Sketching Data Sets for Large-Scale Learning: Keeping only what you need

Rémi Gribonval, Antoine Chatalic, Nicolas Keriven, Vincent Schellekens,
Laurent Jacques, Philip Schniter

► To cite this version:

Rémi Gribonval, Antoine Chatalic, Nicolas Keriven, Vincent Schellekens, Laurent Jacques, et al.. Sketching Data Sets for Large-Scale Learning: Keeping only what you need. IEEE Signal Processing Magazine, 2021, 38 (5), pp.12-36. 10.1109/MSP.2021.3092574 . hal-03350599

HAL Id: hal-03350599

<https://hal.inria.fr/hal-03350599>

Submitted on 21 Sep 2021

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 Datasets for Large-Scale Learning –Keeping Only What You Need–

Rémi Gribonval, Antoine Chatalic, Nicolas Keriven,
Vincent Schellekens, Laurent Jacques, and Philip Schniter

Big data can be a blessing: with very large training datasets it becomes possible to perform complex learning tasks with unprecedented accuracy. Yet, this improved performance comes at the price of enormous computational challenges. Thus, one may wonder: Is it possible to leverage the information content of huge datasets while keeping computational resources under control? Can this also help solve some of the privacy issues raised by large-scale learning? This is the ambition of *compressive learning*, where the dataset is massively compressed *before* learning. Here, a “sketch” is first constructed by computing carefully chosen nonlinear random features (*e.g.*, random Fourier features) and averaging them over the whole dataset. Parameters are then learned from the sketch, without access to the original dataset. This article surveys the current state-of-the-art in compressive learning, including the main concepts and algorithms; their connections with established signal-processing methods; existing theoretical guarantees, on both information preservation and privacy preservation; and important open problems. For an extended version of this article that contains additional references and more in-depth discussions on a variety of topics, see [1].

1 INTRODUCTION TO COMPRESSIVE LEARNING

The overall principle of compressive learning is summarized in Fig. 1. During the *sketching phase*, a potentially huge collection of d -dimensional data vectors $\{\mathbf{x}_i\}_{i=1}^n$ is summarized into a single m -dimensional vector $\tilde{\mathbf{z}}$, called the “sketch”. The sketch is constructed by transforming each data vector and then averaging the results:

$$\tilde{\mathbf{z}} := \frac{1}{n} \sum_{i=1}^n \Phi(\mathbf{x}_i). \quad (1)$$

Next, during the *learning phase*, an estimate $\tilde{\boldsymbol{\theta}}$ of some essential statistical parameters $\boldsymbol{\theta}$ of the dataset are extracted from the sketch $\tilde{\mathbf{z}}$. These parameters are application dependent. For example, as illustrated in Fig. 2, they could represent principal data subspaces (in subspace fitting problems), prediction weights (in estimation and regression problems), distributional parameters (in density estimation problems), or centroids (in clustering problems); we give detailed examples below. Python notebooks that illustrate examples from the paper are available at <https://gitlab.inria.fr/SketchedLearning/spm-notebook>.

The transformation $\Phi(\cdot)$, known as the “feature map,” is generally non-linear and randomized. Although the use of $\Phi(\cdot)$ is related to “kernel methods” in machine learning (*e.g.*, [2, 3]), as will be discussed later, the act of sketching (1) also includes averaging over the n data samples. The advantages of compressive learning are the following.

- 1) By choosing a sketch of dimension $m \ll nd$, the data gets massively compressed. This has obvious advantages for storage and transfer.
- 2) Sketching can speed up the learning phase, whose complexity becomes independent of the cardinality, n , of the original dataset. This enables one to handle massive datasets while keeping computational resources under control.

R. Gribonval is with Univ Lyon, Inria, CNRS, ENS de Lyon, UCB Lyon 1, LIP UMR 5668, F-69342, Lyon, France; remi.gribonval@inria.fr. A. Chatalic is with Univ Rennes, Inria, CNRS, IRISA, F-35042 Rennes Cedex, France; antoine.chatalic@irisa.fr. N. Keriven is with CNRS, GIPSA-lab, Grenoble INP, UGA, Grenoble, France; nicolas.keriven@gipsa-lab.grenoble-inp.fr. V. Schellekens and L. Jacques are with ISPGROUP/ICTEAM, UCLouvain, Belgium; {vincent.schellekens,laurent.jacques}@uclouvain.be. P. Schniter is with The Ohio State University, Columbus, OH, USA; schniter.1@osu.edu

Fig. 1: Overview of sketching and parameter learning.

Fig. 2: Schematic representation of the four running examples covered by this paper. In the four figures, each x_i is associated with a blue colored disk, hence the training collection correspond to a point cloud, and the orange color geometrically represents the learned parameters. PCA learns the principal k -dimensional subspace of the dataset for some $k \leq d$. Linear regression fits observed data (the blue dot heights) as a linear model of the inputs (here, the 2-D horizontal coordinates). For least squares, this amounts to minimizing the square of the differences between these data and the linear predictions. In Gaussian mixture modeling, we learn the set of parameters (mixture weight, mean, and covariance) characterizing each Gaussian term of the mixture, which probability level sets are displayed here as orange ellipses. Clustering methods (such as k -means) learn a set of centroids (the orange squares) defining a Voronoi partition grouping together similar data samples.

- 3) Sketching can preserve privacy: the transformation $\Phi(\cdot)$ can be chosen so that individual-user information is lost while aggregate-user information is preserved.
- 4) The sketching mechanism in (1) is well matched to distributed implementations and streaming scenarios: the sketch of a concatenation of datasets is a simple mean of the sketches of those datasets.

2 ILLUSTRATION USING FOUR WORKED EXAMPLES

To illustrate the compressive-learning framework and discuss the essential aspects of it, we now outline four canonical examples of machine learning tasks to which sketching can be readily applied. See Fig. 2 for an illustration.

a) Principal component analysis (PCA)

PCA seeks to find the linear subspace of a fixed dimension $k < d$ that best fits the d -dimensional data $\{x_i\}_{i=1}^n$ in the least-squares (LS) sense. In this case, we assume the data to be centered so that the target parameters θ can be described by a k -dimensional orthonormal basis $\{u_\ell\}_{\ell=1}^k$ that maximizes $\sum_{\ell=1}^k \sum_{i=1}^n |u_\ell^\top x_i|^2$.

It is well known that one solution is given by the k principal eigenvectors of the empirical autocorrelation matrix $\hat{R} = \frac{1}{n} \sum_{i=1}^n x_i x_i^\top$. This $\hat{R}$ can be interpreted as a sketch of the form (1) that uses the feature map $\Phi(x) = \text{vec}(xx^\top) = (x_1 x_1, x_2 x_1, \dots, x_d x_d)^\top$ of dimension $m = d^2$. Here and in the sequel, the $\text{vec}(\cdot)$ operator vectorizes a matrix by stacking its columns. As soon as the cardinality n of the dataset

exceeds the dimension d , the dimension of this sketch is smaller than nd , the total size of the dataset. Using techniques from matrix completion and compressive matrix sensing, the sketch dimension can be further reduced [4]. For example, one can sketch via (1) using $\Phi(\mathbf{x}) = ((\mathbf{w}_1^\top \mathbf{x})^2, \dots, (\mathbf{w}_m^\top \mathbf{x})^2)^\top$, where the d -dimensional vectors $\mathbf{w}_j$, $1 \leq j \leq m$ are the rows of a tall random matrix $\mathbf{W}$ of size $m \times d$. For m on the order of $kd < d^2$, low-rank matrix reconstruction techniques can estimate the k -dimensional principal subspace of $\hat{\mathbf{R}}$ with provable accuracy [5, 6].

b) LS linear regression

Suppose that the data vectors take the form $\mathbf{x}_i = [\mathbf{x}_{1i}^\top, \mathbf{x}_{2i}^\top]^\top$, and our goal is to linearly predict the d_1 -dimensional vector $\mathbf{x}_{1i}$ from the d_2 -dimensional vector $\mathbf{x}_{2i}$ (with $d = d_1 + d_2$). That is, we want to design a $d_1 \times d_2$ weight matrix $\boldsymbol{\theta}$ such that $\mathbf{x}_{1i} \approx \boldsymbol{\theta} \mathbf{x}_{2i}$ for all samples i . The LS approach to this supervised learning problem chooses $\hat{\boldsymbol{\theta}} = \arg \min_{\boldsymbol{\theta}} \sum_{i=1}^n \|\mathbf{x}_{1i} - \boldsymbol{\theta} \mathbf{x}_{2i}\|^2$.

Although it is possible to compute the LS solution using gradient descent, each iteration would involve the full dataset $\{\mathbf{x}_i\}_{i=1}^n$. A well-known alternative is to first build the empirical auto-correlation matrix $\hat{\mathbf{R}} := \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i^\top$ and then compute $\hat{\boldsymbol{\theta}}$ in closed-form as

$$\hat{\boldsymbol{\theta}} = \hat{\mathbf{R}}_{12} \hat{\mathbf{R}}_{22}^{-1}, \text{ with } \begin{pmatrix} \hat{\mathbf{R}}_{11} & \hat{\mathbf{R}}_{12} \\ \hat{\mathbf{R}}_{21} & \hat{\mathbf{R}}_{22} \end{pmatrix} = \hat{\mathbf{R}}, \quad (2)$$

where each sub-matrix $\hat{\mathbf{R}}_{ij}$ has dimension $d_i \times d_j$. Here we again use the sketch (1) with the feature map $\Phi(\mathbf{x}) = \text{vec}(\mathbf{x} \mathbf{x}^\top)$ and extract the target parameter $\hat{\boldsymbol{\theta}}$ from (2).

c) Gaussian-mixture modeling

Here the objective is to find the parameters $\boldsymbol{\theta} = \{\alpha_\ell, \boldsymbol{\mu}_\ell, \boldsymbol{\Sigma}_\ell\}_{\ell=1}^k \subset \boldsymbol{\Theta}$ that best fit a k -term Gaussian-mixture model $p(\mathbf{x}|\boldsymbol{\theta}) = \sum_{\ell=1}^k \alpha_\ell \mathcal{N}(\mathbf{x}; \boldsymbol{\mu}_\ell, \boldsymbol{\Sigma}_\ell)$ to the data $\{\mathbf{x}_i\}_{i=1}^n$. The parameter space $\boldsymbol{\Theta}$ demands that, for each ℓ : $\alpha_\ell > 0$, $\boldsymbol{\mu}_\ell$ is a d -dimensional centroid, $\boldsymbol{\Sigma}_\ell$ is a $d \times d$ positive definite matrix, and $\sum_\ell \alpha_\ell = 1$. Traditionally, expectation-maximization (EM) is used to approximate the maximum likelihood (ML) estimate $\hat{\boldsymbol{\theta}} = \arg \max_{\boldsymbol{\theta} \in \boldsymbol{\Theta}} \sum_{i=1}^n \log p(\mathbf{x}_i|\boldsymbol{\theta})$, but the EM algorithm processes all n data samples $\{\mathbf{x}_i\}_{i=1}^n$ at each iteration, which can be computationally burdensome when n is very large.

