}$ is compact. Then, for all $K > 1$ the set of GMMs \mathcal{G}_K is compact. Furthermore, for all $\epsilon > 0$, we have*

$$N_{\mathcal{G}_K, \|\cdot\|_{TV}}(\epsilon) \leq \left(\frac{2(B+1)}{\epsilon} \right)^{(2d+1)K}, \quad (37)$$

where B is defined as in Theorem 4.

Unfortunately, unlike the compactness property, Assumption $\mathbf{A}_2(\eta, \Sigma, \Lambda, A)$ cannot be immediately extended from the set of basic distributions \mathcal{G} to the mixture model \mathcal{G}_K , and it is not clear whether doing so would require some additional hypotheses or not. Though we strongly believe that a result similar to the $K = 1$ case holds (see the discussion at the end of this section), here we use the fact that Assumption $\mathbf{A}_2(\eta, \Sigma, \Lambda, A)$ is verified with $A = 2/\eta$ regardless of the model and frequency distribution.

This leads to the following corollary.

Corollary 3. *In the case $\Sigma = \mathcal{G}_K$ of GMMs with a compact set of parameters, for any frequency distribution and constant $0 < \eta \leq 1$, Theorem 3 is verified with (31) replaced by*

$$m \geq 48\eta^{-2} \left(2K(2d+1) \log \left(\frac{C}{\eta} \right) + \log \frac{2}{\rho} \right), \quad (38)$$

where $C = \sqrt{48(B+1)}$, with B defined as in Theorem 4.

Conjecture. Corollary 3 suggests that the reconstruction error η for GMMs decreases as $\mathcal{O} \left(n^{-\frac{1}{2}} + m^{-\frac{1}{2}} \right)$ (up to some inverse exponential factor), which seems to nullify the advantages of the “compressive” approach. This is due to the use of the “worst” bound $A = 2/\eta$ in Assumption $\mathbf{A}_2(\eta, \Sigma, \Lambda, A)$. However, we strongly believe that Assumption $\mathbf{A}_2(\eta, \Sigma, \Lambda, A)$ may hold with a better bound that does not depend on η , similar to the $K = 1$ case.

We support this claim by empirically evaluating reconstruction results of the CL-OMPR algorithm with respect to the MMD in Figure 8. We observe a phase transition pattern similar to the one already noted in Section 5.5 for the KL-divergence, which is inconsistent with an additive error that scales in $\mathcal{O}(m^{-1/2})$. On the contrary, the $\mathcal{O}(n^{-1/2})$ decrease is indeed observed, which supports the theory but also the capacity of CL-OMPR to approximate the ideal decoder (30).

The proof of Assumption $\mathbf{A}_2(\eta, \Sigma, \Lambda, A)$ for general GMMs seems complex and beyond the scope of this paper. A possible strategy would be to be able to directly extend Assumption $\mathbf{A}_2(\eta, \Sigma, \Lambda, A)$ on the basic set of distributions \mathcal{G} to the corresponding mixture model \mathcal{G}_K , while another approach would be to use a different metric other than the total variation norm. The latter is particularly discussed in Appendix B.

8 Conclusion and outlooks

We presented a method for probability mixture estimation on a large database exploiting a *sketch* of the data instead of the data itself. The sketch is an appropriate structure that leads to considerable gain in




Figure 8: Reconstruction results of CL-OMPR for the γ_Λ metric with respect to m (left) and n (right), in dimension $d = 10$ and $K = 5$ components, using the true theoretical sketch $\mathcal{A}P_{\Theta_0, \alpha_0}$ on the left and $m = 5K(2d + 1)$ frequencies on the right. In a similar fashion to the KL-divergence (Section 5.2), the MMD γ_Λ is approximated by drawing $5 \cdot 10^5$ frequencies from Λ and by empirically evaluating (22), using the closed-form expression of the characteristic function of GMMs.

terms of memory. It can be computed in a distributed or streaming manner, and it can fully exploit the advantages of GPU computing.

A typical greedy method for sparse reconstruction was defined, leading to reconstruction algorithms both efficient and stable, even when the dictionary of atoms is infinite and uncountable. In the case of GMM, an additional efficient algorithm based on hierarchical splitting of GMMs was described.

A heuristic to select generalized moments based on a decomposition robust to high dimension and maximal variations of the characteristic function was designed. A procedure to estimate the parameter of this heuristic was described, resulting in a method that is faster than traditional kernel design by cross-validation, has the advantage of being unsupervised and thus is probably suited for other tasks, and yields better reconstruction results.

Excellent results were observed on synthetic data, where the greedy algorithms approach the reconstruction results of EM, using less memory and computation time when the number of database elements is large. The method was successfully applied to a large-scale speaker verification task. The hierarchical approach proved to be the most efficient method for this challenge, illustrating the diversity of the problem and of the proposed solutions. As in usual compressive sensing, limitations of the method when the number of sparse components in the distribution is large were observed.

Finally, information preservation guarantees were developed for the recovery of any compact set of distributions. The proof of Theorem 3 (Appendix B) introduced a weaker variant of the Restricted Isometry Property (RIP) for non-uniform recovery. We then applied this result to GMMs with bounded parameters, and observed a technical bottleneck between the toy $K = 1$ case and general GMMs.

Outlooks. As mentioned earlier, the method can readily be applied to other mixture models, such as mixtures of α -stable distributions which do not have explicit likelihood but whose characteristic function is known [91]. Based on the principle of maximizing the variation of the characteristic function, suitable heuristics for the choice of the sampling pattern may be derived for other models.

The method can also easily incorporate variants of Random Fourier Features that are faster to compute or more precise [69, 106]. Existing methods to learn the kernel [58, 108, 93, 78] may be adapted to our framework, and in return the proposed unsupervised kernel learning procedure and Adapted radius heuristic may be useful for other tasks.

As mentioned earlier, technical difficulties on the domination between certain metrics were observed between the toy $K = 1$ case and general GMMs, pointing to a promising lead for future work. The proof of Theorem 3 also uses innovative variants of several classical tools in compressive sensing, which may be useful in the study of other instances of generalized compressive sensing.

Acknowledgment

This work was supported in part by the European Research Council, PLEASE project (ERC-StG- 2011-277906).

References

References

- [1] Chebyshev Radius. https://www.encyclopediaofmath.org/index.php/Chebyshev_radius, 2016.
- [2] Dimitris Achlioptas. Database-friendly random projections. In *Principles of database systems (PODS)*, pages 274–281, 2001.
- [3] Peter Ahrendt. The Multivariate Gaussian Probability Distribution. Technical report, IMM, Technical University of Denmark, 2005.
- [4] Nachman Aronzajn. Theory of reproducing kernels. *Transactions of the American Mathematical Society*, 68:337–404, 1950.
- [5] Richard Baraniuk, Mark Davenport, Ronald DeVore, and Michael Wakin. A simple proof of the restricted isometry property for random matrices. *Constructive Approximation*, 28(3):253–263, 2008.
- [6] Richard G. Baraniuk. Compressive sensing. *IEEE Signal Processing Magazine*, 24(2):118–121, 2007.
- [7] Eran Bashan, Raviv Raich, and Alfred O. Hero. Optimal two-stage search for sparse targets using convex criteria. *IEEE Transactions on Signal Processing*, 56(11):5389–5402, 2008.
- [8] Stephen Becker. L-BFGSB-C. <https://github.com/stephenbecker/L-BFGS-B-C>, 2013.
- [9] Mikhail Belkin and Kaushik Sinha. Polynomial learning of distribution families. In *IEEE 51st Annual Symposium on Foundations of Computer Science*. Ieee, 2010.
- [10] Mikhail Belkin and Kaushik Sinha. Toward Learning Gaussian Mixtures with Arbitrary Separation. In *Conference On Learning Theory (COLT)*, 2010.
- [11] Karine Bertin, Erwan Le Pennec, and Vincent Rivoirard. Adaptive Dantzig density estimation. *Annales de l'Institut Henri Poincaré, Probabilités et Statistiques*, 47(1):43–74, 2011.
- [12] Thomas Blumensath. Sampling and reconstructing signals from a union of linear subspaces. *IEEE Transactions on Information Theory*, 57(7):4660–4671, 2011.
- [13] Thomas Blumensath and Mike E. Davies. Iterative thresholding for sparse approximations. *Journal of Fourier Analysis and Applications*, 14(5-6):629–654, 2008.
- [14] Thomas Blumensath and Mike E. Davies. Iterative hard thresholding for compressed sensing. *Applied and Computational Harmonic Analysis*, 27(3):265–274, 2009.
- [15] Liefeng Bo and Cristian Sminchisescu. Efficient Match Kernels between Sets of Features for Visual Recognition. In *Advances in Neural Information Processing System (NIPS)*, 2009.
- [16] Karsten M. Borgwardt, Arthur Gretton, Malte J. Rasch, Hans Peter Kriegel, Bernhard Schölkopf, and Alex J. Smola. Integrating structured biological data by Kernel Maximum Mean Discrepancy. *Bioinformatics*, 22(14):49–57, 2006.
- [17] A Bourrier. Compressed sensing and dimensionality reduction for unsupervised learning. 2014.
- [18] Anthony Bourrier, Mike E. Davies, Tomer Peleg, and Rémi Gribonval. Fundamental performance limits for ideal decoders in high-dimensional linear inverse problems. *IEEE Transactions on Information Theory*, 60(12):7928–7946, 2014.
- [19] Anthony Bourrier, Rémi Gribonval, and Patrick Pérez. Compressive gaussian mixture estimation. In *IEEE International Conference on Acoustic, Speech and Signal Processing (ICASSP)*, pages 6024–6028, 2013.
- [20] Anthony Bourrier, Rémi Gribonval, and Patrick Pérez. Compressive Gaussian Mixture estimation. In Haulger Boche, Robert Calderbank, Gitta Kutyniok, and Jan Vybiral, editors, *Compressed Sensing and its Applications - MATHEON Workshop 2013*, pages 6024–6028. Birkhäuser Basel, 2015.
- [21] M. Brooks. VOICEBOX: Speech Processing Toolbox for MATLAB, 2005.
- [22] Florentina Bunea, Alexandre B. Tsybakov, Marten H. Wegkamp, and Adrian Barbu. SPADES and mixture models. *The Annals of Statistics*, 38(4):2525–2558, 2010.
- [23] Richard H Byrd, Peihuang Lu, Jorge Nocedal, and Ciyong Zhu. A Limited Memory Algorithm for Bound Constrained Optimization. *SIAM Journal on Scientific Computing*, 16(5):1190–1208, 1995.
- [24] Robert Calderbank, Sina Jafarpour, and Robert Schapire. Compressed learning: Universal sparse dimensionality reduction and learning in the measurement domain. Technical report, 2009.