[Box 1] Compressive Mixture Modeling for Speaker Verification

Fig. 3: Speaker verification via compressive learning. Unlabeled training speech data is collected in a decentralized manner, pre-processed, and locally sketched. The local sketches are then merged into a global sketch, from which a Universal Background Model is learned.

Given a fragment of speech and a candidate speaker, the goal of speaker verification is to assess whether the fragment was indeed spoken by that person. A classic approach to speaker verification is the GMM-UBM [7] (Gaussian Mixture Model–Universal Background Model). There, the idea is to train a model of a “universal” speaker from unlabeled training data and then compare it to a specialized model for the candidate speaker. Using the speech fragment, the likelihood ratio between the specialized and universal models is computed, and a positive decision is made if the ratio exceeds a certain threshold. In GMM-UBM, the data is modeled using a GMM. That is, $p(\mathbf{x}|\boldsymbol{\theta})$ are multivariate Gaussian distributions applied to a suitable time-frequency transform of the raw audio signal [7] (see Fig. 3).

The training of the UBM is computationally demanding, since it must be done on a large corpus of speech data. Moreover, the latter must be collected in a wide range of situations, so that the final model may be as universal as possible. For this reason, the data collecting process is best performed in a decentralized manner. Finally, speech data collected in real-life situations is known to be sensitive information. For all of these reasons, compressive learning is well suited to speaker verification.

In [8], using an RF feature map $\Phi(\cdot)$, the authors compressed 1000 hours of speech data (50 gigabytes) into a sketch of a few kilobytes on a single laptop. Subsequently, they performed GMM estimation, i.e., learning, using a greedy algorithm, thereby tackling the optimization problem (18) introduced later. They compared this compressive-learning approach to the EM algorithm, which, on the same laptop, could only be trained using 5 hours of speech given the available RAM. They observed that, by enabling the use of more data within a fixed memory budget, sketching produced better results, despite the tremendous compression factor (see Fig. 4).

Fig. 4: Detection Error Tradeoff (DET) curve (False-Positive/False-Negative tradeoff obtained by varying the decision threshold) and Equal Error Rate (EER). On a single laptop: EM was trained on 5h of speech, compressive GMM estimation used 5h in the left figure and 1000h in the right, for different sketch dimensions m .

An alternative [8] is to compute a sketch $\tilde{z}$ of the form (1) using random Fourier (RF) features [9]:

$$\Phi(\mathbf{x}) = \exp(-j2\pi\mathbf{W}\mathbf{x}), \quad (3)$$

where $j = \sqrt{-1}$ and $\mathbf{W} \in \mathbb{R}^{m \times d}$ is a realization of a random matrix (*e.g.*, i.i.d. Gaussian). Here, $\exp(\cdot)$ is the componentwise exponential, not the matrix exponential. We will have a lot to say about the RF feature map later in this article.¹ The parameters can then be extracted by optimizing a cost function as explained later. This approach has been applied to audio source-separation as well as speaker verification [8], where it was shown that 1000 hours of speech can be compressed down to a few kilobytes without loss of verification performance (see Box 1).

d) *k*-means clustering

The goal of clustering is to group together “similar” data samples from $\{\mathbf{x}_i\}_{i=1}^n$. In the *k*-means approach to clustering, one aims to find the set of k centroids $\{\mathbf{c}_\ell\}_{\ell=1}^k$ that minimizes the average squared distance from each sample to its nearest centroid, *i.e.*, $\sum_{i=1}^n \min_{\ell} \|\mathbf{x}_i - \mathbf{c}_\ell\|^2$. The famous Lloyd algorithm[10] is typically used in an attempt to solve this problem. When n is very large, however, Lloyd’s algorithm becomes computationally demanding. Instead, one could sketch the dataset using (1) and extract the centroids from the sketch [11, 12]. For this purpose, one could use RF features (3) and (as explained in the sequel) solve for the centroids $\{\mathbf{c}_\ell\}_{\ell=1}^k$ and the (non-negative, sum-to-one) weights $\{\alpha_\ell\}_{\ell=1}^k$ that minimize $\|\tilde{z} - \sum_{\ell=1}^k \alpha_\ell \Phi(\mathbf{c}_\ell)\|$. In this setting, the weights allow us to model unbalanced clusters, yet only the centroids need to be recovered. On large datasets, this approach can be orders-of-magnitude better than Lloyd’s algorithm in memory and runtime, provided that the sketch dimension m is large enough, *i.e.*, that m is on the order of kd , where kd is the number of free parameters in $\{\mathbf{c}_\ell\}_{\ell=1}^k$ [11, 12]. For example, this method allows us to cluster the MNIST digit dataset, of dimension $d = 784$ and cardinality $n = 70\,000$, using a complex-valued sketch of dimension $m = 400$ (see Box 2).

¹For example, the connection between the RF map and the Gaussian kernel is discussed in Box 8.

[Box 2] Compressive Clustering of MNIST digits

Fig. 5: Clustering of handwritten digits via compressive k -means. First, SIFT descriptors are extracted from each image. Then a similarity graph is constructed to obtain a so-called *spectral embedding* of the dataset (using the first eigenvectors of the Laplacian of the graph). The spectral features are aggregated into a sketch, from which centroids are extracted.

Clustering is a classic machine-learning task and a component of many machine-learning pipelines. The most popular method is Lloyd’s algorithm [10], which aims to solve the k -means problem. In many cases, the raw features are converted to a spectral embedding before clustering. As a proof of concept, the authors in [11] applied this technique to handwritten digits from the MNIST dataset, but used compressive k -means clustering in place of Lloyd’s algorithm (see Fig. 5). Using a $n = 10^7$ -sample augmentation of MNIST, they found that compressive k -means clustering gave approximately the same accuracy as Lloyd’s algorithm but reduced the time- and memory-complexity by 1.5 and 4 orders-of-magnitude, respectively (see Fig. 6).

Fig. 6: Memory, runtime, and error of compressive k -means clustering relative to Lloyd’s algorithm for various sample cardinalities n and sketch lengths m , using $k = 10$ centroids and $d = 10$ dimensional spectral features. Figure from [11].

Although this paper focuses on these examples, the general compressive-learning framework extends beyond them. It has been applied to, for example, Independent Component Analysis [13] (by sketching the cumulant tensor) and subspace clustering [14] (using a polynomial feature map). Extending the framework to a new task raises essential questions such as: How do we choose the feature map $\Phi(\cdot)$? How do we learn the essential parameters θ from the sketch $\tilde{z}$? Are there statistical learning guarantees? Can sketching and learning be made computationally efficient? Can we sketch while respecting privacy? This paper sketches answers to these questions for the worked examples introduced above.

Fig. 7: Left: A streaming scenario, where the data samples are sketched one-by-one, and the mean sketch is updated at each time. Right: A distributed scenario, where each device computes a local sketch, and a centralized entity further averages these local sketches.

[Box 3] Notations and Conventions

Notation	Description
$\mathcal{X} = \{x_i\}_{i=1}^n$	dataset of n training samples $x_i \in \mathbb{R}^d$
$\tilde{z}$	(empirical) sketch of the dataset, a vector of dimension m
$\Phi(\cdot)$	sketching feature map, a function mapping $\mathbb{R}^d$ to either $\mathbb{R}^m$ or $\mathbb{C}^m$
$\theta \in \Theta$	target parameters to learn, of dimension p (e.g., PCA matrix, centroids, mixture model)
W	random matrix associated with certain feature maps $\Phi(\cdot)$, usually drawn i.i.d. Gaussian
$\varrho(\cdot)$	componentwise nonlinearity associated with certain feature maps $\Phi(\cdot)$
w.h.p.	with high probability, i.e., with exponentially decaying failure probability
i.i.d.	independent and identically distributed
w.p.1	with probability 1
$\ \cdot\ , \langle \cdot, \cdot \rangle$	Euclidean norm of a vector, inner product between vectors
$\ \cdot\ _0$	ℓ^0 norm of a vector, i.e., its number of nonzero entries
$\ \cdot\ _F$	Frobenius norm of a matrix
$\ \cdot\ _{L^2}, \langle \cdot, \cdot \rangle_{L^2}$	L^2 norm of a function, L^2 inner product between functions

3 HISTORICAL BACKGROUND

The term “sketch” has different meanings, depending on the field. Our use of the term comes from the literature on relational databases, and more particularly from the subfield of *Approximate Query Processing* (AQP) [15]. The goal of AQP is to build a short description of the content of a massive dataset, called a *synopsis*, by analogy with the synopsis of a movie or a book, such that certain *queries* can be efficiently performed to return answers with controlled error and/or probability of failure. Well-known queries include the frequency of occurrence of a particular element in a stream of data (taken from a discrete collection), and the minimum of several of these frequencies, yielding the celebrated count-min sketch [15, Section 5.3.1] synopsis. In this context, the statistical parameters θ learned from a sketch are interpreted as the result of a particular query on the dataset.

Despite the apparent similarities between AQP and compressive learning, both datatypes and learning tasks differ significantly between the two fields. Typically, AQP focuses on (multi)sets of elements taken from a discrete collection of objects and considers database queries and related operations, while compressive learning focuses on continuous-valued signals (e.g., images or audio signals) and considers machine-learning tasks such as density estimation or regression.

3.1 SKETCHES FOR STREAMING AND DISTRIBUTED METHODS

In AQP, a popular class of synopses is that of *linear sketches* [15, Chap. 5]. A linear sketch, which maps a dataset to a vector, must satisfy the single condition that *the sketch of the concatenation of two datasets*

equals the sum of their sketches. In mathematical terms, if we denote by $\tilde{z}(\mathcal{X})$ the sketch of a dataset $\mathcal{X} = \{\mathbf{x}_i\}_{i=1}^n$, then it is required that $\tilde{z}(\mathcal{X}) = \tilde{z}(\mathcal{X}_1) + \tilde{z}(\mathcal{X}_2)$ whenever $\mathcal{X}$ is the concatenation of $\mathcal{X}_1$ and $\mathcal{X}_2$. It thus follows that a linear sketch must be of the form $\tilde{z}(\mathcal{X}) = \sum_{i=1}^n \Phi(\mathbf{x}_i)$ for some possibly nonlinear feature map $\Phi(\cdot)$. That is the same definition that we adopt in (1), except that we normalize by the number of elements n . Linear sketches are popular in AQP mainly because they are well suited to *streaming* scenarios. That is, inserting or deleting an element $\mathbf{x}_i$ from the dataset corresponds to adding or subtracting $\Phi(\mathbf{x}_i)$ from the sketch, respectively (see Fig. 7). Note that some very simple sketches are not linear: for instance, it is easy to sketch the maximum running value in a stream of scalar data by computing the maximum of the current sketch and each new data point, but this sketching procedure is not linear [15, Chap. 5.2.1]. In particular, this “max” sketch facilitates the insertion of a new element in the database, but not the deletion of an existing one.