- [25] Emmanuel Candes, Yonina C. Eldar, Deanna Needell, and Paige Randall. Compressed sensing with coherent and redundant dictionaries. *Applied and Computational Harmonic Analysis*, 31(1):59–73, 2011.
- [26] Emmanuel Candes and Yaniv Plan. Tight oracle inequalities for low-rank matrix recovery from a minimal number of noisy random measurements. *IEEE Transactions on Information Theory*, 57(4):2342–2359, 2011.
- [27] Emmanuel Candes, Justin K. Romberg, and Terence Tao. Stable signal recovery from incomplete and inaccurate measurements. *Communications on Pure and Applied Mathematics*, 59(8):1207–1223, 2006.
- [28] Emmanuel Candes and Terence Tao. Decoding by linear programming. *IEEE Transactions on Information Theory*, 51(12):4203–4215, 2005.
- [29] Emmanuel Candes and Terence Tao. Near-optimal signal recovery from random projections: Universal encoding strategies? *IEEE Transactions on Information Theory*, 52(12):5406–5425, 2006.
- [30] Olivier Cappé and Eric Moulines. Online EM Algorithm for Latent Data Models. *Journal of the Royal Statistical Society*, 71(3):593–613, 2009.
- [31] Marine Carrasco and Jean-Pierre Florens. Generalization of GMM to a continuum of moment conditions. *Econometric Theory*, 2000.
- [32] Marine Carrasco and Jean-Pierre Florens. Efficient GMM estimation using the empirical characteristic function. *IDEI Working Paper*, 2002.
- [33] Marine Carrasco and Jean-Pierre Florens. On The Asymptotic Efficiency Of Gmm. *Econometric Theory*, 30(02):372–406, 2014.
- [34] Yohann De Castro and Fabrice Gamboa. Exact reconstruction using Beurling minimal extrapolation. *Journal of Mathematical Analysis and Applications*, 395:336–354, 2012.
- [35] Krzysztof Choromanski and Sindhwani Google Com. Recycling Randomness with Structure for Sublinear time Kernel Expansions. *arXiv:1605.09049v1*, 2016.
- [36] Kacper Chwialkowski, Aaditya Ramdas, Dino Sejdinovic, and Arthur Gretton. Fast Two-Sample Testing with Analytic Representations of Probability Measures. In *Advances in Neural Information Processing Systems (NIPS)*, 2015.
- [37] Albert Cohen, Wolfgang Dahmen, and Ronald DeVore. Compressed sensing and best k-term approximation. *Journal of the American mathematical Society*, 22(1):211–231, 2009.
- [38] David A. Cohn, Zoubin Ghahramani, and Michael I. Jordan. Active Learning with Statistical Models. *Journal of Artificial Intelligence Research*, 4:129–145, 1996.
- [39] Graham Cormode and Marios Hadjieleftheriou. Methods for finding frequent items in data streams. *The VLDB Journal*, 19(1):3–20, 2009.
- [40] Graham Cormode, Flip Korn, S. Muthukrishnan, and Srivastava Divesh. Diamond in the rough: Finding hierarchical heavy hitters in multi-dimensional data. In *International Conference on Management of Data*, pages 155–166, 2004.
- [41] Graham Cormode and S. Muthukrishnan. An improved data stream summary: the count-min sketch and its applications. *Journal of Algorithms*, 55(1):58–75, 2005.
- [42] Thomas M. Cover and Joy A. Thomas. *Elements of information theory*. Wiley-Interscience, 2006.
- [43] Felipe Cucker and Steve Smale. On the mathematical foundations of learning. *Bulletin of the American Mathematical Society*, 39(1):1–49, 2002.
- [44] Sanjoy Dasgupta. Learning mixtures of Gaussians. In *IEEE 51st Annual Symposium on Foundations of Computer Science*, number May, 1999.
- [45] Thong T. Do, Lu Gan, Nam H. Nguyen, and Trac D. Tran. Fast and Efficient Compressive Sensing using Structurally Random Matrices. *IEEE Transactions on Signal Processing*, 60(1):139–154, 2012.
- [46] David L. Donoho. Compressed sensing. *IEEE Transactions on Information Theory*, 52(4):1289–1306, 2006.
- [47] John Duchi. Derivations for Linear Algebra and Optimization. Technical report, 2007.
- [48] Alexei A Fedotov, Peter Harremoës, and Flemming Topsøe. Refinements of Pinsker’s Inequality. *IEEE Transactions on Information Theory*, 49(6):1491–1498, 2003.
- [49] A Feuerverger and RA Mureika. The empirical characteristic function and its applications. *The annals of Statistics*, 1977.
- [50] Andrey Feuerverger and Philip McDunnough. On Some Fourier methods for Inference. *Journal of the American Statistical Association*, 76(374):379–387, 1981.
- [51] Tom Fischer. Existence, uniqueness, and minimality of the Jordan measure decomposition. *arXiv:1206.5449*, 2012.
- [52] Simon Foucart and Holger Rauhut. *A Mathematical Introduction to Compressive Sensing*. Applied and Numerical Harmonic Analysis. Springer New York, NY, 2013.

- [53] Dmitriy Fradkin and David Madigan. Experiments with random projections for machine learning. In *International Conference on Knowledge discovery and data mining (KDD)*, pages 517–522, 2003.
- [54] K. Fukumizu, Arthur Gretton, X. Sun, and B. Schölkopf. Kernel Measures of Conditional Dependence. In *Advances in Neural Information Processing System (NIPS)*, 2007.
- [55] Anna C. Gilbert, Yannis Kotidis, S. Muthukrishnan, and Martin J. Strauss. How to summarize the universe: Dynamic maintenance of quantiles. In *International Conference on Very Large Data Bases (VLDB)*, pages 454–465, 2002.
- [56] Raja Giryes, Guillermo Sapiro, and Alex M Bronstein. Deep Neural Networks with Random Gaussian Weights : A Universal Classification Strategy ? *arXiv:1504.08291*, 2015.
- [57] Arthur Gretton, Karsten M. Borgwardt, Malte Rasch, Bernhard Scholkopf, and Alexander J. Smola. A Kernel Method for the Two-Sample Problem. In *Advances in Neural Information Processing Systems (NIPS)*, 2006.
- [58] Arthur Gretton, Bharath Sriperumbudur, Dino Sejdinovic, Heiko Strathmann, and Massimiliano Pontil. Optimal kernel choice for large-scale two-sample tests. *Advances in Neural Information Processing Systems (NIPS)*, pages 1214–1222, 2012.
- [59] Rémi Gribonval, Emmanuel Bacry, Stéphane Mallat, Philippe Depalle, and Xavier Rodet. Analysis of sound signals with high resolution matching pursuit. In *IEEE-SP International Symposium on Time-Frequency and Time-Scale Analysis*, pages 125–128, 1996.
- [60] Alastair R. Hall. *Generalized method of moments*. 2005.
- [61] Jarvis Haupt, Rui Castro, Robert Nowak, and S T May. Distilled Sensing : Adaptive Sampling for Sparse Detection and Estimation. *IEEE Transactions on Information Theory*, 57(9):6222–6235, 2011.
- [62] T Hofmann, B Schölkopf, and a J Smola. Kernel methods in machine learning. *The Annals of Statistics*, 36(3):1171–1220, 2008.
- [63] Daniel Hsu and Sham M. Kakade. Learning mixtures of spherical gaussians: moment methods and spectral decompositions. In *Conference on Innovations in Theoretical Computer Science*, 2013.
- [64] Prateek Jain, Ambuj Tewari, and Inderjit S. Dhillon. Orthogonal matching pursuit with replacement. In *Advances in Neural Information Processing Systems (NIPS)*, pages 1215–1223, 2011.
- [65] Wittawat Jitkrittum, Zoltán Szabó, Kacper Chwialkowski, and Arthur Gretton. Interpretable Distribution Features with Maximum Testing Power. *Advances in Neural Information and Processing Systems (NIPS)*, (Nips):1–21, 2016.
- [66] Motonobu Kanagawa and Kenji Fukumizu. Recovering Distributions from Gaussian RKHS Embeddings. In *17th International Conference on Artificial Intelligence and Statistics*, pages 457–465, 2014.
- [67] Nicolas Keriven. SketchMLbox : a Matlab toolbox for large-scale learning of mixture models. <http://sketchml.gforge.inria.fr>, 2016.
- [68] Charles L. Lawson and Richard J. Hanson. Solving least squares problems. *SIAM classics in applied mathematics*, 1995.
- [69] Quoc Le, Tamás Sarlós, and Alex J Smola. Fastfood — Approximating Kernel Expansions in Loglinear Time. *International Conference on Machine Learning (ICML)*, 28(1), 2013.
- [70] Luc Le Magoarou and Rémi Gribonval. Flexible Multi-layer Sparse Approximations of Matrices and Applications. *IEEE Journal of Selected Topics in Signal Processing*, 10(4):688–700, 2016.
- [71] Boris Mailhé and Rémi Gribonval. LocOMP: algorithme localement orthogonal pour l’approximation parcimonieuse rapide de signaux longs sur des dictionnaires locaux. In *Actes du XXIIe colloque GRETSI*, 2009.
- [72] Oldaric-Ambrym Maillard and Rémi Munos. Compressed Least-Squares Regression. In *Advances in Neural Information and Processing Systems (NIPS)*, 2009.
- [73] Stéphane Mallat and Zhifeng Zhang. Matching Pursuit With Time-Frequency Dictionaries. *IEEE Transactions in Signal Processing*, 41(12):3397–3415, 1993.
- [74] Krikamol Muandet, Kenji Fukumizu, Francesco Dinuzzo, and Bernhard Schölkopf. Learning from distributions via support measure machines. *Advances in Neural Information Processing Systems (NIPS)*, 2012.
- [75] Sangnam Nam, Mike E. Davies, Michael Elad, and Rémi Gribonval. The cosparsity analysis model and algorithms. *Applied and Computational Harmonic Analysis*, 34(1):30–56, 2013.
- [76] Junier B. Oliva, Avinava Dubey, Barnabas Poczos, Jeff Schneider, and Eric P. Xing. Bayesian Nonparametric Kernel-Learning. In *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 2015.
- [77] Junier B. Oliva, Dougal J. Sutherland, and Jeff Schneider. Deep Mean Maps. *arXiv:1511.04150*, 2015.

- [78] Brooks Paige, Dino Sejdinovic, and Frank Wood. Super-Sampling with a Reservoir. *Uncertainty in Artificial Intelligence*, 2016.
- [79] Y.C. Pati, R. Rezaifar, and P.S. Krishnaprasad. Orthogonal matching pursuit: recursive function approximation with applications to wavelet decomposition. In *Asilomar Conference on Signals, Systems and Computers*, 1993.
- [80] Mert Pilanci, Laurent El Ghaoui, and Venkat Chandrasekaran. Recovery of sparse probability measures via convex programming. In *Advances in Neural Information Processing Systems (NIPS)*, pages 2429–2437, 2012.
- [81] Gilles Puy, Mike E. Davies, and Rémi Gribonval. Linear embeddings of low-dimensional subsets of a Hilbert space to \mathbb{R}^m . In *European Signal Processing Conference (EUSIPCO)*, pages 469–473, 2015.
- [82] Ali Rahimi and Benjamin Recht. Random Features for Large Scale Kernel Machines. *Advances in Neural Information Processing Systems (NIPS)*, 2007.
- [83] Ali Rahimi and Benjamin Recht. Weighted sums of random kitchen sinks: Replacing minimization with randomization in learning. *Advances in Neural Information Processing Systems (NIPS)*, 2009.
- [84] Hugo Reberedo, Francesco Renna, Robert Calderbank, and Miguel D. Rodrigues. Compressive Classification. In *IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, pages 1029–1032, 2013.
- [85] Sashank J. Reddi, Aaditya Ramdas, Barnabás Póczos, Aarti Singh, and Larry Wasserman. On the Decreasing Power of Kernel and Distance based Nonparametric Hypothesis Tests in High Dimensions. In *AAAI Conference on Artificial Intelligence*, 2015.
- [86] Douglas A. Reynolds, Thomas F. Quatieri, and Robert B. Dunn. Speaker Verification Using Adapted Gaussian Mixture Models. *Digital Signal Processing*, 10(1-3):19–41, 2000.
- [87] Douglas A. Reynolds and Richard C. Rose. Robust text-independent speaker identification using Gaussian mixture speaker models. *IEEE Transactions on Speech and Audio Processing*, 3(1):72–83, 1995.
- [88] Walter Rudin. *Fourier Analysis on Groups*. Interscience Publishers, 1962.
- [89] Walter Rudin. *Real and complex analysis*. McGraw-hill edition, 1987.
- [90] S.O. Sadjadi, M. Slaney, and L. Heck. MSR Identity Toolbox v1.0: A MATLAB toolbox for speaker-recognition Research. *Speech and Language Processing Technical Committee Newsletter*, 2013.
- [91] Diego Salas-Gonzalez, Ercan E. Kuruoglu, and Diego P. Ruiz. Finite mixture of α -stable distributions. *Digital Signal Processing*, 19(2):250–264, 2009.
- [92] Bernhard Schölkopf. Computing Functions of Random Variables via Reproducing Kernel Hilbert Space Representations. *Statistics and Computing*, 25(4):755–766, 2015.
- [93] Aman Sinha and John Duchi. Learning Kernels with Random Features. In *Advances in Neural Information and Processing Systems (NIPS)*, 2016.
- [94] Alexander J. Smola, Arthur Gretton, Le Song, and Bernhard Schölkopf. A Hilbert Space Embedding for Distributions. In *International Conference on Algorithmic Learning Theory*, pages 13–31, 2007.
- [95] Karthik Sridharan. A Gentle Introduction to Concentration Inequalities. Technical report, 2002.
- [96] Bharath K. Sriperumbudur. Mixture density estimation via hilbert space embedding of measures. In *IEEE International Symposium on Information Theory*, pages 1027–1030, 2011.
- [97] Bharath K. Sriperumbudur, Kenji Fukumizu, Arthur Gretton, Gert R. G. Lanckriet, and Bernhard Schölkopf. Kernel choice and classifiability for RKHS embeddings of probability distributions. In *Advances in Neural Information and Processing Systems (NIPS)*, 2009.
- [98] Bharath K. Sriperumbudur, Arthur Gretton, Kenji Fukumizu, Bernhard Schölkopf, and Gert R. G. Lanckriet. Hilbert space embeddings and metrics on probability measures. *The Journal of Machine Learning Research*, 11:1517–1561, 2010.
- [99] Dougal J Sutherland, Junier B. Oliva, P Barnabas, and Jeff Schneider. Linear-time Learning on Distributions with Approximate Kernel Embeddings. *arXiv:1509.07553*, 2015.
- [100] Nitin Thaper, Sudipto Guha, Piotr Indyk, and Nick Koudas. Dynamic multidimensional histograms. In *International conference on Management of data (SIGMOD)*, pages 428–439, 2002.
- [101] Kien C. Tran. Estimating mixtures of normal distributions via empirical characteristic function. *Economic Reviews*, 17(2):167–183, 1998.
- [102] Andrea Vedaldi and Brian Fulkerson. VLFeat - An open and portable library of computer vision algorithms. Technical report, 2010.
- [103] Andrea Vedaldi and Andrew Zisserman. Efficient additive kernels via explicit feature maps. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 34(3):480–492, 2012.

- [104] Vasily Volkov and James Demmel. Benchmarking g GPUs to Tune Dense Linear Algebra. In *High Performance Computing, Networking, Storage and Analysis*, 2008.
- [105] A Wilson and R Adams. Gaussian process kernels for pattern discovery and extrapolation. In *International Conference on Machine Learning (ICML)*, volume 28, pages 1067–1075, 2013.
- [106] Felix Xinnan, Yu Ananda, Theertha Suresh, Krzysztof Choromanski, Daniel Holtmann-Rice, and Sanjiv Kumar. Orthogonal Random Features. In *Advances in Neural Information and Processing Systems (NIPS)*, 2016.
- [107] Dinghai Xu and John Knight. Continuous Empirical Characteristic Function Estimation of Mixtures of Normal Parameters. *Econometric Reviews*, 30(1):25–50, 2010.
- [108] Zichao Yang, Alexander J Smola, Le Song, and Andrew Gordon Wilson. A la Carte — Learning Fast Kernels. *International Conference on Artificial Intelligence and Statistics (AISTATS)*, 38:1098–1106, 2015.
- [109] Peng Zhang and Yuxiang Gao. Matrix Multiplication on High-Density Multi-GPU Architectures: Theoretical and Experimental Investigations Peng. In *ISC High Performance*, volume 1, pages 17–30. 2015.
- [110] Ji Zhao and Deyu Meng. FastMMD: Ensemble of Circular Discrepancy for Efficient Two-Sample Test. *Journal Neural Computation*, 27(6):1345–1372, 2015.
- [111] Xi Zhou, Xiaodan Zhuang, Shuicheng Yan, Shih-Fu Chnag, Mark Hasegawa-Johnson, and Thomas S. Huang. SIFT-Bag Kernel for Video Event Analysis. In *ACM International Conference on Multimedia*, pages 229–238, 2008.

A Definitions, preliminary results

In this section, we group some definitions and useful results.

A.1 Positive definite kernels

We recall the definition of p.d. kernels.

Definition A.1 (Positive definite kernel). *Let X be an arbitrary set. A symmetric function (or **kernel**) $\kappa : X \times X \rightarrow \mathbb{C}$ is called **positive definite (p.d.)** if, for all $n \in \mathbb{N}$, $c_1, \dots, c_n \in \mathbb{C}$ and all $\mathbf{x}_1, \dots, \mathbf{x}_n \in X$, we have*

$$\sum_{i,j=1}^n c_i \bar{c}_j \kappa(\mathbf{x}_i, \mathbf{x}_j) \geq 0.$$

Note that strict positivity is not mandatory in the above equation. In terms of vocabulary, p.d. kernels bear connections with, *e.g.*, positive *semi*-definite matrices (however they are indeed called positive definite kernels in the literature).

Definition A.2 (Positive definite function). *A function $\mathbf{K} : \mathbb{R}^d \rightarrow \mathbb{C}$ is called **positive definite** if the kernel defined by $\kappa(\mathbf{x}, \mathbf{y}) = \mathbf{K}(\mathbf{x} - \mathbf{y})$ is positive definite.*

A.2 Measures

Definition A.3 (Nonnegative measure). *A measure $\nu \in E$ over a measurable space (X, \mathcal{B}) is said **nonnegative** if:*

$$\forall B \in \mathcal{B}, \nu(B) \geq 0.$$

Definition A.4 (Support of a measure). *The **support** of a signed measure $\nu \in E$ over a measurable, topological space X is defined to be the closed set,*

$$\text{supp}(\nu) := X \setminus \bigcup \{U \subset X : U \text{ is open, } \nu(U) = 0\}.$$

Definition A.5 (Total variation norm, Finite measure). *Let $\nu \in E$ be a signed measure over a measurable space (X, \mathcal{B}) . Define the Jordan decomposition (ν^-, ν^+) of ν where ν^+ and ν^- are positive measures (see [51] and [89] Chap. 6 for more details). Denote $|\nu| = \nu^+ + \nu^-$. The **total variation norm** of ν is defined as:*

$$\|\nu\|_{TV} = |\nu|(X) = \int_X d|\nu|(\mathbf{x}).$$

The measure ν is said **finite** if $\|\nu\|_{TV} < \infty$.

Note that if ν is totally continuous with respect to the Lebesgue measure, *i.e.* if there exists an integrable function f such that $d\nu(\mathbf{x}) = f(\mathbf{x})d\mathbf{x}$, then the total variation norm is the classic L^1 -norm of this function: $\|\nu\|_{TV} = \|f\|_{L^1}$.

We have the following bounds.

Lemma A.2. *For any sketching operator \mathcal{A} obtained by sampling the characteristic function, and any finite signed measure $\nu \in E$, we have*

$$\|\mathcal{A}\nu\|_2 \leq \|\nu\|_{TV}. \quad (\text{A.39})$$

For any frequency sampling distribution Λ and any pair of probability distributions $P, Q \in \mathcal{P}$, we have

$$\gamma_\Lambda(P, Q) \leq \|P - Q\|_{TV}. \quad (\text{A.40})$$

Proof. For any $\nu \in E$, we have

$$\|\mathcal{A}\nu\|_2^2 = \frac{1}{m} \sum_{j=1}^m \left| \int_{\mathbb{R}^d} e^{-i\omega_j^T \mathbf{x}} d\nu(\mathbf{x}) \right|^2 \leq \frac{1}{m} \sum_{j=1}^m \left(\int_{\mathbb{R}^d} d|\nu|(\mathbf{x}) \right)^2 = \|\nu\|_{TV}^2.$$

For all $P, Q \in \mathcal{P}$ a simple reformulation of $\gamma_\Lambda(P, Q)$, see [98], is:

$$\gamma_\Lambda(P, Q)^2 = \iint \kappa(\mathbf{x}, \mathbf{y}) d(P - Q)(\mathbf{x}) d(P - Q)(\mathbf{y}),$$

where the kernel κ is defined by (19) and (20). Since $|\kappa| \leq 1$, we immediately obtain the result. \square

A.3 Covering numbers

Definition A.6 (Ball, ϵ -covering, Covering number). *Let (X, d) be a metric space. For any $\epsilon > 0$ and $x \in X$, we denote $B_X(x, \epsilon)$ the **ball** of radius ϵ centered at the point x :*

$$B_X(x, \epsilon) = \{y \in X, d(x, y) \leq \epsilon\}.$$

*Let $Y \subseteq X$ be a subset of X . A subset $Z \subseteq Y$ is an ϵ -**covering** of Y if $Y \subseteq \bigcup_{z \in Z} B_X(z, \epsilon)$.*

*The **covering number** $N_{Y,d}(\epsilon) \in \mathbb{N} \cup \{+\infty\}$ is the smallest number of points $y_i \in Y$ such that the set $\{y_i\}$ is an ϵ -covering of Y .*

Remark A.1. *A subset Y of a topological space (X, d) that has finite covering numbers for any $\epsilon > 0$ is called totally bounded and is not necessarily compact: a set is in fact compact if and only if it is totally bounded and complete. Hence, though in the rest of the paper we often focus on explicitly bounding the covering numbers of certain sets, if compactness of these sets is required it will have to be proved independently.*

Our definition of covering numbers is that of *internal* covering numbers, meaning that the centers of the covering balls are required to be included in the set being covered. Somewhat counter-intuitively these covering numbers (for a fixed radius ϵ) are not necessarily increasing with the inclusion of sets⁸. We have instead the following property:

Lemma A.3. *Let $A \subseteq B \subseteq X$ be subsets of a metric space (X, d) , and $\epsilon > 0$. Then,*

$$N_{A,d}(\epsilon) \leq N_{B,d}(\epsilon/2). \quad (\text{A.41})$$

Proof. Let b_1, \dots, b_N be a $\epsilon/2$ -covering of B . We construct a ϵ -covering a_i of A in the following way. Each b_i is either: a) in the set A , in which case we take $a_i = b_i$, b) at distance less than $\epsilon/2$ of a point $a \in A$, in which case we take $a_i = a$ and note that the ball centered on a_i covers at least as much as the ball centered in b_i , i.e. $B_X(b_i, \epsilon/2) \subset B_X(a_i, \epsilon)$, c) in none of these cases and we discard it. There are less a_i 's than b_i 's, and the union of balls of radius ϵ with centers a_i covers at least as much as the balls of radius $\epsilon/2$ with centers b_i , and therefore the set of a_i 's is a ϵ -covering of B and of A . \square

Another useful property is related to the embedding of sets by a Lipschitz function.

⁸For instance, consider a set A formed by two points, included in set B which is a ball of radius ϵ . Suppose those two points diametrically opposed in B . We have $A \subset B$, but two balls of radius ϵ are required to cover A (since their centers have to be in A), while only one such ball is sufficient to cover B .

Lemma A.4. Let (X, d) and (X', d') be two metric spaces, and $Y \subseteq X$, $Y' \subseteq X'$. If there exists a surjective function $f : Y \rightarrow Y'$ which is L -Lipschitz with $L > 0$, i.e. such that

$$\forall x, y \in Y, d'(f(x), f(y)) \leq Ld(x, y),$$

then for all $\epsilon > 0$ we have

$$N_{Y', d'}(\epsilon) \leq N_{Y, d}(\epsilon/L). \quad (\text{A.42})$$

Proof. Define $\epsilon_2 = \epsilon/L$, denote $N = N_{Y, d}(\epsilon_2)$, and let $y_i \in Y$, $i = 1, \dots, N$ be an ϵ_2 -covering of Y . Let any $y' \in Y'$. There exists $y \in Y$ such that $f(y) = y'$ since f is surjective. Let y_i be a center of a ball in the ϵ_2 -covering of Y , we have

$$d'(y', f(y_i)) = d'(f(y), f(y_i)) \leq Ld(y, y_i) \leq L\epsilon_2 = \epsilon.$$

Thus $\{f(y_i)\}_{i=1, \dots, N}$ is an ϵ -covering of Y' , and we have $N_{Y', d'}(\epsilon) \leq N$. \square

Finally, we report a property from [43]:

Lemma A.5 ([43], Prop. 5). Let $(X, \|\cdot\|)$ be a Banach space of finite dimension d . Then for any $\epsilon > 0$, $x \in X$ and $R > 0$ we have

$$N_{B_X(x, R), \|\cdot\|}(\epsilon) \leq \left(\frac{4R}{\epsilon}\right)^d. \quad (\text{A.43})$$

A.4 Concentration of averages

We will use Bernstein's inequality in the following simple version [95]:

Lemma A.6 (Bernstein's inequality ([95], Thm. 6)). Let $x_i \in \mathbb{R}$, $i = 1, \dots, n$ be i.i.d. bounded random variables such that $\mathbb{E}x_i = 0$, $|x_i| \leq M$ and $\text{Var}(x_i) \leq \sigma^2$ for all i 's.

Then for all $t > 0$ we have

$$P\left(\frac{1}{n} \sum_{i=1}^n x_i \geq t\right) \leq \exp\left(-\frac{nt^2}{2\sigma^2 + 2Mt/3}\right). \quad (\text{A.44})$$

We also report a concentration result in Hilbert spaces from [83].

Lemma A.7 ([83], Lemma 4). Let $\mathbf{x}_i \in \mathcal{H}$, $i = 1, \dots, n$ be i.i.d. random variables in a Hilbert Space $(\mathcal{H}, \|\cdot\|)$ such that $\|\mathbf{x}_i\| \leq M$ with probability one. Denote $\bar{\mathbf{x}}$ their empirical average $\bar{\mathbf{x}} = (\sum_{i=1}^n \mathbf{x}_i)/n$. Then for any $\rho > 0$, with probability at least $1 - \rho$,

$$\|\bar{\mathbf{x}} - \mathbb{E}\bar{\mathbf{x}}\| \leq \frac{M}{\sqrt{n}} \left(1 + \sqrt{2 \log \frac{1}{\rho}}\right). \quad (\text{A.45})$$

B Proof of Theorem 3

B.1 Lower RIP

A measurement operator \mathcal{A} satisfies the usual generalized Lower Restricted Isometry Property (LRIP) [17] on the model Σ with constant $\beta > 0$ if:

$$\forall P, Q \in \Sigma, \beta\gamma_\Lambda(P, Q)^2 \leq \|\mathcal{A}P - \mathcal{A}Q\|_2^2. \quad (\text{B.46})$$

For a measurement operator \mathcal{A} drawn at random (in our case by randomly drawing a set of frequencies $\Omega = \{\omega_j\}_{j=1, \dots, m}$), the usual approach from compressive sensing theory is to prove that, with high probability, (B.46) is satisfied:

$$\mathbb{P}_\Omega \left(\forall P, Q \in \Sigma, \beta\gamma_\Lambda(P, Q)^2 \leq \|\mathcal{A}P - \mathcal{A}Q\|_2^2 \right) \geq 1 - \rho, \quad (\text{B.47})$$

where \mathbb{P}_Ω indicates probability with respect to the set of frequencies Ω .