Remark. *In the subsections that follow, we describe linear sketches from other points of view. Before doing that, however, we want to clarify that the terms “sketching” and “linear sketching” appear in many other fields, although with meanings that differ considerably from ours. For instance, sketching is often used to indicate **linear** dimensionality reduction in n or d , as attained by multiplying the data matrix $[\mathbf{x}_1, \dots, \mathbf{x}_n]^\top \in \mathbb{R}^{n \times d}$ by a random matrix on the left or the right [16–19]. Sketching may also refer to the use of sampling-based approaches [20], such as coresets [21] or the Nyström method [22]. These latter methods differ from linear sketches in AQP, which is our notion of sketching, in that these latter methods generally do not respect the concatenation condition described earlier and are therefore less directly amenable to streaming scenarios. Moreover, these methods typically do not involve a **nonlinear** feature map $\Phi(\cdot)$, which is a key component of our sketch.*

3.2 SKETCHES AS (RANDOMIZED) GENERALIZED MOMENTS

In signal processing and machine learning, the data samples $\mathbf{x}_i$ generally live in the vector space $\mathbb{R}^d$ and are often modeled as independent and identically distributed (i.i.d.) random vectors having a *probability distribution* with density p_X . Consider what happens when the number of samples n goes to infinity. The strong law of large numbers says that

$$\lim_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \Phi(\mathbf{x}_i) \stackrel{\text{w.p.1}}{=} \mathbb{E}[\Phi(X)] = \int p_X(\mathbf{x}) \Phi(\mathbf{x}) d\mathbf{x}, \quad (4)$$

where $\mathbb{E}[\cdot]$ denotes expectation with respect to the probability density p_X . If we consider the simple case of dimension $d = 1$ and the scalar transformation $\Phi(x) = x^k$ (so that $m = 1$), then $\mathbb{E}[\Phi(X)]$ is the (uncentered) k -th *moment* of the random variable X , a quantity that has a long history in statistics. By analogy, with a generic vector-valued feature map $\Phi(\cdot)$ and in dimension $d > 1$, quantities of the form $\mathbb{E}[\Phi(X)]$ are known as *generalized moments* of the random vector $X \in \mathbb{R}^d$.

Performing inference from generalized moments is often referred to as the *Generalized Method of Moments* (GeMM) [23]. This method to “learn from a sketch” is very popular in, *e.g.*, the field of econometrics [23, Chap. 1]. The GeMM can be seen as an alternative to maximum likelihood estimation that avoids the need to work with the full likelihood function, which can have computational benefits. Indeed, for many classes of probability distributions, such as heavy-tailed α -stable distributions, the likelihood function is not given in closed form, but generalized moments are given in closed form for appropriately chosen feature maps $\Phi(\cdot)$ (see Box 1).

GeMM differs from compressive learning in several aspects. In GeMM, the feature map $\Phi(\cdot)$ is typically constructed to make the parameter estimates computable in closed form. This narrows the range of learning tasks that GeMM can handle. In compressive learning, $\Phi(\cdot)$ is designed with information-preservation in mind. Consequently, the range of learning tasks is much broader in compressive learning, although the estimation procedure (*i.e.* *learning from a sketch*) may be algorithmic in nature. Another difference is that, in compressive learning, $\Phi(\cdot)$ is typically randomized. This results in *randomized generalized moments*, which are rarely seen in GeMM.

3.3 COMPRESSIVE LEARNING AND COMPRESSIVE SENSING

As we will show in this section, the sketching mechanism (1) can be interpreted as a dimensionality-reducing linear “projection” of the probability distribution underlying the data set $\{x_i\}_{i=1}^n$. This differs from the traditional approach in signal processing, where dimensionality reduction is performed on the features $x_i \in \mathbb{R}^d$ themselves, rather than the distribution that generates them. Still, many of the intuitions, analyses, and tools designed for feature-based dimensionality reduction can be extended to distribution-based dimensionality reduction.

In the field of *inverse problems* and *compressive sensing* (CS) [4, 24], a signal-of-interest is modeled as a high-dimensional vector $x \in \mathbb{R}^d$, and a physical measurement of that signal is approximated by a linear transformation plus additive noise: $y = Ax + e \in \mathbb{R}^m$ (see Box 4). Here, the linear measurement operator is represented by the $m \times d$ matrix A . In many applications, there is a great motivation to make the measurement dimension much less than the signal dimension, *i.e.*, $m \ll d$. In this case, the recovery of x from y is generally ill-posed. Still, it is possible to accurately recover x from y when both x and A obey certain properties and the noise is small enough.

The celebrated Johnson Lindenstrauss (JL) lemma and its variations specifies, for instance, that with high probability one can nearly preserve the pairwise distances between N features in $\mathbb{R}^d$ by linearly projecting them in a $O(\log N)$ -dimensional domain. This is used, for instance, to accelerate nearest-neighbor searches in large databases. An important extension is to compressive sensing: if x is *sparse*, in that relatively few of its coefficients deviate significantly from zero, and if A satisfies the restricted isometry property (RIP)—an extension of the JL lemma to the continuous set of sparse signals—then one can pose a *regularized* inverse problem whose solution is close to the true x (see Box 4).

[Box 4] Compressive Sensing and the Restricted Isometry Property

In compressive sensing, one observes a linear measurement $y = Ax \in \mathbb{R}^m$ of signal $x \in \mathbb{R}^d$ with $m \ll d$, *i.e.*, with significantly reduced dimension. (For simplicity, we focus on the noiseless case for now.) To distinguish between different signals x in a given signal class, one desires that the distances between all signals in that class are preserved by the measurement operator A . For the class of k -sparse signals Σ_k , *i.e.*, signals with at most k non-zero entries, this property is satisfied up to a tolerance $\delta \in [0, 1)$ when A obeys the restricted isometry property (RIP) [4, 24]

$$(1 - \delta)\|x - x'\|^2 \leq \|Ax - Ax'\|^2 \leq (1 + \delta)\|x - x'\|^2, \quad \forall x, x' \in \Sigma_k. \quad (5)$$

One way to create a RIP-satisfying A is to draw it randomly. For example, if the coefficients of A are drawn i.i.d. zero-mean Gaussian, then A will satisfy the RIP with high probability (w.h.p.) when the sketch dimension m is at least on the order of $k \log(d/k)$ [4].

To recover k -sparse x from y , one might attempt to search for the sparsest signal among all of those that agree with the measurements, *i.e.*, within the set $\mathcal{C}_y := \{u : Au = y\}$. The complexity of this search, however, grows exponentially in k . Fortunately, when A satisfies the RIP, one can provably recover the true x using polynomial-complexity methods [4]. One approach is to solve the convex problem of finding the signal with the smallest ℓ_1 -norm within $\mathcal{C}_y$. Another is to use a greedy algorithm, like orthogonal matching pursuit (OMP), which estimates x by progressively removing from y the k columns of A that best “align” with it, using a least-squares fit.

The RIP also provides guarantees on robust recovery. Suppose that we have noisy measurements $y = Ax + e$, with noise e of bounded norm $\|e\| \leq \varepsilon$. In addition, suppose that x is only *approximately* k -sparse, and use x_k to denote the best k -sparse approximation of x . Finally, consider recovering an estimate $\hat{x}$ of x by searching for the signal with smallest ℓ_1 -norm that agrees with the measurements up to a tolerance of ε , *i.e.*, within the set $\mathcal{C}_{y,\varepsilon} := \{u : \|Au - y\| \leq \varepsilon\}$. Then, if the RIP (5) holds, the estimation error $\hat{x} - x$ satisfies [4]

$$\|\hat{x} - x\| \leq C \frac{\|x - x_k\|_1}{\sqrt{k}} + D\varepsilon, \quad (6)$$

where constants $C, D > 0$ depend only on the value of δ that appears in (5). Thus, the estimation error increases linearly with the noise level ε and the deviation $\|x - x_k\|_1$ from perfect sparsity.

Now that CS has been described, we can clearly connect it to compressive learning. Recall that the sketch (1) $\tilde{z}$ converges to the generalized moment $z := \mathbb{E}[\Phi(X)] = \int p_X(x)\Phi(x)dx$ as the number of data samples n tends to infinity, as per (4).

A crucial observation is the following: *due to the linearity of integration, the sketch z depends linearly on the probability density p_X .* To see it another way, consider a mixture of two densities, $p_X = \alpha p_{X_1} + (1 - \alpha)p_{X_2}$. The corresponding expectation satisfies

$$\mathbb{E}_{X \sim p_X}[\Phi(X)] = \alpha \mathbb{E}_{X_1 \sim p_{X_1}}[\Phi(X_1)] + (1 - \alpha) \mathbb{E}_{X_2 \sim p_{X_2}}[\Phi(X_2)], \quad (7)$$

which implies that the generalized moment z is linear in p_X . With this understanding, we can write

$$z = \mathcal{A}(p_X) := \mathbb{E}_{X \sim p_X}[\Phi(X)], \quad (8)$$

where $\mathcal{A}$ is a linear operator mapping the probability distribution p_X to the m -dimensional sketch vector z .

Remark. Although $\mathcal{A}$ is a linear function of p_X , we emphasize that the feature map $\Phi(x)$ is generally not a linear function of x . This makes compressive learning concretely very different from the vast majority of existing “sketching” mechanisms, which use random linear projections of the data for dimensionality reduction.

With a small modification of the above arguments, we can handle the finite-sample case. Consider the difference between the true generalized moment z and the empirical moment $\tilde{z}$, i.e.,

$$\frac{1}{n} \sum_{i=1}^n \Phi(x_i) - \mathbb{E}[\Phi(X)] =: e. \quad (9)$$

As we discussed earlier, e converges to zero as n tends to infinity by the law of large numbers. Combining (1), (8), and (9), we obtain

$$\tilde{z} = \mathcal{A}(p_X) + e, \quad (10)$$

which shows that the sketch (1) can be interpreted as a “noisy” observation of the data distribution p_X through the linear measurement operator $\mathcal{A}$. Under mild conditions², the central limit theorem can be used to show that $\|e\|$ decays as $1/\sqrt{n}$ w.h.p. [4, Chapter 8]. With the above interpretation of compressive learning, one recovers all of the traditional ingredients of CS:

- The measurement operator $\mathcal{A}$ is linear.
- The measurements $\tilde{z}$ are drastically dimension-reduced. In mathematical terms, probability distributions p_X belong to the infinite-dimensional vector space of so-called *finite* measures [25]. The operator $\mathcal{A}$ maps these infinite-dimensional objects to vectors of finite dimension m .
- The measurement operator $\mathcal{A}$ is typically designed using randomness. This is accomplished by choosing an appropriate randomized feature map $\Phi(\cdot)$, such as one based on RF features, as in (3).

²From the assumed independence of the data samples, this happens for instance if $\mathbb{P}[\|\Phi(X)\| \geq t]$ decays exponentially fast when t increases.

Fig. 9: The analogy between compressive learning, which uses the dimensionality reducing linear measurement $\mathcal{A}(p_X) = \mathbb{E}[\Phi(X)]$ of a distribution p_X , and compressive sensing, which uses the dimensionality reducing linear measurement Ax of a signal x .

- The measurements $\tilde{z}$ are noisy, as per (10).

The analogy between compressive learning and CS is illustrated in Fig. 9. The CS analog to the *sketching phase* (1) is the *signal sensing phase*, where the signal x is *linearly* mapped to the observation vector y . The analog to *sparse recovery*, where an estimate of x is computed from the observation y by solving an optimization problem, is to *learn from a sketch*, where an estimate of the data distribution p_X (or of distributional parameters θ of interest) is computed from the sketch $\tilde{z}$. Its expression as an optimization problem will be further discussed in Section 4.

3.4 RANDOM FOURIER SAMPLING AND SUPER-RESOLUTION RECOVERY

In this section, we consider the specific case of compressive clustering, which allows us to forge a concrete connection between CS and compressive learning using random Fourier sampling and super-resolution recovery.

We begin by considering the goal of recovering k centroids $\{c_\ell\}_{\ell=1}^k$ in $\mathbb{R}^d$ with Euclidean norm $\leq r$ that are separated from each other in Euclidean distance by $\geq \epsilon$. A naive approach would be to discretize the d -dimensional cube of side-length $2r$ with a grid spacing of ϵ , leading to $N = (2r/\epsilon)^d$ bins. A valid sketch of the data $\mathcal{X}$ is obtained by simply computing the histogram $\hat{p} \in \mathbb{R}_+^N$ over these bins (*i.e.*, by using the “binning” feature map). However, the dimension of this sketch, N , grows exponentially in the feature dimension, d . To construct a smaller sketch, one might reason that, if the data clusters tightly around k points, then $\hat{p}$ is close to a k -sparse vector. In this case, ideas from CS can be directly exploited. In particular, one could use a sketched histogram [15] of the form $\tilde{z} = A\hat{p}$, where matrix $A \in \mathbb{R}^{m \times N}$ is randomly drawn with i.i.d. Gaussian components. Here, “learning from a sketch” means recovering the centroids. For this, one would first search for the best non-negative, k -sparse, sum-to-one vector using

$$\tilde{p} = \arg \min_{p \in \Sigma_k^+} \|\tilde{z} - Ap\|^2, \quad (11)$$

(or a convex or greedy relaxation of this problem) where Σ_k^+ here denotes the set of k -sparse, non-negative, sum-to-one vectors. Then, one would identify the k grid locations in $\mathbb{R}^d$ corresponding to the non-zero indices of $\tilde{p}$. CS theory [4] says that the support of $\tilde{p}$ will be accurate for sketch dimensions m at least on the order of $k \log N$, *i.e.*, at least on the order of $kd \log(r/\epsilon)$. Although this latter approach substantially reduces the dimension of the sketch, practical challenges remain when the feature dimension d is large. For example, the number of columns, N , in the compression matrix A grows exponentially in d , making storage and multiplication by A impractical.

An alternative approach could be to construct A using m rows of the (d -dimensional, in this case) discrete Fourier transform (DFT) matrix, *i.e.*, by sampling the DFT at m (d -dimensional) frequencies. In this case, multiplication-by- A could be implemented by the fast Fourier transform (FFT) algorithm, and thus the matrix A would not need to be explicitly stored. Fourier-domain sampling is a familiar operation

in the context of signal processing, as it forms the cornerstone for radar, medical imaging, and radio interferometry, see *e.g.*, [26, 27]. When the m frequencies are drawn uniformly at random, CS theory [4, 24] has established that accurate recovery of an N -length k -sparse signal can be accomplished (with high probability) when m is on the order of $k \log^3(k) \log N$ [4, Corollary 12.38]. Since $N = (2r/\epsilon)^d$ here, this would mandate sketch dimensions m at least on the order of $kd \log^3(k) \log(r/\epsilon)$. Although the FFT avoids the need to store $\mathbf{A}$ as an explicit matrix and allows efficient computation of the sketch, the cost of solving the optimization problem (11) using existing convex relaxations or greedy approaches is impractical due to the need to manipulate N -dimensional vectors, where N grows exponentially in d .

Until now, we considered discretizing the d -dimensional feature space on an ϵ -spaced hypergrid within a $2r$ -sidelength hypercube, but found that this requires manipulating vectors (*e.g.*, a histogram $\hat{\mathbf{p}}$) whose dimension grows exponentially in d . We can avoid this discretization (*i.e.*, take $\epsilon \rightarrow 0$ and $r \rightarrow \infty$) by replacing the DFT with the *continuous* Fourier transform, in which case we are Fourier-transforming the empirical distribution $\hat{p}_{\mathcal{X}}(\mathbf{x}) := \frac{1}{n} \sum_{i=1}^n \delta(\mathbf{x} - \mathbf{x}_i)$ rather than its N -bin histogram $\hat{\mathbf{p}}$. The sketch $\tilde{\mathbf{z}}$ then has components

$$\tilde{z}_j = \int_{\mathbb{R}^d} \hat{p}_{\mathcal{X}}(\mathbf{x}) \exp(-j2\pi \mathbf{w}_j^\top \mathbf{x}) d\mathbf{x} \quad (12)$$

$$= \frac{1}{n} \sum_{i=1}^n \exp(-j2\pi \mathbf{w}_j^\top \mathbf{x}_i) \quad j = 1, \dots, m, \quad (13)$$

where $\{\mathbf{w}_j\}_{j=1}^m$ are d -dimensional frequencies that are drawn at random. Typically, $\mathbf{w}_j$ are drawn i.i.d. Gaussian, but other distributions can be used, as discussed in Sections 5.6 and 6.1. Note that (13) corresponds precisely to the sketch (1) with the random Fourier (RF) feature map $\Phi(\cdot)$ from (3). Taking a statistical perspective, the components $\tilde{z}_j$ in (12) can also be recognized as samples of the *characteristic function* of $\hat{p}_{\mathcal{X}}$. Recall that, for a density p , the characteristic function Ψ_p is defined as

$$\Psi_p(\mathbf{w}) := \int_{\mathbb{R}^d} p(\mathbf{x}) \exp(j2\pi \mathbf{w}^\top \mathbf{x}) d\mathbf{x} = \mathbb{E}_{\mathbf{X} \sim p}[\exp(j2\pi \mathbf{w}^\top \mathbf{X})]. \quad (14)$$

[Box 5] Super-Resolution Recovery

Super-resolution (SR) is a general class of techniques to enhance the resolution of a sensing system, *e.g.*, to observe sub-wavelength features in astronomy or medical imaging [25]. The problem addressed by SR is to recover a continuous-time (when dimension $d = 1$), or continuous-space (when dimension $d \geq 2$) sparse signal $s(\mathbf{t})$ from a few, possibly noisy Fourier measurements $\{y_j\}_{j=1}^m$. This amounts to recovering a weighted sum $s(\mathbf{t}) = \sum_{\ell=1}^k \alpha_\ell \delta(\mathbf{t} - \mathbf{t}_\ell)$ of k Diracs with amplitudes α_ℓ and locations $\mathbf{t}_\ell \in \mathbb{R}^d$ from

$$y_j = \int_{\mathbb{R}^d} s(\mathbf{t}) \exp(-j2\pi \mathbf{w}_j^\top \mathbf{t}) d\mathbf{t} + e_j, \quad j = 1, \dots, m, \quad (15)$$

with measurement noise e_j and frequency vectors $\{\mathbf{w}_j\}_{j=1}^m$ in $\mathbb{R}^d$. Recovery of $s(\cdot)$ can be posed as an infinite-dimensional convex problem on measures [25]. However, most reconstruction algorithms involve non-convex steps [28]. When the frequencies $\mathbf{w}_j$ are drawn randomly, the signal can be accurately recovered with high probability when m is of the order of at least kd^3 , up to log factors. Proving this usually requires additional assumptions, such as a minimal separation between the locations $\mathbf{t}_\ell$ [29] or positivity of the amplitudes α_ℓ [28].

The link between super-resolution and compressive clustering follows from rewriting (12) as

$$\tilde{z}_j = \int_{\mathbb{R}^d} p_{\mathcal{X}}(\mathbf{x}) \exp(-j2\pi \mathbf{w}_j^\top \mathbf{x}) d\mathbf{x} + e_j, \quad j = 1, \dots, m, \quad (16)$$

where e_j captures the “noise” due to finite-sample effects (recall (9)). Comparing (16) to (15), we see that they are mathematically equivalent when $p_{\mathcal{X}}(\mathbf{x}) = \sum_{\ell=1}^k \alpha_\ell \delta(\mathbf{x} - \mathbf{c}_\ell)$, except for the fact that, in the case of compressive clustering, α_ℓ are non-negative and sum to one.

Intuitively, when a probability distribution p has a “simple” structure, one can recover it (with high probability) from enough randomly chosen samples of its Fourier transform. Centroid recovery from the

sketch (12) is premised on the empirical distribution $\hat{p}_X$ being well approximated by a mixture of k Diracs, *i.e.*, $\hat{p}_X(\mathbf{x}) \approx \sum_{\ell=1}^k \alpha_\ell \delta(\mathbf{x} - \mathbf{c}_\ell)$. In this case, centroid recovery parallels the “super-resolution” recovery problem (see Box 5) through an optimization problem that is the continuous analog to (11) (or a convex or greedy relaxation for the continuous case) and will be further elaborated in the next section.

In both problems, recovery guarantees are possible when the frequencies w_j are randomly drawn. For example, when the centroids $\{\mathbf{c}_\ell\}_{\ell=1}^k$ are ϵ -separated and r -bounded, centroid recovery guarantees have been established provided the sketch dimension m is on the order of $k^2 d \log(r/\epsilon)$, omitting—for simplicity—some log factors involving k and d [5, 6]. Similar guarantees hold when $\hat{p}_X$ is approximately a sum of spatially localized components (*e.g.*, in compressive GMM) [5, 6].