Defining the *normalized secant set*

$$\mathcal{S} := \left\{ \frac{P - Q}{\gamma_\Lambda(P, Q)}; P, Q \in \Sigma, \gamma_\Lambda(P, Q) \neq 0 \right\} \subset E, \quad (\text{B.48})$$

the LRIP (B.46) is equivalent to : $\forall \nu \in \mathcal{S}, \beta \leq \|\mathcal{A}\nu\|_2^2$, and (B.47) is equivalent to

$$\mathbb{P}_\Omega (\forall \nu \in \mathcal{S}, \beta \leq \|\mathcal{A}\nu\|_2^2) \geq 1 - \rho. \quad (\text{B.49})$$

Hence, a typical proof of the (L)RIP [5, 81] consists in defining an ϵ -covering of the normalized secant set, proving a pointwise LRIP at the center of each ball using concentration results, then uniformly extending the result to the whole normalized secant set using Lipschitz continuity of the measurement operator.

Semi-uniform LRIP. In our framework, we introduce a “non-uniform” version of the LRIP, in which the inequality (B.47) will be verified for a *given* $P \in \Sigma$ with high probability, *uniformly* for all $Q \in \Sigma$. It is expressed as:

$$\forall P \in \Sigma, \mathbb{P}_\Omega (\forall Q \in \Sigma, \beta \gamma_\Lambda (P, Q)^2 \leq \|\mathcal{A}P - \mathcal{A}Q\|_2^2) \geq 1 - \rho. \quad (\text{B.50})$$

We refer to this version of the LRIP as *semi-uniform* in probability. It holds with a smaller number of measurements m than the uniform case, and we show in the next section that it is *sufficient* to obtain recovery guarantees with *joint* probability on the drawing of frequencies $\{\omega_j\}$ and items $\{\mathbf{x}_i\}$. For more details on non-uniform compressive sensing results, we refer the reader to the book by Foucart and Rauhut [52], Chaps. 9 and 11.

Similar to the uniform case, we introduce a family of “non-uniform” normalized secant sets, defined for each $P \in \Sigma$ as:

$$\mathcal{S}_P := \left\{ \frac{P - Q}{\gamma_\Lambda (P, Q)}; Q \in \Sigma, \gamma_\Lambda (P, Q) \neq 0 \right\} \subset E. \quad (\text{B.51})$$

The semi-uniform LRIP (B.50) is then equivalent to

$$\forall P \in \Sigma, \mathbb{P}_\Omega (\forall \nu \in \mathcal{S}_P, \beta \leq \|\mathcal{A}\nu\|_2^2) \geq 1 - \rho. \quad (\text{B.52})$$

A typical proof would therefore follow the exact same pattern than the uniform case, using non-uniform normalized secant sets instead of the normalized secant set.

Restricted, semi-uniform LRIP. Unlike finite dimensional frameworks, where normalized secant sets are contained in a unit ball that is necessarily compact, here it is in general challenging to prove the existence of finite covering numbers for this set. Under Assumption $\mathbf{A}_1(\Sigma)$, the model Σ *itself* is compact, which suggests using the embedding $Q \in \Sigma \rightarrow \frac{P - Q}{\gamma_\Lambda (P, Q)} \in \mathcal{S}_P$. However the behavior of this function when Q gets close to P may be delicate to analyze. Thus for all $\eta \geq 0$ and $P \in \Sigma$ we define the *restricted* non-uniform normalized secant set:

$$\mathcal{S}_P^\eta := \left\{ \frac{P - Q}{\gamma_\Lambda (P, Q)}; Q \in \Sigma, \gamma_\Lambda (P, Q) > \eta \right\} \subset E. \quad (\text{B.53})$$

Note that, when we let $\eta = 0$ the restricted non-uniform normalized secant set \mathcal{S}_P^0 is just the previous non-uniform normalized secant set \mathcal{S}_P .

Hypotheses to establish the restricted, semi-uniform LRIP We are going to prove the restricted semi-uniform LRIP (B.50) under two hypotheses. The first hypothesis depends on a model Σ , a frequency distribution Λ , a non-negative constant $\eta \geq 0$, and a choice of metric $d(\cdot, \cdot)$.

Hypothesis $\mathbf{H}_1(\eta, \Sigma, \Lambda, d)$ (Covering numbers of the secant set). *For all $P \in \Sigma$, the restricted non-uniform normalized secant set \mathcal{S}_P^η has finite covering numbers $N_{\mathcal{S}_P^\eta, d}(\epsilon) < \infty$.*

In the case where the constant $\eta > 0$ is positive and the metric $d = \|\cdot\|$ is a norm, the covering numbers of the secant set are controlled by those of the model.

Lemma B.8. *Let $\|\cdot\|$ be a norm on the space of finite signed measure E , and Λ a frequency distribution such that the model Σ has finite covering numbers with respect to some metric \bar{d} which satisfies*

$$\forall P, Q \in \Sigma, \bar{d}(P, Q) \geq \max \left(\|P - Q\|, \gamma_\Lambda (P, Q) \right).$$

The model is in particular bounded for the norm $\|\cdot\|$, denote $C = \max \left(1, \sup_{P, Q \in \Sigma} \|P - Q\| \right)$.

Then, for any strictly positive constant $1 \geq \eta > 0$, Assumption $\mathbf{H}_1(\eta, \Sigma, \Lambda, d)$ holds with $d = \|\cdot\|$. Furthermore, for any $P \in \Sigma$ and $\epsilon > 0$ we have

$$N_{\mathcal{S}_P^\eta, \|\cdot\|}(\epsilon) \leq N_{\Sigma, \bar{d}}\left(\frac{\epsilon \eta^2}{2(C+1)}\right). \quad (\text{B.54})$$

Proof. Let $P \in \Sigma$ be any distribution in the model. Consider the complement of the ball $B_{\Sigma, \gamma_\Lambda}(P, \eta)$:

$$\mathcal{Q}_P^\eta = B_{\Sigma, \gamma_\Lambda}^c(P, \eta) = \{Q \in \Sigma, \gamma_\Lambda(P, Q) > \eta\} \subset \Sigma, \quad (\text{B.55})$$

and the function $f_P : \mathcal{Q}_P^\eta \rightarrow \mathcal{S}_P^\eta$ such that $f_P(Q) = \frac{P-Q}{\gamma_\Lambda(P, Q)}$, which is surjective by definition of \mathcal{S}_P^η . Let us show that f_P is $(C+1)/\eta^2$ -Lipschitz continuous for the metric \bar{d} , and conclude with Lemma A.4.

For any $Q_1, Q_2 \in \mathcal{Q}_P^\eta$, we have

$$\begin{aligned} \|f_P(Q_1) - f_P(Q_2)\| &= \left\| \frac{P - Q_1}{\gamma_\Lambda(P, Q_1)} - \frac{P - Q_2}{\gamma_\Lambda(P, Q_2)} \right\|, \\ &\leq \left\| \frac{P - Q_1}{\gamma_\Lambda(P, Q_1)} - \frac{P - Q_2}{\gamma_\Lambda(P, Q_1)} \right\| + \left\| \frac{P - Q_2}{\gamma_\Lambda(P, Q_1)} - \frac{P - Q_2}{\gamma_\Lambda(P, Q_2)} \right\|, \\ &\leq \frac{1}{\eta} \|Q_2 - Q_1\| + \|P - Q_2\| \left| \frac{1}{\gamma_\Lambda(P, Q_1)} - \frac{1}{\gamma_\Lambda(P, Q_2)} \right|, \quad \text{since } \gamma_\Lambda(P, Q_1) > \eta \\ &\leq \frac{1}{\eta} \|Q_1 - Q_2\| + \frac{C}{\eta^2} \left| \gamma_\Lambda(P, Q_2) - \gamma_\Lambda(P, Q_1) \right|, \quad \text{since } \|P - Q_2\| \leq C \\ &\leq \frac{1}{\eta} \|Q_1 - Q_2\| + \frac{C}{\eta^2} \gamma_\Lambda(Q_1, Q_2), \quad \text{by the triangle inequality,} \\ &\leq \frac{C+1}{\eta^2} \bar{d}(Q_1, Q_2). \quad \text{since } \eta \leq 1 \end{aligned}$$

Hence the function f_P is Lipschitz continuous with constant $L = (C+1)/\eta^2$, and therefore for all $\epsilon > 0$:

$$N_{\mathcal{S}_P^\eta, \|\cdot\|}(\epsilon) \stackrel{\text{Lemma A.4}}{\leq} N_{\mathcal{Q}_P^\eta, \bar{d}}(\epsilon/L) \stackrel{\text{Lemma A.3}}{\leq} N_{\Sigma, \bar{d}}\left(\frac{\epsilon}{2L}\right). \quad \square$$

To formulate the second hypothesis, we denote f a function from $\mathbb{N} \times \mathbb{R}_+$ to \mathbb{R}_+ .

Hypothesis $\mathbf{H}_2(\eta, \Sigma, \Lambda, f)$ (Probability of the pointwise LRIP). For any $P \in \Sigma$, any $\nu \in \mathcal{S}_P^\eta$, any $t \geq 0$ and any integer $m > 0$, we have

$$\mathbb{P}_\Omega(1 - \|\mathcal{A}\nu\|_2^2 \geq t) \leq f(m, t), \quad (\text{B.56})$$

where $\mathcal{A} : E \rightarrow \mathbb{C}^m$ is a sketching operator built by independently drawing m frequencies according to Λ .

In Section B.2, under hypotheses $\mathbf{H}_1(\eta, \Sigma, \Lambda, d)$ and $\mathbf{H}_2(\eta, \Sigma, \Lambda, f)$, we prove an extended version of Theorem 3 (referred to as Theorem 3 bis). Unlike Theorem 3 this extended version covers the case $\eta = 0$.

Then, in Section B.3 we prove that under the Assumptions $\mathbf{A}_1(\Sigma)$ and $\mathbf{A}_2(\eta, \Sigma, \Lambda, \Lambda)$ used to state Theorem 3, $\mathbf{H}_1(\eta, \Sigma, \Lambda, d)$ holds with $d(\nu, \nu') = \|\nu - \nu'\|_{TV}$ provided that $\eta > 0$, and $\mathbf{H}_2(\eta, \Sigma, \Lambda, f)$ holds for an appropriate choice of function f .

Remark B.2. Ideally, one would like to exploit Theorem 3 bis with $\eta = 0$ to obtain performance guarantees without the extra additive term η . However, putting this into practice would require characterizing covering numbers for the normalized secant set $\mathcal{S}_P = \mathcal{S}_P^0$, which can be tricky and is left to future work. It is indeed already not trivial to determine when this set has finite covering numbers.

We can now state our version of the LRIP.

Theorem B.6. Consider a model Σ , a frequency distribution Λ , a non-negative constant $\eta \geq 0$, a metric $d(\cdot, \cdot)$ and a function f such that Assumptions $\mathbf{H}_1(\eta, \Sigma, \Lambda, d)$ and $\mathbf{H}_2(\eta, \Sigma, \Lambda, f)$ are satisfied.

Let $P^* \in \Sigma$ be any distribution in the model.

Assume that for all $\nu, \nu' \in \mathcal{S}_{P^*}^\eta$

$$\sup_{\mathcal{A}} \left| \|\mathcal{A}\nu\|_2 - \|\mathcal{A}\nu'\|_2 \right| \leq d(\nu, \nu'), \quad (\text{B.57})$$

where the supremum is over all possible frequencies ω defining the sketching operator \mathcal{A} .

Define

$$\rho = N_{\mathcal{S}_{P^*,d}^\eta} \left(\frac{1}{4} \right) \cdot f \left(m, \frac{7}{16} \right).$$

Then, with probability at least $1 - \rho$ on the drawing of the m frequencies ω_j 's, we have

$$\forall P \in \Sigma \text{ s.t. } \gamma_\Lambda(P^*, P) \geq \eta: \quad \frac{1}{4} \gamma_\Lambda(P^*, P)^2 \leq \|\mathcal{A}P^* - \mathcal{A}P\|_2^2. \quad (\text{B.58})$$

Proof. The idea is to define an ϵ -covering of the restricted non-uniform normalized secant set $\mathcal{S}_{P^*}^\eta$ with respect to the metric d , to apply the concentration result of Assumption $\mathbf{H}_2(\eta, \Sigma, \Lambda, f)$ at the center of each ball and to conclude with a union bound.

Let $\epsilon > 0$, $1 > t > 0$ and denote $N = N_{\mathcal{S}_{P^*,d}^\eta}(\epsilon)$ for simplicity, which is finite by Assumption $\mathbf{H}_1(\eta, \Sigma, \Lambda, d)$. We consider ν_1, \dots, ν_N an ϵ -covering of $\mathcal{S}_{P^*}^\eta$ with respect to the metric d . Considering Assumption $\mathbf{H}_2(\eta, \Sigma, \Lambda, f)$, a union bound yields that, with probability greater than $1 - N \cdot f(m, t)$,

$$\forall i \in [1, N], \quad 1 - \|\mathcal{A}\nu_i\|_2^2 < t. \quad (\text{B.59})$$

Assuming (B.59) is satisfied, let any distribution $P \in \Sigma$ such that $\gamma_\Lambda(P^*, P) > \eta$. Denote $\nu := \frac{P^* - P}{\gamma_\Lambda(P^*, P)} \in \mathcal{S}_{P^*}^\eta$, and let $i \in [1, N]$ such that ν_i is the center of the ball closest to ν in the covering. We have

$$\begin{aligned} 1 - \|\mathcal{A}\nu\|_2 &= 1 - \|\mathcal{A}\nu_i\|_2 + \|\mathcal{A}\nu_i\|_2 - \|\mathcal{A}\nu\|_2 \\ &\leq 1 - \sqrt{1-t} + \|\mathcal{A}\nu_i\|_2 - \|\mathcal{A}\nu\|_2 && \text{since (B.59) is verified} \\ &\stackrel{(a)}{\leq} 1 - \sqrt{1-t} + d(\nu_i, \nu) && \text{by (B.57)} \\ &\leq 1 - \sqrt{1-t} + \epsilon. \end{aligned}$$

Choosing $t > 0$ and $\epsilon = \sqrt{1-t} - 1/2$ (for example: $t = 7/16$, $\epsilon = 1/4$), we obtain $\|\mathcal{A}\nu\|_2^2 \geq 1/4$, that is to say

$$\frac{1}{4} \gamma_\Lambda(P^*, P)^2 \leq \|\mathcal{A}P^* - \mathcal{A}P\|_2^2. \quad (\text{B.60})$$

This shows that the LRIP (B.58) is verified except with probability at most

$$N_{\mathcal{S}_{P^*,d}^\eta}(\sqrt{1-t} - 1/2) \cdot f(m, t).$$

Specializing to $t = 7/16$ yields the desired result. \square

Remark B.3 (LRIP without the restricted Lipschitz property (B.57) ?). *In (a) we used Property (B.57), which could be called a restricted Lipschitz property for the function $\nu \mapsto \|\mathcal{A}\nu\|_2$. It is restricted to $\mathcal{S}_{P^*}^\eta$, but assumed to hold uniformly over all possible draws of the sketching operator \mathcal{A} . Thanks to Lemma A.2 and the triangle inequality,*

$$\left| \|\mathcal{A}\nu\|_2 - \|\mathcal{A}\nu'\|_2 \right| \leq \|\mathcal{A}(\nu - \nu')\|_2 \leq \|\nu - \nu'\|_{TV},$$

hence property (B.57) indeed holds when $d(\nu, \nu') = \|\nu - \nu'\|_{TV}$, even without the restriction to $\mathcal{S}_{P^*}^\eta$. In the rest of this paper we primarily concentrate on this setting.

It would however be interesting for future work to consider metrics d with much larger balls than those of the total variation norm, since they may lead to substantially smaller covering numbers $N_{\mathcal{S}_{P^*,\|\cdot\|}^\eta} \ll N_{\mathcal{S}_{P^*,\|\cdot\|_{TV}}^\eta}$, and thus LRIP guarantees for much smaller m . Theorem B.6 still yields guarantees with such metrics, provided the restricted Lipschitz property (B.57) holds. An analog of Theorem B.6 for even weaker metrics, that do not satisfy (B.57), can also be envisioned using chaining arguments [81] provided that

$$\mathbb{P}_\Omega(\|\mathcal{A}\nu - \mathcal{A}\nu'\|_2 \geq (1+t)d(\nu, \nu')) \leq g(m, t)$$

with an appropriately decaying function $g(m, t)$.

B.2 A version of Theorem 3 with weaker assumptions

In this section, we formulate a version of Theorem 3, referred to as Theorem 3 bis, that allows for an additive error $\eta = 0$, under the Hypotheses \mathbf{H}_1 . In the next section we deduce from it the version given in Section 7.1, that uses Assumptions \mathbf{A}_1 .

Theorem 3 bis. *Consider a model Σ , a frequency distribution Λ , a non-negative constant $\eta \geq 0$ and a function f such that Assumptions $\mathbf{H}_1(\eta, \Sigma, \Lambda, d)$ and $\mathbf{H}_2(\eta, \Sigma, \Lambda, f)$ hold with $d(\nu, \nu') = \|\nu - \nu'\|_{TV}$.*

Consider $\tilde{d}(\cdot, \cdot)$ a metric such that, for any $P_1, P_2 \in \mathcal{P}$ we have

$$\tilde{d}(P_1, P_2) \geq \max \left(\gamma_\Lambda(P_1, P_2), \sup_{\mathcal{A}} \|\mathcal{A}P_1 - \mathcal{A}P_2\|_2 \right), \quad (\text{B.61})$$

where the supremum is over all possible frequencies ω defining the sketching operator \mathcal{A} .

Assume that Σ is compact⁹ with respect to \tilde{d} , and note that under this assumption the decoder Δ is still well-defined by (30), since one can replace the total variation norm by the metric \tilde{d} in the r.h.s. of (29).

Let $\mathbf{x}_i \in \mathbb{R}^d$, $i = 1 \dots n$ be n points drawn i.i.d. from an arbitrary distribution $P^* \in \mathcal{P}$, and $\omega_j \in \mathbb{R}^d$, $j = 1 \dots m$ be m frequencies drawn i.i.d. from Λ . Denote $\bar{P} = \Delta(\tilde{\mathbf{z}}, \mathcal{A})$ the distribution reconstructed from the empirical sketch $\tilde{\mathbf{z}}$.

Define $P_{\text{proj}} \in \Sigma$ as (one of) the projection(s) of the probability P^* onto the model:

$$P_{\text{proj}} \in \underset{P \in \Sigma}{\operatorname{argmin}} \tilde{d}(P^*, P), \quad (\text{B.62})$$

which indeed exists since Σ is assumed to be compact.

Define

$$\rho = 2N_{S_{P_{\text{proj}}}, \|\cdot\|_{TV}} \left(\frac{1}{4} \right) \cdot f \left(m, \frac{7}{16} \right). \quad (\text{B.63})$$

Then, with probability at least $1 - \rho$ on the drawing of the items \mathbf{x}_i and sampling frequencies ω_j , we have

$$\gamma_\Lambda(P^*, \bar{P}) \leq 5 \tilde{d}(P^*, \Sigma) + \frac{4(1 + \sqrt{2 \log(2/\rho)})}{\sqrt{n}} + \eta, \quad (\text{B.64})$$

where $\tilde{d}(P^*, \Sigma) = \inf_{P \in \Sigma} \tilde{d}(P^*, P)$ is the distance from P^* to the model.

Proof. Recall that the target distribution P^* and its projection P_{proj} are fixed. By Lemma A.2, the restricted Lipschitz property (B.57) holds with $d(\nu, \nu') = \|\nu - \nu'\|_{TV}$. Considering (B.63), Theorem B.6 yields that since Assumptions $\mathbf{H}_1(\eta, \Sigma, \Lambda, d)$ and $\mathbf{H}_2(\eta, \Sigma, \Lambda, f)$ hold, the LRIP applied to P_{proj} is satisfied with probability at least $1 - \rho/2$ on the drawing of frequencies:

$$\forall P_\Sigma \in \Sigma \text{ s.t. } \gamma_\Lambda(P_{\text{proj}}, P_\Sigma) \geq \eta : \frac{1}{4} \gamma_\Lambda(P_{\text{proj}}, P_\Sigma)^2 \leq \|\mathcal{A}P_{\text{proj}} - \mathcal{A}P_\Sigma\|_2^2,$$

which can be reformulated in:

$$\forall P_\Sigma \in \Sigma, \gamma_\Lambda(P_{\text{proj}}, P_\Sigma) \leq \max \left(2\|\mathcal{A}P^* - \mathcal{A}P_\Sigma\|_2, \eta \right) \leq 2\|\mathcal{A}P^* - \mathcal{A}P_\Sigma\|_2 + \eta. \quad (\text{B.65})$$

Let $P \in \mathcal{P}$ be any distribution. For some random draw of the operator \mathcal{A} , denote $\bar{P} = \Delta(\mathcal{A}P, \mathcal{A}) \in \Sigma$ which, by definition of the decoder, belongs to the model. We have:

$$\begin{aligned} \gamma_\Lambda(\bar{P}, P^*) &\leq \gamma_\Lambda(\bar{P}, P_{\text{proj}}) + \gamma_\Lambda(P_{\text{proj}}, P^*) && \text{triangle ineq.} \\ &\leq 2\|\mathcal{A}\bar{P} - \mathcal{A}P_{\text{proj}}\|_2 + \gamma_\Lambda(P_{\text{proj}}, P^*) + \eta && \text{using (B.65)} \\ &\leq 2\|\mathcal{A}\bar{P} - \mathcal{A}P\|_2 + 2\|\mathcal{A}P - \mathcal{A}P^*\|_2 + 2\|\mathcal{A}P^* - \mathcal{A}P_{\text{proj}}\|_2 + \gamma_\Lambda(P_{\text{proj}}, P^*) + \eta && \text{triangle ineq.} \end{aligned}$$

Given the definition (30) of the decoder Δ and of the distribution $\bar{P} = \Delta(\mathcal{A}P, \mathcal{A})$, we have

$$\|\mathcal{A}\bar{P} - \mathcal{A}P\|_2 = \min_{P_\Sigma \in \Sigma} \|\mathcal{A}P_\Sigma - \mathcal{A}P\|_2. \quad (\text{B.66})$$

⁹Compactness of the model is still required here, in order to define a projection operator onto it as well as to ensure at least one solution to the problem (26). This assumption could be relaxed, with the addition of an arbitrary small additive error, similar to [18].

Since P_{proj} is in the model Σ , we thus have $\|\mathcal{A}\bar{P} - \mathcal{A}P\|_2 \leq \|\mathcal{A}P_{\text{proj}} - \mathcal{A}P\|_2$. Hence

$$\begin{aligned} \gamma_\Lambda(\bar{P}, P^*) &\leq 2\|\mathcal{A}P_{\text{proj}} - \mathcal{A}P\|_2 + 2\|\mathcal{A}P - \mathcal{A}P^*\|_2 + 2\|\mathcal{A}P^* - \mathcal{A}P_{\text{proj}}\|_2 + \gamma_\Lambda(P_{\text{proj}}, P^*) + \eta \\ &\leq 4\|\mathcal{A}P - \mathcal{A}P^*\|_2 + 4\|\mathcal{A}P^* - \mathcal{A}P_{\text{proj}}\|_2 + \gamma_\Lambda(P_{\text{proj}}, P^*) \quad \text{triangle ineq.} \\ &\stackrel{(b)}{\leq} 4\|\mathcal{A}P - \mathcal{A}P^*\|_2 + 5\tilde{d}(P^*, \Sigma) + \eta. \quad \text{using (B.61) and (B.62)} \end{aligned}$$

Thus we proved that, with probability at least $1 - \rho/2$ on the drawing of frequencies, we have:

$$\forall P \in \mathcal{P}, \quad \gamma_\Lambda(P^*, \Delta(\mathcal{A}P, \mathcal{A})) \leq 5\tilde{d}(P^*, \Sigma) + 4\|\mathcal{A}P - \mathcal{A}P^*\|_2 + \eta.$$

In particular, with a joint probability of at least $1 - \rho/2$ on the drawing of frequencies $\boldsymbol{\omega}_j$ and items \mathbf{x}_i , we have

$$\gamma_\Lambda(P^*, \Delta(\mathcal{A}\hat{P}, \mathcal{A})) \leq 5\tilde{d}(P^*, \Sigma) + 4\|\mathcal{A}\hat{P} - \mathcal{A}P^*\|_2 + \eta, \quad (\text{B.67})$$

where $\hat{P} = \frac{1}{n} \sum_{i=1}^n \delta_{\mathbf{x}_i}$.

We now show that with high probability the term $\|\mathcal{A}\hat{P} - \mathcal{A}P^*\|_2$ is bounded by $\epsilon := \frac{(1 + \sqrt{2 \log(2/\rho)})}{\sqrt{n}}$. Denote \mathbb{P}_Ω (resp. $\mathbb{P}_\mathcal{X}$) the probability distribution of the set of frequencies $\Omega = \{\boldsymbol{\omega}_j\}_{j=1, \dots, m}$ (resp. of the set of items in the database $\mathcal{X} = \{\mathbf{x}_i\}_{i=1, \dots, n}$). Their joint distribution is denoted $\mathbb{P}_{\Omega, \mathcal{X}}$, and is such that $d\mathbb{P}_{\Omega, \mathcal{X}}(\Omega, \mathcal{X}) = d\mathbb{P}_\Omega(\Omega)d\mathbb{P}_\mathcal{X}(\mathcal{X})$ by independence.