4 LEARNING FROM A SKETCH

Until now, we have primarily focused on the first stage of the compressive-learning pipeline (see Fig. 1), where the dataset X is sketched down to $\tilde{z}$, a compressed and noisy representation of the underlying data-generating distribution p_X . We now discuss the second stage of the pipeline, where the distributional parameters of interest, θ , are recovered from the sketch $\tilde{z}$. The close analogy between sketching and CS allows us to cast this “parameter learning” stage as an optimization problem, *i.e.*,

$$\tilde{\theta} = \arg \min_{\theta} C(\theta | \tilde{z}), \quad (17)$$

where the cost function $C(\cdot | \tilde{z})$ is adapted to the considered learning task. As in CS, many candidate distributions p_X (and hence many candidate parameters θ) can yield the same sketch $\tilde{z}$. Thus, to make the inverse problem well-posed, one needs to employ *concrete modeling assumptions* and *regularization*, both of which can take several forms. As in CS, we will assume that the sketched quantity p_X is of low intrinsic complexity, *i.e.*, close to some family of “simple” probability distributions.

As a first example, we consider the problem of learning a mixture model from a sketch. Similar to a sparse vector $\mathbf{x}$, which is a linear combination of a few elements of the standard basis, a *mixture model* p_X is a linear combination of a few “simple” densities $\{p_{\theta_\ell}\}_{\ell=1}^k$. Concretely, $p_X = \sum_{\ell=1}^k \alpha_\ell p_{\theta_\ell}$, where the mixture weights $\{\alpha_\ell\}_{\ell=1}^k$ are non-negative and sum to one. For example, with a Gaussian Mixture Model (GMM), we have that $p_{\theta_\ell} = \mathcal{N}(\mu_\ell, \Sigma_\ell)$, where $\theta_\ell = \{\mu_\ell, \Sigma_\ell\}$ contains the mean μ_ℓ and covariance Σ_ℓ . If p_X is well approximated by the mixture model $\sum_{\ell=1}^k \alpha_\ell p_{\theta_\ell}$, then, according to (10), the sketch $\tilde{z}$ is well approximated by the linear combination $\sum_{\ell=1}^k \alpha_\ell \mathcal{A}(p_{\theta_\ell})$. Hence, one could try to extract the mixture parameters, $\theta = \{\alpha_\ell, \theta_\ell\}_{\ell=1}^k$, from the sketch $\tilde{z}$ by solving the (nonconvex) optimization problem (17) with

$$C(\theta | \tilde{z}) := \left\| \tilde{z} - \sum_{\ell=1}^k \alpha_\ell \mathcal{A}(p_{\theta_\ell}) \right\|^2. \quad (18)$$

The cost $C(\theta | \tilde{z})$ can be interpreted as the negative log-likelihood (up to a shift and scale) of θ given the sketch $\tilde{z}$, under the classic modeling assumption of i.i.d. Gaussian measurement noise e in (10).

When the RF feature map $\Phi(\cdot)$ is used to compute the sketch $\tilde{z}$ and the component densities p_{θ_ℓ} are Gaussian or α -stable, there exist analytic expressions for $\mathcal{A}(p_{\theta_\ell})$ and for the gradient of $\mathcal{A}(p_{\theta_\ell})$ with respect to the mixture parameters in θ_ℓ [8]. These expressions are convenient when numerically optimizing (18). For instance, greedy approaches, similar to the OMP algorithm for CS (recall Box 4), can be used [8] to estimate the parameters $\{\alpha_\ell, \theta_\ell\}_{\ell=1}^k$. These approaches sequentially estimate and subtract, from $\tilde{z}$, each of the k components $\alpha_\ell \mathcal{A}(p_{\theta_\ell})$ that best align with it, where “best” is measured via the correlation between $\tilde{z}$ and $\mathcal{A}(p_{\theta_\ell})$. As another example, an iterative approach [12] was proposed that exploits the log-likelihood interpretation of $C(\theta | \tilde{z})$ in (18) and the i.i.d. random nature of the linear transform W in the RF map (3). In [8], applications of compressive mixture-modeling are demonstrated on speaker verification (see Box 1) and source separation.

As a second example, we consider compressive k -means clustering. Here, the goal is to recover, from the sketch $\tilde{z}$, the centroids $\theta = \{\mathbf{c}_\ell\}_{\ell=1}^k$ that minimize the average squared Euclidean distance from each sample to its nearest centroid, *i.e.*, $\frac{1}{n} \sum_{i=1}^n \min_{\ell} \|\mathbf{x}_i - \mathbf{c}_\ell\|^2$. To tackle this k -means problem, we view it as an approximation of a particular GMM fitting problem. In particular, suppose that the probability distribution p_X is a GMM with weights α_ℓ , mean vectors $\mathbf{c}_\ell$, and covariance matrices Σ_ℓ .

Then, in the special case that $\alpha_\ell = 1/k$ and $\Sigma_\ell = \sigma^2 \mathbf{I}$ for all components $1 \leq \ell \leq k$, we can write the likelihood as $\prod_{i=1}^n p(\mathbf{x}_i | \boldsymbol{\theta}) \propto \prod_{i=1}^n \sum_{\ell=1}^k \exp(-\frac{1}{2\sigma^2} \|\mathbf{x}_i - \mathbf{c}_\ell\|^2)$, and so the negative log-likelihood becomes $-\sum_{i=1}^n \log \sum_{\ell=1}^k \exp(-\frac{1}{2\sigma^2} \|\mathbf{x}_i - \mathbf{c}_\ell\|^2)$ up to an additive constant. We can then use the log-sum-exp approximation $\log \sum_{\ell} \exp(f_\ell(\mathbf{x})) \approx \max_{\ell} f_\ell(\mathbf{x})$ to approximate this latter expression as $\frac{1}{2\sigma^2} \sum_{i=1}^n \min_{\ell} \|\mathbf{x}_i - \mathbf{c}_\ell\|^2$, which agrees with the k -means cost up to a scaling. If we furthermore consider the case of a vanishing variance $\sigma^2 \rightarrow 0$, then the component density $p_{\boldsymbol{\theta}_\ell}$ reduces to a point mass, *i.e.*, $p_{\boldsymbol{\theta}_\ell}(\mathbf{x}) \rightarrow \delta(\mathbf{x} - \mathbf{c}_\ell)$. In this limiting case, the linear measurement of the point-mass $p_{\boldsymbol{\theta}_\ell}$ is $\mathcal{A}(p_{\boldsymbol{\theta}_\ell}) = \mathbb{E}_{X \sim p_{\boldsymbol{\theta}_\ell}}[\Phi(X)] = \int \Phi(\mathbf{x}) p_{\boldsymbol{\theta}_\ell}(\mathbf{x}) d\mathbf{x} = \Phi(\mathbf{c}_\ell)$ according to (8).

Thus, with these justifications, the cost function (18) for compressive GMM would change to

$$C(\boldsymbol{\theta} | \tilde{\mathbf{z}}) := \left\| \tilde{\mathbf{z}} - \frac{1}{k} \sum_{\ell=1}^k \Phi(\mathbf{c}_\ell) \right\|^2 \quad (19)$$

for compressive k -means. If we do not want to assume that $\alpha_\ell = 1/k$ for each ℓ , we could instead estimate $\{\alpha_\ell\}_{\ell=1}^k$ from the sketch, leading to the cost function suggested in [11]:

$$C(\boldsymbol{\theta} | \tilde{\mathbf{z}}) := \min_{\boldsymbol{\alpha}} \left\| \tilde{\mathbf{z}} - \sum_{\ell=1}^k \alpha_\ell \Phi(\mathbf{c}_\ell) \right\|^2. \quad (20)$$

Similar to the compressive GMM problem described earlier, minimization of the cost function (20) for compressive k -means can be tackled by greedy approaches, as described in [11]. Despite the fact that (20) does not directly minimize the k -means cost, it has been shown empirically [11] that the centroids estimated by such greedy algorithms nearly minimize this cost, see, *e.g.*, Box 2 for an example on MNIST data. This claim is also supported by theoretical results guaranteeing that the minimizer of (20) is endowed with statistical-learning guarantees with respect to the original k -means cost [5, 6]. Such guarantees will be discussed shortly.

Depending on the choice of parameters $\boldsymbol{\theta}$ and the feature map $\Phi(\cdot)$, the form of the optimization problem posed to learn-from-a-sketch can differ considerably from that for GMMs in (18) and that for k -means in (19)-(20). Consider, for example, learning-from-a-sketch for PCA. As described earlier, the parameter $\boldsymbol{\theta}$ of interest is the k -dimensional subspace that best fits the d -dimensional data $\{\mathbf{x}_i\}_{i=1}^n$ in a least-squares sense. It is well known that this subspace is spanned by k principal eigenvectors of the empirical autocorrelation matrix $\hat{\mathbf{R}} = \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i^\top$, or—equivalently—the column space of the (symmetric positive semi-definite) matrix $\hat{\mathbf{R}}_k$ that is closest to $\hat{\mathbf{R}}$ in the Frobenius norm:

$$\hat{\mathbf{R}}_k := \arg \min_{\mathbf{R}: \text{rank}(\mathbf{R}) \leq k} \|\hat{\mathbf{R}} - \mathbf{R}\|_F = \arg \min_{\mathbf{R}: \text{rank}(\mathbf{R}) \leq k} \|\text{vec}(\hat{\mathbf{R}}) - \text{vec}(\mathbf{R})\|^2. \quad (21)$$

When sketching using quadratic features $\Phi(\mathbf{x}) = ((\mathbf{w}_1^\top \mathbf{x})^2, \dots, (\mathbf{w}_m^\top \mathbf{x})^2)$ with random $\mathbf{w}_j \in \mathbb{R}^d$, the j th component of the sketch becomes $\tilde{z}_j = \frac{1}{n} \sum_{i=1}^n \mathbf{w}_j^\top \mathbf{x}_i \mathbf{x}_i^\top \mathbf{w}_j = \mathbf{w}_j^\top \hat{\mathbf{R}} \mathbf{w}_j$. Importantly, this $\tilde{z}_j$ is a linear function of $\hat{\mathbf{R}}$, and so there exists an $m \times d^2$ matrix $\mathbf{A}$ such that $\tilde{\mathbf{z}} = \mathbf{A} \text{vec}(\hat{\mathbf{R}})$. Thus, by analogy with (21), one could first fit a symmetric positive semi-definite low-rank matrix to the sketch $\tilde{\mathbf{z}}$ via

$$\tilde{\mathbf{R}} = \arg \min_{\mathbf{R}: \text{rank}(\mathbf{R}) \leq k, \mathbf{R}^\top = \mathbf{R}, \mathbf{R} \succeq 0} \|\tilde{\mathbf{z}} - \mathbf{A} \text{vec}(\mathbf{R})\|^2, \quad (22)$$

and then set the parameter estimate $\tilde{\boldsymbol{\theta}}$ equal to the column space of $\tilde{\mathbf{R}}$. While the optimum of (21) is automatically symmetric and positive semi-definite, this property needs to be enforced explicitly in (22).

The low-rank matrix recovery problem (22) has been thoroughly investigated by the signal processing and machine learning communities (*e.g.*, [30]). It arises, for example, in applications such as collaborative filtering for recommender systems and signal reconstruction from phaseless measurements. Early approaches to solving the non-convex problem (22) involved convex relaxation via nuclear-norm regularization [4, Chapter 4]. More recent approaches exploit the non-convex geometry of (22) [31].

TABLE I: Loss functions for our four running examples

Running example	Parameters θ	Loss function $L(\theta x_i)$
Principal component analysis	Orthonormal basis $\theta = \{u_\ell\}_{\ell=1}^k$	$\sum_{\ell=1}^k x_i^\top u_\ell ^2$
Least-squares linear regression	A weight matrix θ	$\ x_{1i} - \theta x_{2i}\ ^2$, $x_i := (x_{1i}^\top, x_{2i}^\top)^\top$
Gaussian mixture modeling	GMM parameters $\theta = \{\alpha_\ell, \mu_\ell, \Sigma_\ell\}_{\ell=1}^k$	$-\log \sum_{\ell=1}^k \alpha_\ell \mathcal{N}(x_i; \mu_\ell, \Sigma_\ell)$
k -means clustering	A set of centroids $\theta = \{c_\ell\}_{\ell=1}^k$	$\min_\ell \ x_i - c_\ell\ ^2$

5 COMPRESSIVE LEARNING WITH THEORETICAL GUARANTEES

In the previous section, learning-from-a-sketch was posed as the optimization problem (17), repeated here for convenience:

$$\tilde{\theta} = \arg \min_{\theta} C(\theta|\tilde{z}). \quad (23)$$

The quantity $C(\cdot|\tilde{z})$ is a real-valued *cost function* whose minimizer $\tilde{\theta}$ is the “best” (in some sense) estimate of θ from the sketch $\tilde{z}$. This approach contrasts with traditional statistical learning [32, 33], which computes the parameter estimate

$$\hat{\theta} := \arg \min_{\theta} R(\theta|\mathcal{X}) \quad \text{with } R(\theta|\mathcal{X}) := \frac{1}{n} \sum_{i=1}^n L(\theta|x_i), \quad (24)$$

where $R(\cdot|\mathcal{X})$ is the “empirical risk” and $L(\cdot|x_i)$ is the “loss function” for the i th sample. Table I presents typical loss functions for our four running examples. In most cases, performing the minimization in (24) involves querying the full training dataset $\mathcal{X}$ many times, *e.g.*, for stochastic gradient descent. In compressive learning, the cost $C(\cdot|\tilde{z})$, which depends only on the low-dimensional sketch $\tilde{z}$, is used as a *surrogate* for the empirical risk. Because the dimension of the sketch is so much smaller than the cardinality of the full dataset $\mathcal{X}$, the minimization in (17) can be made much more efficient than that in (24).

Of course, the estimate $\tilde{\theta}$, which minimizes the cost $C(\cdot|\tilde{z})$ is not, in general, the same as $\hat{\theta}$, which minimizes the empirical risk $R(\cdot|\mathcal{X})$. Both, however, are approximations of the *ideal* estimate

$$\theta^* := \arg \min_{\theta} R^*(\theta) \quad \text{with } R^*(\theta) := \mathbb{E}[L(\theta|X)], \quad (25)$$

where $R^*(\theta)$ is known as the “true risk.” Indeed, the feature map $\Phi(\cdot)$ and cost $C(\cdot|\tilde{z})$ are designed precisely to ensure that the surrogate minimization (23) approximates the true-risk minimization (25).

5.1 EXCESS RISK GUARANTEES

To establish a guarantee on the goodness of an estimate θ , one must prove that the “excess risk” $R^*(\theta) - R^*(\theta^*) \geq 0$ is small. Indeed, the true risk can be interpreted as the *expected* loss on test samples that have the same generating distribution p_X as the training samples $\mathcal{X}$, but that are drawn independently and not accessible at training time. In statistical learning tasks, the true risk $R^*(\cdot)$ is the primary metric by which one judges the quality of an arbitrary estimate θ . Note that controlling the excess risk is different from proving that θ is close to the ideal estimate θ^* in, *e.g.*, the Euclidean distance $\|\theta - \theta^*\|$. Consider, for example, the problem of fitting a one-dimensional linear subspace to a dataset (*i.e.*, PCA with $k = 1$). When the two largest eigenvalues of the autocorrelation matrix R are equal, any nontrivial linear combination of the corresponding eigenvectors generates a one-dimensional subspace θ with minimum true risk. The problem of estimating a subspace “close to the optimal one” is thus ill-defined, yet finding a subspace with close-to-optimal performance is well-defined and achievable, both by classic PCA and compressive PCA.

[Box 6] Traditional Statistical Learning versus Compressive Learning

Fig. 10: k -means clustering of one-dimensional data $\mathcal{X}$, showing cost function $C(\theta | \tilde{z})$, risk $R^*(\theta)$, ideal estimate θ^* , compressive learning estimate $\tilde{\theta}$, and excess risk $R^*(\tilde{\theta}) - R^*(\theta^*)$, for a sketch of small dimension m (top) and for a sketch of moderate size m (bottom).

In statistical learning, the ideal parameter θ^* minimizes the true risk $R^*(\theta) = \mathbb{E}_{X \sim p_X} [L(\theta | X)]$. The performance of any estimate θ is measured according to its excess risk, *i.e.*, $R^*(\theta) - R^*(\theta^*)$. In compressive learning, one obtains $\tilde{\theta}$ by minimizing a cost function $C(\theta | \tilde{z})$, where the sketch $\tilde{z}$ is a compressed version of a finite-size dataset $\mathcal{X}$ with samples drawn from distribution p_X .

To gain intuition into how these quantities manifest in compressive learning, Fig. 10 shows a simple example: compressive k -means clustering of one-dimensional data $\mathcal{X} = \{x_i\}_{i=1}^n$ with two centroids, θ_1 and θ_2 . The true risk $R^*(\cdot)$ and the cost $C(\cdot | \tilde{z})$ are plotted versus $\theta = (\theta_1, \theta_2)$ and versus a 1-D slice in the θ plane. As the sketch dimension m increases, it can be seen that the excess risk $R^*(\tilde{\theta}) - R^*(\theta^*)$ decreases. Moreover, although the cost function $C(\cdot | \tilde{z})$ is nonconvex, it is approximately quadratic in a large basin of attraction around its global minimizer, suggesting that gradient-descent algorithms will behave well when properly initialized.

Despite the fact that the estimates $\tilde{\theta}$, $\hat{\theta}$, and θ^* minimize different objective functions, they often yield similar true-risks, *i.e.*, the excess risks of $\tilde{\theta}$ and $\hat{\theta}$ are provably small. For $\hat{\theta}$, the proofs use classic results from statistical learning [32, 33], under assumptions that we will briefly discuss in the sequel. For $\tilde{\theta}$, with an appropriately designed feature map $\Phi(\cdot)$ and cost function $C(\cdot | \tilde{z})$, the proofs informally follow from the fact that $C(\cdot | \tilde{z})$ and $R^*(\cdot)$ have a similar shape. This fact is illustrated in Box 6 for a toy example of compressive clustering. There it can be seen that, with a properly designed feature map $\Phi(\cdot)$, a well-chosen cost function $C(\cdot | \tilde{z})$, and a sufficiently large sketch dimension m , the minimizer $\tilde{\theta}$ of the cost yields a nearly minimal true risk $R^*(\tilde{\theta})$ and lives in a large basin-of-attraction in $C(\cdot | \tilde{z})$. If the sketch dimension m is chosen too small, however, then the excess risk $R^*(\tilde{\theta}) - R^*(\theta^*)$ increases.

A theory of compressive learning [5, 6] has been developed to better understand how (random) feature maps $\Phi(\cdot)$ and cost functions $C(\cdot | \tilde{z})$ can be designed to ensure that the sketch $\tilde{z}$ captures sufficient information to control the excess risk on $\tilde{\theta}$. **In the following four subsections, we aim to summarize the key aspects of this theory using broadly accessible language.** For those interested in the full technical details, they can be found in [5, 6].

In a nutshell, the theory relies on interpreting $\tilde{\theta}$ as the minimizer of the risk attributed to a *surrogate* probability distribution $\tilde{p}$ with the following properties:

- $\tilde{p}$ is “close” to the distribution p_X , as measured by a task-driven distance $d(\tilde{p}, p_X)$;

- $\tilde{p}$ is a “simple” distribution (e.g., for compressive GMM, $\tilde{p}$ is a Gaussian mixture);
- $\tilde{p}$ minimizes $\|\tilde{z} - \mathcal{A}(p)\|$ among all simple distributions p .

For example, in compressive GMM, this holds when $\tilde{p} = p_{\tilde{\theta}}$, i.e., the Gaussian mixture distribution parameterized by $\tilde{\theta}$. Given such a $\tilde{p}$, the theory can be explained in the following step-by-step manner:

- Just as in traditional statistical learning, excess risk bounds can be established if the task-driven distance $d(\tilde{p}, p_X)$ is controlled (see Section 5.2).
- Given a family of “simple” probability distributions (see Section 5.3), this distance can indeed be controlled provided a certain “lower restricted isometry property” (LRIP) holds (see Section 5.4). This is analogous to traditional CS, where the RIP allows one to control the performance of sparse signal recovery.
- To establish the LRIP for a random feature maps $\Phi(\cdot)$, one can build on connections between kernel methods and the JL Lemma (see Section 5.5).
- Finally, to design task-specific feature maps $\Phi(\cdot)$, one can appeal to the problem of kernel design (see Section 5.6).

5.2 TASK-DRIVEN DISTANCES AND EXCESS-RISK BOUNDS

In traditional statistical learning, it is common to define a distance between the true distribution p_X and an arbitrary surrogate p'_X as

$$d(p'_X, p_X) = \sup_{\theta} |R^*(\theta|p'_X) - R^*(\theta|p_X)|, \quad (26)$$

where $R^*(\theta|p) := \mathbb{E}_{X \sim p}[L(\theta|X)]$ denotes the risk under an arbitrary distribution p . This distance is task specific, since the loss function $L(\cdot|\cdot)$ depends on the estimation task. The utility of $d(p'_X, p_X)$ for assessing the goodness of parameter estimation follows from the inequalities

$$\begin{aligned} R^*(\theta^*|p_X) &\leq R^*(\theta^{*'}|p_X) \leq R^*(\theta^{*'}|p'_X) + d(p'_X, p_X) \\ &\leq R^*(\theta^*|p'_X) + d(p'_X, p_X) \\ &\leq R^*(\theta^*|p_X) + 2d(p'_X, p_X), \end{aligned} \quad (27)$$

where $\theta^* := \arg \min_{\theta} R^*(\theta|p_X)$ and $\theta^{*'} := \arg \min_{\theta} R^*(\theta|p'_X)$. In particular, these inequalities yield the following bounds on the excess risk of θ^{*} :

$$0 \leq R^*(\theta^{*'}|p_X) - R^*(\theta^*|p_X) \leq 2d(p'_X, p_X). \quad (28)$$

For example, if p'_X equals the empirical distribution $\hat{p}_X(x) := \frac{1}{n} \sum_{i=1}^n \delta(x - x_i)$, then (28) bounds the “generalization error” of *empirical risk minimization*, i.e., the error incurred when training with $\mathcal{X} = \{x_i\}_{i=1}^n$ but testing with independent samples drawn from p_X . This is the reasoning behind many classic statistical-learning guarantees, where so-called “uniform convergence” results establish that, under appropriate conditions, $d(\hat{p}_X, p_X) = O(1/\sqrt{n})$ with high probability on the draw of i.i.d. training samples $\mathcal{X} = \{x_i\}_{i=1}^n$. The term “uniform convergence” stems from the analogy between the distance (26) and the ℓ_{∞} norm.