Consider the set:

$$A := \left\{ (\Omega, \mathcal{X}) \text{ s.t. } \|\mathcal{A}\hat{P} - \mathcal{A}P^*\|_2 \leq \epsilon \right\}.$$

For a *fixed* measurement operator \mathcal{A} , we use Lemma A.7 on the random variables $\mathcal{A}\delta_{\mathbf{x}_i}$ in \mathbb{C}^m . We observe that $\|\mathcal{A}\delta_{\mathbf{x}_i}\|_2 = 1$, $\mathcal{A}\hat{P} = (\sum_{i=1}^n \mathcal{A}\delta_{\mathbf{x}_i})/n$ and $\mathcal{A}P^* = \frac{1}{\sqrt{m}} \left[\mathbb{E}_{\mathbf{x} \sim P^*} e^{-i\boldsymbol{\omega}_j^T \mathbf{x}} \right]_{j=1 \dots m} = \mathbb{E}_{\mathbf{x} \sim P^*} \mathcal{A}\delta_{\mathbf{x}}$. Hence, with probability at least $1 - \rho/2$ on the drawing of items:

$$\|\mathcal{A}\hat{P} - \mathcal{A}P^*\|_2 \leq \epsilon.$$

In other words, we obtain the following result:

$$\forall \Omega, \quad \int_{\mathcal{X}} \mathbf{1}_A(\Omega, \mathcal{X}) d\mathbb{P}_\mathcal{X}(\mathcal{X}) \geq 1 - \rho/2. \quad (\text{B.68})$$

Hence,

$$\begin{aligned} \iint_{\Omega, \mathcal{X}} \mathbf{1}_A(\Omega, \mathcal{X}) d\mathbb{P}_{\Omega, \mathcal{X}}(\Omega, \mathcal{X}) &= \int_{\Omega} \left(\int_{\mathcal{X}} \mathbf{1}_A(\Omega, \mathcal{X}) d\mathbb{P}_\mathcal{X}(\mathcal{X}) \right) d\mathbb{P}_\Omega(\Omega) \\ &\geq (1 - \rho/2) \int_{\Omega} d\mathbb{P}_\Omega(\Omega) = 1 - \rho/2, \end{aligned}$$

meaning that, with probability at least $1 - \rho/2$ on the drawing of frequencies and items, we have

$$\|\mathcal{A}\hat{P} - \mathcal{A}P^*\|_2 \leq \epsilon. \quad (\text{B.69})$$

We can now conclude: a union bound yields that (B.67) and (B.69) are simultaneously satisfied with probability at least $1 - \rho$, which leads to the desired result. \square

Remark B.4. In (b) we used inequality (B.61), in particular the assumption that the inequality

$$\tilde{d}(P_1, P_2) \geq \|\mathcal{A}P_1 - \mathcal{A}P_2\|_2 \quad (\text{B.70})$$

holds uniformly over all possible choices of frequencies defining \mathcal{A} . Thanks to Lemma A.2, the inequality (B.61) indeed holds with $\tilde{d} = d_{TV}$. In the rest of this paper, we concentrate on this setting. It would however be interesting for future work to consider metrics \tilde{d} much closer to the natural choice $\gamma_\Lambda(\cdot, \cdot)$, since they may lead to sharper upper bounds. A possibility would be to relax (B.61) and only assume that inequality (B.70) (possibly up to multiplicative constants) holds with high probability on the draw of \mathcal{A} , given a pair $P_1 = P^*$ and $P_2 = P_{\text{proj}}$.

B.3 Proof of Theorem 3

We now turn to the proof of Theorem 3, which consists in proving that the Assumptions \mathbf{A}_i imply the hypotheses \mathbf{H}_i in the $\eta > 0$ case, and applying Theorem 3 bis.

For Hypothesis $\mathbf{H}_1(\eta, \Sigma, \Lambda, d)$, we apply Lemma B.8 with the total variation norm $\|\cdot\| = \bar{d} = \|\cdot\|_{TV}$ to obtain the following corollary.

Corollary B.4. *Let a model Σ such that Assumption $\mathbf{A}_1(\Sigma)$ is satisfied, and a positive constant $\eta > 0$. Then, for any frequency distribution Λ , Hypothesis $\mathbf{H}_1(\eta, \Sigma, \Lambda, d)$ is satisfied with $d = \|\cdot\|_{TV}$, and we have*

$$N_{\mathcal{S}_P^\eta, \|\cdot\|_{TV}}(\epsilon) \leq N_{\Sigma, \|\cdot\|_{TV}}\left(\frac{\epsilon\eta^2}{6}\right). \quad (\text{B.71})$$

Hypothesis $\mathbf{H}_2(\eta, \Sigma, \Lambda, f)$ is an application of Bernstein's inequality.

Lemma B.9. *Consider a model Σ , a frequency distribution Λ , a non-negative constant $\eta \geq 0$ and a constant A such that Assumption $\mathbf{A}_2(\eta, \Sigma, \Lambda, A)$ is satisfied.*

Then Assumption $\mathbf{H}_2(\eta, \Sigma, \Lambda, f)$ is satisfied with the function f defined as

$$f(m, t) = \exp\left(-\frac{m}{2A^2} \cdot \frac{t^2}{1+t/3}\right). \quad (\text{B.72})$$

Proof. Fix $P^* \in \Sigma$. Suppose ω_j , $j = 1..m$ are drawn *i.i.d.* from Λ and let \mathcal{A} be the corresponding sketching operator. Let any $\nu \in \mathcal{S}_{P^*}^\eta$, denote $P \in \Sigma$ such that $\nu = \frac{P^* - P}{\gamma_\Lambda(P^*, P)}$. Denote $Z_j = 1 - \frac{|\psi_{P^*}(\omega_j) - \psi_P(\omega_j)|^2}{\gamma_\Lambda(P^*, P)^2}$. Since Assumption $\mathbf{A}_2(\eta, \Sigma, \Lambda, A)$ is verified and $|\psi_{P^*}(\omega_j) - \psi_P(\omega_j)| \leq \|P^* - P\|_{TV}$ for all frequencies, the Z_j 's are *i.i.d.* random variables verifying $Z_j \in [1 - A^2; 1]$. Furthermore, according to Lemma A.2 we have necessarily $A \geq 1$ and thus we have

$$|Z_j| \leq A^2. \quad (\text{B.73})$$

The Z_j 's are also centered:

$$\mathbb{E}_{\omega_j} Z_j = 1 - \frac{\mathbb{E}_{\omega \sim \Lambda} |\psi_{P^*}(\omega) - \psi_P(\omega)|^2}{\gamma_\Lambda(P^*, P)^2} = 0. \quad \text{using (22)}$$

Furthermore, we have

$$\begin{aligned} \text{Var}(Z_j) &= \text{Var}\left(\frac{|\psi_{P^*}(\omega) - \psi_P(\omega)|^2}{\gamma_\Lambda(P^*, P)^2}\right) \leq \frac{\mathbb{E}|\psi_{P^*}(\omega) - \psi_P(\omega)|^4}{\gamma_\Lambda(P^*, P)^4} \leq \frac{\|P^* - P\|_{TV}^2 \cdot \mathbb{E}|\psi_{P^*}(\omega) - \psi_P(\omega)|^2}{\gamma_\Lambda(P^*, P)^2 \cdot \gamma_\Lambda(P^*, P)^2} \\ &\leq A^2 \cdot \frac{\mathbb{E}|\psi_{P^*}(\omega) - \psi_P(\omega)|^2}{\gamma_\Lambda(P^*, P)^2} = A^2. \end{aligned}$$

Since $\frac{1}{m} \sum_{j=1}^m Z_j = 1 - \|\mathcal{A}\nu\|_2^2$, applying Bernstein's inequality (Lemma A.6) we get : for all $t > 0$,

$$\forall \nu \in \mathcal{S}_{P^*}^\eta, \mathbb{P}_\Omega(1 - \|\mathcal{A}\nu\|_2^2 \geq t) = \mathbb{P}_\Omega\left(\frac{1}{m} \sum_{j=1}^m Z_j \geq t\right) \leq \exp\left(-\frac{m}{2A^2} \cdot \frac{t^2}{1+t/3}\right). \quad (\text{B.74})$$

□

We can finally prove Theorem 3, by combining Corollary B.4, Lemma B.9 and Theorem 3 bis. For simplification, we also use the following bound on the function f defined in Lemma B.9:

$$f(m, 7/16) = \exp\left(-\frac{147m}{1760A^2}\right) \leq \exp\left(-\frac{m}{12A^2}\right).$$

C Application to GMMs

C.1 Compactness of mixture models

In this section we prove Lemma 1, which extends the compactness of the set \mathcal{G} of basic distributions to the corresponding mixture model \mathcal{G}_K .

Proof of Lemma 1. Recall that we have assumed compactness of the set \mathcal{G} with respect to some norm $\|\cdot\|$. In particular, it is bounded, and we note $C = \max_{P \in \mathcal{G}} \|P\|$. Let $\epsilon > 0$ and $\tau \in]0; 1[$. Denote $\epsilon_1 = \tau\epsilon$ and $\epsilon_2 = (1 - \tau)\epsilon/C$. Also denote $N_1 = N_{\mathcal{G}, \|\cdot\|}(\epsilon_1)$ and let $\mathcal{N}_1 = \{P_1, \dots, P_{N_1}\}$ be an ϵ_1 -covering of \mathcal{G} .

Similarly, let $B_1^+ = \{\alpha \in \mathbb{R}_+^K, \sum_{k=1}^K \alpha_k = 1\}$, denote $N_2 = N_{B_1^+, \|\cdot\|_1}(\epsilon_2)$, let $\mathcal{N}_2 = \{\alpha_1, \dots, \alpha_{N_2}\}$ be an ϵ_2 -covering of B_1^+ . Denote $B_1 := B_{\mathbb{R}, \|\cdot\|_1}(0, 1)$ the unit ℓ_1 -ball in \mathbb{R}^K , note that $B_1^+ \subset B_1$, such that we have

$$N_2 = N_{B_1^+, \|\cdot\|_1}(\epsilon_2) \stackrel{\text{Lemma A.3}}{\leq} N_{B_1, \|\cdot\|_1}(\epsilon_2/2) \stackrel{\text{Lemma A.5}}{\leq} (8/\epsilon_2)^K. \quad (\text{C.75})$$

Define the following set:

$$\Gamma := \{P_{\Theta, \alpha} \in \mathcal{G}_K; \forall k, P_{\theta_k} \in \mathcal{N}_1, \alpha \in \mathcal{N}_2\}. \quad (\text{C.76})$$

The cardinality of this set verifies $|\Gamma| \leq (|\mathcal{N}_1|)^K |\mathcal{N}_2| = (N_1)^K N_2$.

Let us show that Γ is an ϵ -covering of \mathcal{G}_K . Let $P_{\Theta, \alpha} \in \mathcal{G}_K$ be any K -sparse distribution. For all $k = 1 \dots K$, let $P_{\theta_k} \in \mathcal{N}_1$ be the distribution in \mathcal{N}_1 which is the closest to P_{θ_k} , and let $\bar{\alpha} \in \mathcal{N}_2$ be the weight vector in \mathcal{N}_2 that is the closest to α . Denote $\bar{\Theta} = \{\bar{\theta}_1, \dots, \bar{\theta}_K\}$, and note that $P_{\bar{\Theta}, \bar{\alpha}} \in \Gamma$. We have

$$\begin{aligned} \|P_{\Theta, \alpha} - P_{\bar{\Theta}, \bar{\alpha}}\| &= \left\| \sum_{k=1}^K \alpha_k P_{\theta_k} - \sum_{k=1}^K \bar{\alpha}_k P_{\bar{\theta}_k} \right\|, \\ &\leq \left\| \sum_{k=1}^K \alpha_k P_{\theta_k} - \sum_{k=1}^K \alpha_k P_{\bar{\theta}_k} \right\| + \left\| \sum_{k=1}^K \alpha_k P_{\bar{\theta}_k} - \sum_{k=1}^K \bar{\alpha}_k P_{\bar{\theta}_k} \right\|, \\ &\leq \sum_{k=1}^K \alpha_k \|P_{\theta_k} - P_{\bar{\theta}_k}\| + \sum_{k=1}^K |\alpha_k - \bar{\alpha}_k| \|P_{\bar{\theta}_k}\|, \\ &\leq \sum_{k=1}^K \alpha_k \|P_{\theta_k} - P_{\bar{\theta}_k}\| + C \|\alpha - \bar{\alpha}\|_1, \\ &\leq \epsilon_1 \sum_{k=1}^K \alpha_k + C \epsilon_2 = \epsilon_1 + C \epsilon_2 = \epsilon, \end{aligned} \quad (\text{C.77})$$

and Γ is indeed an ϵ -covering of \mathcal{G}_K . Therefore, we have the bound (for all τ)

$$N_{\mathcal{G}_K, \|\cdot\|}(\epsilon) \leq |\Gamma| \leq (N_1)^K N_2 \stackrel{\text{by (C.75)}}{\leq} \left(\frac{8C \cdot N_{\mathcal{G}, \|\cdot\|}(\tau\epsilon)}{(1 - \tau)\epsilon} \right)^K.$$

Furthermore, in equation (C.77), we have shown in particular that for all $P_{\Theta, \alpha}, P_{\Theta', \alpha'} \in \mathcal{G}_K$,

$$\|P_{\Theta, \alpha} - P_{\Theta', \alpha'}\| \leq \sum_{k=1}^K \alpha_k \|P_{\theta_k} - P_{\theta'_k}\| + \|\alpha - \alpha'\|_1$$

and therefore the embedding $(P_{\theta_1}, \dots, P_{\theta_K}, \alpha) \rightarrow P_{\Theta, \alpha}$ from $(\mathcal{G})^K \times B_1^+$ to \mathcal{G}_K is continuous. Hence \mathcal{G}_K is the continuous image of the set $(\mathcal{G})^K \times B_1^+$, which is compact since \mathcal{G} is compact, and therefore \mathcal{G}_K is compact. \square

C.2 Covering numbers of Gaussians

Proof of Theorem 4. Consider the embedding $\phi : \mathcal{T} \rightarrow \mathcal{G}$ defined as $\phi(\theta) = P_{\theta}$, which is surjective by definition of \mathcal{G} . We show that ϕ is Lipschitz continuous, for the Euclidean norm on $\mathcal{T} \subset \mathbb{R}^{2n}$ and total variation norm on $\mathcal{G} \subset E$.

We begin by the classical Pinsker's inequality [48]:

$$\|P - Q\|_{TV} \leq \sqrt{2D_{KL}(P||Q)}, \quad (\text{C.78})$$

where D_{KL} is the Kullback-Leibler divergence. By symmetry, we have:

$$\|P - Q\|_{TV}^2 \leq D_{KL}(P||Q) + D_{KL}(Q||P). \quad (\text{C.79})$$

The Kullback-Leibler divergence has a closed-form expression in the case of multivariate Gaussians [47]:

$$D_{KL}(P_{\theta_1}||P_{\theta_2}) = \frac{1}{2} \left[\log \frac{|\Sigma_2|}{|\Sigma_1|} + \text{tr}(\Sigma_2^{-1}\Sigma_1) - d + (\boldsymbol{\mu}_2 - \boldsymbol{\mu}_1)^T \Sigma_2^{-1} (\boldsymbol{\mu}_2 - \boldsymbol{\mu}_1) \right]. \quad (\text{C.80})$$

In our case, with diagonal Gaussians and bounded parameters, we have

$$\begin{aligned} D_{KL}(P_{\theta_1}||P_{\theta_2}) + D_{KL}(P_{\theta_1}||P_{\theta_2}) &= \frac{\text{tr}(\Sigma_2^{-1}\Sigma_1) + \text{tr}(\Sigma_1^{-1}\Sigma_2)}{2} - d + (\boldsymbol{\mu}_2 - \boldsymbol{\mu}_1)^T \frac{\Sigma_2^{-1} + \Sigma_1^{-1}}{2} (\boldsymbol{\mu}_2 - \boldsymbol{\mu}_1) \\ &= \frac{1}{2} \sum_{\ell=1}^d \left(\frac{\sigma_{1,\ell}^2}{\sigma_{2,\ell}^2} + \frac{\sigma_{2,\ell}^2}{\sigma_{1,\ell}^2} - 2 \right) + \sum_{\ell=1}^d \frac{\sigma_{1,\ell}^{-2} + \sigma_{2,\ell}^{-2}}{2} (\mu_{2,\ell} - \mu_{1,\ell})^2, \\ &\leq \frac{1}{2} \sum_{\ell=1}^d \left(\frac{\sigma_{2,\ell}^4 + \sigma_{1,\ell}^4 - 2\sigma_{1,\ell}^2\sigma_{2,\ell}^2}{\sigma_{1,\ell}^2\sigma_{2,\ell}^2} \right) + \frac{1}{\sigma_{\min}^2} \|\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2\|_2^2, \\ &\leq \frac{1}{2\sigma_{\min}^4} \sum_{\ell=1}^d (\sigma_{1,\ell}^2 - \sigma_{2,\ell}^2)^2 + \frac{1}{\sigma_{\min}^2} \|\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2\|_2^2, \\ &\leq \frac{1}{2\sigma_{\min}^4} \|\boldsymbol{\sigma}_2 - \boldsymbol{\sigma}_1\|_2^2 + \frac{1}{\sigma_{\min}^2} \|\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2\|_2^2 \leq L^2 \|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\|_2^2, \end{aligned}$$

where $L := \max(\sigma_{\min}^{-1}, \sigma_{\min}^{-2}/\sqrt{2})$. Hence

$$\|P_{\theta_1} - P_{\theta_2}\|_{TV} \leq L \|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\|_2, \quad (\text{C.81})$$

and the embedding ϕ is L -Lipschitz. Hence \mathcal{G} is the continuous image of a compact set and is compact.

Since ϕ is also surjective, we can apply Lemma A.4 and conclude: denote $\mathcal{B} \subseteq \mathbb{R}^{2d}$ a ball of radius radius (\mathcal{T}) for the Euclidean norm such that $\mathcal{T} \subseteq \mathcal{B}$, and we have

$$N_{\mathcal{G}, \|\cdot\|_{TV}}(\epsilon) \stackrel{\text{Lemma A.4}}{\leq} N_{\mathcal{T}, \|\cdot\|_2}(\epsilon/L) \stackrel{\text{Lemma A.3}}{\leq} N_{\mathcal{B}, \|\cdot\|_2} \left(\frac{\epsilon}{2L} \right) \stackrel{\text{Lemma A.5}}{\leq} \left(\frac{8L \text{radius}(\mathcal{T})}{\epsilon} \right)^{2d} = \left(\frac{B}{\epsilon} \right)^{2d}, \quad (\text{C.82})$$

which is the desired result. \square

Proof of Corollary 2. Combining Theorem 4 and Lemma 1 proves that the set of GMMs \mathcal{G}_K is compact. Furthermore, we obtain the following bound, for all $0 < \tau < 1$:

$$N_{\mathcal{G}_K, \|\cdot\|_{TV}}(\epsilon) \leq \frac{B^{2dK} 8^K}{\tau^{2dK} (1-\tau)^K \epsilon^{(2d+1)K}},$$

where B is defined as in Theorem 4. By choosing¹⁰ $\tau = \frac{B}{B+1}$, we obtain

$$N_{\mathcal{G}_K, \|\cdot\|_{TV}}(\epsilon) \leq \left(\frac{B+1}{\epsilon} \right)^{(2d+1)K} 8^K \leq \left(\frac{2(B+1)}{\epsilon} \right)^{(2d+1)K},$$

since $8^{\frac{1}{2d+1}} \leq 8^{1/3} = 2$. \square

C.3 Domination between metrics on the set of Gaussians

In this section, we aim at proving Theorem 5. We begin by an intermediate result.

Lemma C.10. *Suppose that $\mathcal{T} \subset \mathbb{R}^{2d}$ is such that $\|\boldsymbol{\mu}\|_2 \leq M$ and $0 < \sigma_{\min}^2 \leq \sigma_i^2 \leq \sigma_{\max}^2$ for all $[\boldsymbol{\mu}, \boldsymbol{\sigma}] \in \mathcal{T}$. For all $P_{\theta_1}, P_{\theta_2} \in \mathcal{G}$,*

$$\|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\|_2^2 \leq D' \|p_{\theta_1} - p_{\theta_2}\|_{L^2(\mathbb{R}^d)}^2 \quad (\text{C.83})$$

where

$$D' := \frac{(4\pi\sigma_{\max}^2)^{d/2+1} D_1}{\pi(1-e^{-D_1})} \text{ with } D_1 = \frac{M^2}{\sigma_{\min}^2} + \frac{d}{2} \log \left(\frac{\sigma_{\max}^2}{\sigma_{\min}^2} \right)$$

¹⁰Note that the choice of τ is not optimal (indeed the minimum is attained for $\tau = \frac{2d}{2d+1}$), however we choose this value for the simplicity and clarity of the resulting bound.

Proof. We use a property from [3] on product of Gaussians:

$$\int \mathcal{N}(\mathbf{x}; \boldsymbol{\mu}_a, \boldsymbol{\Sigma}_a) \mathcal{N}(\mathbf{x}; \boldsymbol{\mu}_b, \boldsymbol{\Sigma}_b) d\mathbf{x} = \frac{1}{(2\pi)^{d/2} |\boldsymbol{\Sigma}_a + \boldsymbol{\Sigma}_b|^{1/2}} \exp\left(-\frac{1}{2}(\boldsymbol{\mu}_a - \boldsymbol{\mu}_b)^T (\boldsymbol{\Sigma}_a + \boldsymbol{\Sigma}_b)^{-1} (\boldsymbol{\mu}_a - \boldsymbol{\mu}_b)\right) \quad (\text{C.84})$$

Hence we have

$$\begin{aligned} \|p_{\boldsymbol{\theta}_1} - p_{\boldsymbol{\theta}_2}\|_{L^2(\mathbb{R}^d)}^2 &= \int (p_{\boldsymbol{\theta}_1}(\mathbf{x}) - p_{\boldsymbol{\theta}_2}(\mathbf{x}))^2 d\mathbf{x} \\ &= \int p_{\boldsymbol{\theta}_1}(\mathbf{x})^2 d\mathbf{x} + \int p_{\boldsymbol{\theta}_2}(\mathbf{x})^2 d\mathbf{x} - 2 \int p_{\boldsymbol{\theta}_1}(\mathbf{x}) p_{\boldsymbol{\theta}_2}(\mathbf{x}) d\mathbf{x} \\ &= \frac{1}{(2\pi)^{d/2}} \left[|2\boldsymbol{\Sigma}_1|^{-\frac{1}{2}} + |2\boldsymbol{\Sigma}_2|^{-\frac{1}{2}} - 2|\boldsymbol{\Sigma}_1 + \boldsymbol{\Sigma}_2|^{-\frac{1}{2}} e^{-\frac{1}{2}(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2)^T (\boldsymbol{\Sigma}_1 + \boldsymbol{\Sigma}_2)^{-1} (\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2)} \right] \\ &= \frac{|2\boldsymbol{\Sigma}_1|^{-\frac{1}{2}} + |2\boldsymbol{\Sigma}_2|^{-\frac{1}{2}}}{(2\pi)^{d/2}} \left[1 - e^{-\left(\frac{1}{2}(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2)^T (\boldsymbol{\Sigma}_1 + \boldsymbol{\Sigma}_2)^{-1} (\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2) + \log\left(\frac{|2\boldsymbol{\Sigma}_1|^{-\frac{1}{2}} + |2\boldsymbol{\Sigma}_2|^{-\frac{1}{2}}}{2|\boldsymbol{\Sigma}_1 + \boldsymbol{\Sigma}_2|^{-\frac{1}{2}}}\right)\right)} \right] \\ &\geq \frac{2}{(4\pi\sigma_{\max}^2)^{d/2}} \left[1 - e^{-\left(\frac{1}{2}(\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2)^T (\boldsymbol{\Sigma}_1 + \boldsymbol{\Sigma}_2)^{-1} (\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2) + \log\left(\frac{|2\boldsymbol{\Sigma}_1|^{-\frac{1}{2}} + |2\boldsymbol{\Sigma}_2|^{-\frac{1}{2}}}{2|\boldsymbol{\Sigma}_1 + \boldsymbol{\Sigma}_2|^{-\frac{1}{2}}}\right)\right)} \right] \quad (\text{C.85}) \end{aligned}$$

On the one hand, we have

$$0 \leq \frac{1}{4\sigma_{\max}^2} \|\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2\|_2^2 \leq \frac{1}{2} (\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2)^T (\boldsymbol{\Sigma}_1 + \boldsymbol{\Sigma}_2)^{-1} (\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2) \leq \frac{1}{2} \frac{4M^2}{2\sigma_{\min}^2} = \frac{M^2}{\sigma_{\min}^2} \quad (\text{C.86})$$