Similarly, excess risk bounds for compressive learning are achieved by interpreting the minimizer $\tilde{\theta}$ of (23) as the minimizer of the risk $R^*(\theta|\tilde{p})$, for some “simple” probability distribution $\tilde{p}$ that minimizes $\|\tilde{z} - \mathcal{A}(\tilde{p})\|$, and by controlling the distance $d(\tilde{p}, p_X)$. To control this distance, key geometric intuitions exploit the connections between compressive learning and CS, as illustrated in Fig. 9.

5.3 EXPLOITING SIMPLIFIED MODELS TO LEARN FROM A SKETCH

Recall that, in CS, the goal is to recover the best k -sparse approximation of an unknown vector $x_0 \in \mathbb{R}^d$ given the noisy linear measurement $y = Ax_0 + e \in \mathbb{R}^m$, where $m \ll d$. Ideally, recovery aims to find the k -sparse vector whose (noiseless) measurement is closest to the observed y , i.e.,

$$\hat{x} = \arg \min_{x \in \Sigma_k} \|y - Ax\|^2. \quad (29)$$

In practice, convex or greedy approximations of this combinatorial approach are often used. Recovery guarantees can be established using the *restricted isometry property* (RIP) in (5), which says that the Euclidean distance $\|\mathbf{A}\mathbf{x} - \mathbf{A}\mathbf{x}'\|$ between the (noiseless) measurements of two k -sparse vectors is almost the same as the Euclidean distance $\|\mathbf{x} - \mathbf{x}'\|$ between the vectors themselves (see Box 4). These guarantees require that the sparsity k is sufficiently small for a given m and d .

To connect compressive learning to CS, we first reframe the goal of learning-from-a-sketch as that of recovering a parametric model *distribution* $p_\theta \in \Sigma_\Theta$ from the sketch $\tilde{\mathbf{z}}$, rather than recovering the model parameters $\theta \in \Theta$ themselves. Here, $\Sigma_\Theta := \{p_\theta : \theta \in \Theta\}$ denotes some set of admissible distributions p_θ . For example, in compressive GMM, the goal becomes that of recovering a k -term GMM $p_\theta = \sum_{\ell=1}^k \alpha_\ell \mathcal{N}(\mu_\ell, \Sigma_\ell)$ rather than the GMM parameters $\theta = \{\alpha_\ell, \mu_\ell, \Sigma_\ell\}_{\ell=1}^k$. Likewise, in compressive PCA, the goal becomes that of recovering a distribution p_θ whose autocorrelation matrix has rank of at most k , rather than a k -dimensional orthonormal basis θ for the column space of that autocorrelation matrix. Then, mirroring the CS recovery approach (29), compressive learning recovery aims to find a parametric distribution $p_\theta \in \Sigma_\Theta$ whose (noiseless) sketch is closest to the observed sketch $\tilde{\mathbf{z}}$, *i.e.*,

$$\tilde{p} = p_{\tilde{\theta}} := \arg \min_{p_\theta \in \Sigma_\Theta} \|\tilde{\mathbf{z}} - \mathcal{A}(p_\theta)\|^2. \quad (30)$$

The construction of Σ_Θ implies that the parameters $\tilde{\theta}$ defining $p_{\tilde{\theta}}$ minimize a certain cost function, *i.e.*,

$$\tilde{\theta} = \arg \min_{\theta \in \Theta} C(\theta | \tilde{\mathbf{z}}) \quad \text{with } C(\theta | \tilde{\mathbf{z}}) = \|\tilde{\mathbf{z}} - \mathcal{A}(p_\theta)\|^2. \quad (31)$$

Note the similarity between (31) and (17)-(18). Moreover, $\tilde{\theta}$ also minimizes the risk $R^*(\theta | \tilde{p})$ [5, 6].

From the description above, we see that a central theme of both CS and compressive learning is that *measurement compression makes it impossible to recover the full object-of-interest* (*i.e.*, $\mathbf{x}_0$ in CS or p_X in compressive learning) without additional side-information. For this reason, both seek to recover the parameters of a *simplified model* of the object-of-interest. In CS, this is accomplished by seeking to recover the best k -sparse approximation to $\mathbf{x}_0$, rather than $\mathbf{x}_0$ itself. In compressive learning, this is accomplished by seeking to recover the best model distribution $p_\theta \in \Sigma_\Theta$, rather than the true distribution p_X . In both cases, if the linear operator (*i.e.*, $\mathbf{A}$ in CS or $\mathcal{A}(\cdot)$ in compressive learning) is well designed, then the model parameters (in Σ_k for CS or in Σ_Θ for compressive learning) can be accurately recovered.

5.4 THE LRIP AND EXCESS-RISK CONTROL

Recovery guarantees can be established using a tool analogous to the RIP (5) in CS.

Take, for example, the case of compressive PCA. As discussed around (22), one could use a sketch of the form $\tilde{\mathbf{z}} = \mathbf{A} \text{vec}(\hat{\mathbf{R}}) \in \mathbb{R}^m$, where $\hat{\mathbf{R}} = \frac{1}{n} \sum_{i=1}^n \mathbf{x}_i \mathbf{x}_i^\top$ is the $d \times d$ empirical autocorrelation matrix and $m \ll d^2$, and then search for $\tilde{\mathbf{R}}$, the symmetric matrix with rank at-most- k that minimizes the cost $C(\mathbf{R} | \tilde{\mathbf{z}}) = \|\tilde{\mathbf{z}} - \mathbf{A} \text{vec}(\mathbf{R})\|^2$. Recovery guarantees can be established using an RIP of the form

$$(1 - \delta) \|\mathbf{R} - \mathbf{R}'\|_F^2 \leq \|\mathbf{A} \text{vec}(\mathbf{R}) - \mathbf{A} \text{vec}(\mathbf{R}')\|^2 \leq (1 + \delta) \|\mathbf{R} - \mathbf{R}'\|_F^2, \quad \forall \mathbf{R}, \mathbf{R}' \in \Sigma_k, \quad (32)$$

where $\delta \in (0, 1)$ is a tolerance and Σ_k is now the set of $d \times d$ symmetric rank- k matrices. It is known [4, Chapter 9] that i.i.d. Gaussian $\mathbf{A}$ satisfy the RIP (32) with high probability when m is on the order of kd times a constant that depends on the tolerance δ . Furthermore, the cost-minimizing $\tilde{\mathbf{R}}$ is near-optimal in the sense of minimizing the excess risk, even when a convex relaxation of the cost is used [5, 6].

These compressive-PCA guarantees hold even under *model mismatch*: although Σ_k constrains $\tilde{\mathbf{R}}$ to have rank at-most- k , the empirical autocorrelation matrix $\hat{\mathbf{R}}$ used to construct $\tilde{\mathbf{z}}$ tends to have full rank d in practice. To explore this idea in more detail, suppose for the moment that the data $\mathcal{X}$ lies in a rank- k subspace. In this case, $\hat{\mathbf{R}}$ would have rank at-most- k , and there exists a symmetric $\mathbf{R}$ with rank at-most- k that drives the cost $C(\mathbf{R} | \tilde{\mathbf{z}}) = \|\mathbf{A} \text{vec}(\hat{\mathbf{R}}) - \mathbf{A} \text{vec}(\mathbf{R})\|^2$ to zero. Furthermore, when $\mathbf{A}$ satisfies the RIP (32), the left inequality in (32) implies that this cost-minimizing $\tilde{\mathbf{R}}$ must equal $\hat{\mathbf{R}}$. When the data $\mathcal{X}$ does not lie in a rank- k subspace, the minimal cost will be non-zero. In this case, an upper bound on the error $\|\hat{\mathbf{R}} - \tilde{\mathbf{R}}\|_F$ of the cost-minimizing estimate $\tilde{\mathbf{R}}$, as well as an upper bound on the excess risk of the principal subspace $\tilde{\theta}$ of $\tilde{\mathbf{R}}$, can both be obtained as a consequence of the RIP (32).

In place of the RIP, compressive learning uses the so-called “*lower* RIP” (LRIP),

$$d(p_{\theta}, p_{\theta'}) \leq C_0 \|\mathcal{A}(p_{\theta}) - \mathcal{A}(p_{\theta'})\|, \quad (33)$$

assumed to be valid for each pair $p_{\theta}, p_{\theta'} \in \Sigma_{\Theta}$, where C_0 is a positive constant. The LRIP says that the Euclidean distance $\|\mathcal{A}(p_{\theta}) - \mathcal{A}(p_{\theta'})\|$ between the (noiseless) sketches of two distributions is—up to a scaling—controlling the distance $d(p_{\theta}, p_{\theta'})$ between the distributions themselves (26). Note that the lower bound $d(\cdot, \cdot)$ in (33) is task-specific, and not Euclidean as in the CS case (5).

The first main theoretical guarantee about compressive learning [5, 6] is that, when the LRIP (33) holds, the estimate $\tilde{\theta}$ obtained by minimizing (31) automatically has controlled excess risk. In particular, using $\tilde{\theta} = \arg \min_{\theta} R^*(\theta | p_{\tilde{\theta}})$ [5, 6], the excess risk bound (28) can be combined with the LRIP (33) and the definition of $\tilde{\theta}$ (31) as follows:

$$\begin{aligned} R^*(\tilde{\theta} | p_X) - R^*(\theta^* | p_X) &\leq 2d(p_{\tilde{\theta}}, p_X) \\ &\leq 2 \left[d(p_{\tilde{\theta}}, p_{\theta'}) + d(p_{\theta'}, p_X) \right] \\ &\leq 2 \left[C_0 \|\mathcal{A}(p_{\tilde{\theta}}) - \mathcal{A}(p_{\theta'})\| + d(p_{\theta'}, p_X) \right] \\ &\leq 2 \left[C_0 \|\mathcal{A}(p_{\tilde{\theta}}) - \tilde{z}\| + C_0 \|\tilde{z} - \mathcal{A}(p_{\theta'})\| + d(p_{\theta'}, p_X) \right] \\ &\leq 2 \left[2C_0 \|\mathcal{A}(p_{\theta'}) - \tilde{z}\| + d(p_{\theta'}, p_X) \right] \\ &\leq 2 \left[2C_0 \|\mathcal{A}(p_X) - \tilde{z}\| + 2C_0 \|\mathcal{A}(p_{\theta'}) - \mathcal{A}(p_X)\| + d(p_{\theta'}, p_X) \right], \end{aligned} \quad (34)$$

for any distribution $p_{\theta'}$ in Σ_{Θ} (because $d(\cdot, \cdot)$ is a valid distance metric, *i.e.*, $d(p_1, p_2) \leq d(p_1, p_3) + d(p_3, p_2) \forall p_3$). One option is to choose $p_{\theta'} = p_{\theta^*}$, in which case the term $2C_0 \|\mathcal{A}(p_{\theta^*}) - \mathcal{A}(p_X)\| + d(p_{\theta^*}, p_X)$ in the upper bound reflects the excess risk due to modeling and $2C_0 \|\mathcal{A}(p_X) - \tilde{z}\|$ reflects the excess risk due to sketching from a finite dataset $\mathcal{X}$. For a tighter bound, we could choose $p_{\theta'}$ as the distribution in Σ_{Θ} that minimizes the right side of (34). These approaches, and more refined variants, have been applied to analyze various learning tasks, in *e.g.*, [5, 6, 13, 14].