On the other hand,

$$\begin{aligned} \log\left(\frac{|2\boldsymbol{\Sigma}_1|^{-\frac{1}{2}} + |2\boldsymbol{\Sigma}_2|^{-\frac{1}{2}}}{2|\boldsymbol{\Sigma}_1 + \boldsymbol{\Sigma}_2|^{-\frac{1}{2}}}\right) &= \log\left(\frac{|2\boldsymbol{\Sigma}_1|^{-\frac{1}{2}} + |2\boldsymbol{\Sigma}_2|^{-\frac{1}{2}}}{2}\right) + \frac{1}{2} \log|\boldsymbol{\Sigma}_1 + \boldsymbol{\Sigma}_2| \\ &\geq \frac{1}{2} \log|2\boldsymbol{\Sigma}_1|^{-\frac{1}{2}} + \frac{1}{2} \log|2\boldsymbol{\Sigma}_2|^{-\frac{1}{2}} + \frac{1}{2} \log|\boldsymbol{\Sigma}_1 + \boldsymbol{\Sigma}_2| \text{ by concavity of the log} \\ &= \frac{1}{4} \left(\sum_{\ell=1}^d \log(\sigma_{1,\ell}^2 + \sigma_{2,\ell}^2)^2 - \sum_{\ell=1}^d \log(2\sigma_{1,\ell}^2) - \sum_{\ell=1}^d \log(2\sigma_{2,\ell}^2) \right) \\ &= \frac{1}{4} \sum_{\ell=1}^d \log\left(\frac{(\sigma_{1,\ell}^2 + \sigma_{2,\ell}^2)^2}{4\sigma_{1,\ell}^2 \sigma_{2,\ell}^2}\right) \\ &\geq \frac{1}{4} \sum_{\ell=1}^d \left(1 - \frac{4\sigma_{1,\ell}^2 \sigma_{2,\ell}^2}{(\sigma_{1,\ell}^2 + \sigma_{2,\ell}^2)^2}\right) \quad \text{since } -\log(x) \geq 1 - x \\ &= \sum_{\ell=1}^d \frac{(\sigma_{1,\ell}^2 - \sigma_{2,\ell}^2)^2}{4(\sigma_{1,\ell}^2 + \sigma_{2,\ell}^2)} \geq \frac{1}{8\sigma_{\max}^2} \|\boldsymbol{\sigma}_1 - \boldsymbol{\sigma}_2\|_2^2 \end{aligned}$$

and therefore

$$0 \leq \frac{1}{8\sigma_{\max}^2} \|\boldsymbol{\sigma}_1 - \boldsymbol{\sigma}_2\|_2^2 \leq \log\left(\frac{|2\boldsymbol{\Sigma}_1|^{-\frac{1}{2}} + |2\boldsymbol{\Sigma}_2|^{-\frac{1}{2}}}{2|\boldsymbol{\Sigma}_1 + \boldsymbol{\Sigma}_2|^{-\frac{1}{2}}}\right) \leq \log\left(\left(\frac{2\sigma_{\max}^2}{2\sigma_{\min}^2}\right)^{d/2}\right) = \frac{d}{2} \log\left(\frac{\sigma_{\max}^2}{\sigma_{\min}^2}\right). \quad (\text{C.87})$$

We can now bound

$$0 \leq \frac{1}{8\sigma_{\max}^2} \|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\|_2^2 \leq \left[\frac{1}{2} (\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2)^T (\boldsymbol{\Sigma}_1 + \boldsymbol{\Sigma}_2)^{-1} (\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2) + \log\left(\frac{|2\boldsymbol{\Sigma}_1|^{-\frac{1}{2}} + |2\boldsymbol{\Sigma}_2|^{-\frac{1}{2}}}{2|\boldsymbol{\Sigma}_1 + \boldsymbol{\Sigma}_2|^{-\frac{1}{2}}}\right) \right] \leq D_1, \quad (\text{C.88})$$

where

$$D_1 = \frac{M^2}{\sigma_{\min}^2} + \frac{d}{2} \log\left(\frac{\sigma_{\max}^2}{\sigma_{\min}^2}\right).$$

By concavity of $x \mapsto 1 - e^{-x}$, the function $\varphi : x \mapsto \frac{1 - e^{-x}}{x}$ is decreasing. Hence we have:

$$\forall x \in [0; D_1], \quad 1 - e^{-x} \geq \frac{1 - e^{-D_1}}{D_1} x.$$

Therefore, given (C.85) and (C.88), we have

$$\|p_{\boldsymbol{\theta}_1} - p_{\boldsymbol{\theta}_2}\|_{L^2(\mathbb{R}^d)}^2 \geq D_2 \left(\frac{1}{2} (\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2)^T (\boldsymbol{\Sigma}_1 + \boldsymbol{\Sigma}_2)^{-1} (\boldsymbol{\mu}_1 - \boldsymbol{\mu}_2) + \log \left(\frac{|2\boldsymbol{\Sigma}_1|^{-\frac{1}{2}} + |2\boldsymbol{\Sigma}_2|^{-\frac{1}{2}}}{2|\boldsymbol{\Sigma}_1 + \boldsymbol{\Sigma}_2|^{-\frac{1}{2}}} \right) \right),$$

where $D_2 := \frac{2(1 - e^{-D_1})}{D_1(4\pi\sigma_{\max}^2)^{d/2}}$.

And we have, using (C.88) again:

$$\|p_{\boldsymbol{\theta}_1} - p_{\boldsymbol{\theta}_2}\|_{L^2(\mathbb{R}^d)}^2 \geq \frac{D_2}{8\sigma_{\max}^2} \|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\|_2^2,$$

which leads to the desired result. \square

Proof of Theorem 5. We denote \mathcal{F} and \mathcal{F}^{-1} the Fourier and inverse Fourier transform:

$$\begin{aligned} \mathcal{F}(f)(\boldsymbol{\omega}) &= \int_{\mathbb{R}^d} e^{-i\boldsymbol{\omega}^T \mathbf{x}} f(\mathbf{x}) d\mathbf{x}, \\ \mathcal{F}^{-1}(F)(\mathbf{x}) &= \frac{1}{(2\pi)^d} \int_{\mathbb{R}^d} e^{i\boldsymbol{\omega}^T \mathbf{x}} F(\boldsymbol{\omega}) d\boldsymbol{\omega}. \end{aligned}$$

We recall the classical Plancherel's Theorem:

$$\|f\|_2^2 = \frac{1}{(2\pi)^d} \|\mathcal{F}(f)\|_2^2. \quad (\text{C.89})$$

Let $P_1, P_2 \in \mathcal{G}$. Recall that $\Lambda = \mathcal{N}(0, \frac{a}{n} \mathbf{I})$, denote $\sigma_\Lambda^2 := \frac{a}{d}$. The MMD is expressed:

$$\begin{aligned} \gamma_\Lambda^2(P_1, P_2) &= \int |\psi_1(\boldsymbol{\omega}) - \psi_2(\boldsymbol{\omega})|^2 \mathcal{N}(\boldsymbol{\omega}; 0, \sigma_\Lambda^2 \mathbf{I}) d\boldsymbol{\omega} \\ &= \int |\psi_1(\boldsymbol{\omega}) - \psi_2(\boldsymbol{\omega})|^2 \frac{1}{(2\pi\sigma_\Lambda^2)^{d/2}} e^{-\frac{\|\boldsymbol{\omega}\|_2^2}{2\sigma_\Lambda^2}} d\boldsymbol{\omega} \\ &= \frac{1}{(2\pi\sigma_\Lambda^2)^{d/2}} \int \left| e^{-i\boldsymbol{\omega}^T \boldsymbol{\mu}_1} e^{-\frac{1}{2}\boldsymbol{\omega}^T \left(\boldsymbol{\Sigma}_1 + \frac{\mathbf{I}}{2\sigma_\Lambda^2}\right) \boldsymbol{\omega}} - e^{-i\boldsymbol{\omega}^T \boldsymbol{\mu}_2} e^{-\frac{1}{2}\boldsymbol{\omega}^T \left(\boldsymbol{\Sigma}_2 + \frac{\mathbf{I}}{2\sigma_\Lambda^2}\right) \boldsymbol{\omega}} \right|^2 d\boldsymbol{\omega} \\ &= \frac{1}{(2\pi\sigma_\Lambda^2)^{d/2}} \int |\psi_{P'_1}(\boldsymbol{\omega}) - \psi_{P'_2}(\boldsymbol{\omega})|^2 d\boldsymbol{\omega}, \end{aligned}$$

where P'_i is a Gaussian with the same mean than P_i and dilated variance $\boldsymbol{\Sigma}'_i = \boldsymbol{\Sigma}_i + \frac{\mathbf{I}}{2\sigma_\Lambda^2}$. Since $\psi_{P'_i} = \mathcal{F}(p'_i)$, by Plancherel's Theorem we have:

$$\gamma_\Lambda^2(P_1, P_2) = \left(\frac{2\pi}{\sigma_\Lambda^2} \right)^{d/2} \int |p'_1(\mathbf{x}) - p'_2(\mathbf{x})|^2 d\mathbf{x} = \left(\frac{2\pi}{\sigma_\Lambda^2} \right)^{d/2} \|p'_1 - p'_2\|_2^2.$$

The parameters of the Gaussians P'_i belong to a compact set $\mathcal{T}' = \left\{ \left(\boldsymbol{\mu}, \boldsymbol{\Sigma} + \frac{\mathbf{I}}{2\sigma_\Lambda^2} \right); (\boldsymbol{\mu}, \boldsymbol{\Sigma}) \in \mathcal{T} \right\}$. We can therefore apply Lemma C.10, such that:

$$D' \|p'_1 - p'_2\|_2^2 \geq \|\boldsymbol{\theta}'_1 - \boldsymbol{\theta}'_2\|_2^2 = \|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\|_2^2. \quad (\text{C.90})$$

The last equality comes from the fact that the variance between $\boldsymbol{\theta}_i$ and $\boldsymbol{\theta}'_i$ are just translated. The constant D' is:

$$D' = \frac{(4\pi(\sigma_{\max}^2 + 1/(2\sigma_\Lambda^2)))^{d/2+1} D_1}{\pi(1 - e^{-D_1})} \quad \text{with} \quad D_1 = \frac{M^2}{\sigma_{\min}^2 + 1/(2\sigma_\Lambda^2)} + \frac{d}{2} \log \left(\frac{\sigma_{\max}^2 + 1/(2\sigma_\Lambda^2)}{\sigma_{\min}^2 + 1/(2\sigma_\Lambda^2)} \right)$$

and thus

$$\left(\frac{\sigma_\Lambda^2}{2\pi}\right)^{d/4} D' \gamma_\Lambda(P_1, P_2) \geq \|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\|_2.$$

Considering eq. (C.81) in the proof of Theorem 4, we have

$$\|P_1 - P_2\|_{TV} \leq \max(\sigma_{\min}^{-1}, \sigma_{\min}^{-2}/\sqrt{2}) \|\boldsymbol{\theta}_1 - \boldsymbol{\theta}_2\|_2 \leq \bar{D} \gamma_\Lambda(P_1, P_2), \quad (\text{C.91})$$

with $\bar{D} := \max(\sigma_{\min}^{-1}, \sigma_{\min}^{-2}/\sqrt{2}) \left(\frac{\sigma_\Lambda^2}{2\pi}\right)^{d/4} D' = \max(\sigma_{\min}^{-1}, \sigma_{\min}^{-2}/\sqrt{2}) \sqrt{\frac{2D_1 \cdot (2\sigma_\Lambda^2 \sigma_{\max}^2 + 1)^{d/2+1}}{\sigma_\Lambda^2 (1 - e^{-D_1})}}$.

We now use the fact that $\sigma_\Lambda^2 = \frac{a}{d}$. We have

$$\begin{aligned} D_1 &= \frac{M^2}{\sigma_{\min}^2 + d/(2a)} + \frac{d}{2} \log \left(\frac{2a\sigma_{\max}^2/d + 1}{2a\sigma_{\min}^2/d + 1} \right) \\ &\leq \frac{M^2}{\sigma_{\min}^2 + d/(2a)} + \frac{d}{2} \log(2a\sigma_{\max}^2/d + 1) \\ &\leq \frac{M^2}{\sigma_{\min}^2 + d/(2a)} + \frac{d}{2} 2a\sigma_{\max}^2/d = \frac{M^2}{\sigma_{\min}^2 + d/(2a)} + a\sigma_{\max}^2 \quad \text{since } \log(1+x) \leq x \\ &\leq \sigma_{\max}^2 a \left(1 + \frac{2M^2}{d}\right) := D_2 \end{aligned} \quad (\text{C.92})$$

and similarly

$$\begin{aligned} (2\sigma_\Lambda^2 \sigma_{\max}^2 + 1)^{d/2+1} &= (2a\sigma_{\max}^2/d + 1) \exp \left(\frac{d}{2} \log(2a\sigma_{\max}^2/d + 1) \right) \\ &\leq (2a\sigma_{\max}^2/d + 1) e^{a\sigma_{\max}^2} \\ &\leq e^{a\sigma_{\max}^2(1+\frac{1}{d})} \leq e^{3a\sigma_{\max}^2}. \end{aligned} \quad (\text{C.93})$$

Using (C.92) with the fact that the function $\varphi : x \mapsto \frac{1-e^{-x}}{x}$ is decreasing, and (C.93), we obtain

$$\|P_1 - P_2\|_{TV} \leq D \gamma_\Lambda(P_1, P_2), \quad (\text{C.94})$$

with $D := \max(\sigma_{\min}^{-1}, \sigma_{\min}^{-2}/\sqrt{2}) \sqrt{\frac{2dD_2 \cdot e^{3a\sigma_{\max}^2}}{a(1 - e^{-D_2})}}$. □