It should be emphasized that recovery guarantees based on the RIP (5) or LRIP (33) hold *even in the presence of measurement noise e and/or modeling errors*. In CS, measurement noise arises due to, *e.g.*, thermal noise or interference, while modeling errors arise when x_0 is not truly k -sparse. Many sparse recovery techniques are provably robust to such noise and modeling errors, *cf.* (6) in Box 4. In compressive learning, measurement noise arises due to the finite cardinality of the dataset $\mathcal{X}$ (recall (9)), while modeling errors arise when $p_X \notin \Sigma_{\Theta}$. For example, GMM-recovery guarantees can be established even when p_X is not truly a GMM.

5.5 ESTABLISHING THE LRIP VIA KERNELS AND THE JL LEMMA

In light of the fact that the LRIP (33) yields statistical learning guarantees (34) for compressive learning, a key question becomes: How can we choose the sketch dimension m so that the LRIP (33) holds? Similar to how the RIP is proven in CS (see, *e.g.*, [4, Lemma 9.33]), refinements of a mathematical tool called the Johnson-Lindenstrauss (JL) lemma [34] can be used to obtain a value of m sufficient for the LRIP to hold *with high probability on the draw of the random feature map $\Phi(\cdot)$* . We now summarize this approach.

To begin, we establish connections to kernel methods (see Box 7 for a brief review of kernels). Some key observations are that any feature map explicitly defines a positive definite *kernel*, and that the *expectation* of this kernel defines another useful positive definite kernel. To see why, consider, for example, the RF feature map (3). First, for a fixed set of frequency vectors $\{\mathbf{w}_j\}_{j=1}^m$, we have

$$\left\langle \frac{1}{\sqrt{m}} \Phi(\mathbf{x}), \frac{1}{\sqrt{m}} \Phi(\mathbf{x}') \right\rangle = \frac{1}{m} \sum_{j=1}^m \exp(-j2\pi \mathbf{w}_j^\top (\mathbf{x} - \mathbf{x}')), \quad \forall \mathbf{x}, \mathbf{x}' \quad (35)$$

where $\langle \cdot, \cdot \rangle$ is the standard Euclidean inner product in $\mathbb{R}^m$ or $\mathbb{C}^m$. The left-hand side defines a positive definite kernel according to the terminology reviewed in Box 7. This kernel is also “shift-invariant,” in that it depends only on the difference $\mathbf{x} - \mathbf{x}'$. Next, imagine that the d -dimensional frequency vectors $\mathbf{w}_j$,

which constitute the m rows of the random matrix $\mathbf{W}$, are drawn i.i.d. from some probability distribution p_W . In the limit of large sketch dimension m , the law of large numbers combined with (35) says

$$\left\langle \frac{1}{\sqrt{m}} \Phi(\mathbf{x}), \frac{1}{\sqrt{m}} \Phi(\mathbf{x}') \right\rangle \xrightarrow{w.p.1} \mathbb{E}_W \exp(-j2\pi \mathbf{W}^\top (\mathbf{x} - \mathbf{x}')). \quad (36)$$

The right hand side of (36) is the Fourier transform (FT) of the probability density p_W evaluated at $\mathbf{x} - \mathbf{x}'$, i.e., the characteristic function $\Psi_{p_W}(\mathbf{x}' - \mathbf{x})$ using the notation from (14). Because $\Psi_{p_W}(\cdot)$ is the FT of a non-negative function, it is a so-called *positive definite function*, which means that, for any $\{\mathbf{x}_i\}_{i=1}^n$, the $n \times n$ matrix Ψ defined with elements $\Psi_{ij} = \Psi_{p_W}(\mathbf{x}_i - \mathbf{x}_j)$ for $1 \leq i, j \leq n$ will be positive semi-definite. Thus, if we construct a kernel as $\kappa(\mathbf{x}, \mathbf{x}') := \Psi_{p_W}(\mathbf{x}' - \mathbf{x})$, then it will be a positive definite kernel. When $p_W = \mathcal{N}(\mathbf{0}, \sigma_w^2 \mathbf{I}_d)$, this approach yields the familiar Gaussian kernel (a particular type of “radial basis function”), i.e., $\kappa_\sigma(\mathbf{x}, \mathbf{x}') := \exp(-\|\mathbf{x} - \mathbf{x}'\|^2 / 2\sigma^2)$, here of width $\sigma = 1/\sigma_w$.

More generally, by considering any parametric feature map of the form $\Phi(\mathbf{x} | \mathbf{W})$, where the parameter $\mathbf{W}$ is drawn at random according to some probability distribution, one can define³ the “expected kernel”

$$\kappa(\mathbf{x}, \mathbf{x}') := \mathbb{E}_{\mathbf{W}} \left\langle \frac{1}{\sqrt{m}} \Phi(\mathbf{x} | \mathbf{W}), \frac{1}{\sqrt{m}} \Phi(\mathbf{x}' | \mathbf{W}) \right\rangle, \quad (37)$$

not to be confused with the “mean kernel” defined in (38). This setting includes RF features (3) with i.i.d. frequencies w_j (as above), or with a frequency matrix $\mathbf{W}$ that includes structured blocks of rows, as we will soon discuss.

[Box 7] Kernel Methods and Kernel Embeddings of Probability Distributions

Sketching shares connections with *kernel methods* [3], a family of machine learning techniques that produce decisions or insights using a kernel function, $\kappa(\mathbf{x}, \mathbf{x}') \in \mathbb{R}$, which measures the “similarity” between $\mathbf{x}$ and $\mathbf{x}'$. A kernel is said to be “positive definite” if the $n \times n$ matrix $\mathbf{K}$, constructed with entries $\kappa(\mathbf{x}_i, \mathbf{x}_j)$ for $1 \leq i, j \leq n$, is positive *semi*-definite for every possible $\{\mathbf{x}_i\}_{i=1}^n$. The celebrated “kernel trick” states that any positive definite kernel *implicitly* amounts to an inner product in some higher-dimensional (and potentially infinite-dimensional) feature space $\mathcal{H}$, and vice-versa. That is, $\kappa(\mathbf{x}, \mathbf{x}') = \langle \Phi(\mathbf{x}), \Phi(\mathbf{x}') \rangle$ for some (not necessarily explicitly known) mapping $\Phi(\cdot)$ from the signal space to $\mathcal{H}$. Any machine learning method that relies only on the evaluation of inner products—such as ridge regression, support-vector-machine classification, PCA [32], and dictionary learning [35]—can be “kernelized” by using a kernel in place of the inner product. Kernelizing a method is thus tantamount to applying that method in a transformed, higher-dimensional feature space. In this way, more complex estimation and/or decision functions can be implemented.

Given a positive definite kernel $\kappa(\mathbf{x}, \mathbf{x}')$ operating on signals $\mathbf{x}$ and $\mathbf{x}'$ in some set, it is possible to “lift” $\kappa(\cdot, \cdot)$ to a positive definite kernel *operating on probability distributions* over this set by defining the so-called *mean kernel*

$$k(p, q) := \mathbb{E}_{X \sim p, X' \sim q} [\kappa(X, X')]. \quad (38)$$

This defines an embedding of probability distributions into a kernel space, which is analogous to the finite-dimensional embedding $\mathcal{A}(p)$ of probability distribution p from (8). The *maximum mean discrepancy* (MMD) [36, 37]

$$\text{MMD}(p, q) := \sqrt{k(p, p) + k(q, q) - 2k(p, q)} \quad (39)$$

is the Euclidean metric naturally induced by the mean kernel. It is analogous to the Euclidean distance between sketches, $\|\mathcal{A}(p) - \mathcal{A}(q)\|$. The MMD, originally introduced in the context of two-sample hypothesis testing [36], is now well-known in machine learning. When the MMD behaves as a true metric, i.e., when $\text{MMD}(p, q) = 0 \Leftrightarrow p = q$, the mean kernel $k(\cdot, \cdot)$ is said to be “characteristic”. In $\mathbb{R}^d$, many classic kernels $\kappa(\cdot, \cdot)$, such as the Gaussian and Laplace kernels, yield characteristic mean kernels [37].

Now that the kernel connections have been established, we return to our original objective of *understanding when the LRIP (33) holds*. Importantly, the law of large numbers shows that, in the limit of large sketch dimension m , the right side of (33) is related to the maximum mean discrepancy (MMD) from

³In many papers, the $1/\sqrt{m}$ scaling in (35)-(37) is subsumed in the feature map $\Phi(\cdot)$.

(39), which is a kernel-based distance between distributions (see Box 7). Indeed, for arbitrary probability distributions p and q , we have

$$\lim_{m \rightarrow \infty} \frac{1}{\sqrt{m}} \|\mathcal{A}(p) - \mathcal{A}(q)\| \stackrel{\text{w.p.1}}{=} \text{MMD}(p, q), \quad (40)$$

where the maximum mean discrepancy is implicitly the one obtained from the kernel used to build the sketching operator $\mathcal{A}$. Thus, when the LRIP (33) holds, it must also be true that, for sufficiently large sketch dimension m ,

$$d(p_\theta, p_{\theta'}) \leq C'_0 \text{MMD}(p_\theta, p_{\theta'}), \quad \text{for all } p_\theta, p_{\theta'} \in \Sigma_\Theta, \quad (41)$$

where C'_0 is a positive constant. Property (41) connects two different metrics on probability distributions: the left hand side of (41) is defined by the learning task (recall (26)), while the right hand side of (41) is defined by the expected kernel κ , from (37), associated with the randomized feature map. Note that (41) is a deterministic property; it does not depend on the draw of the randomized feature map, unlike the LRIP (33). In the literature, (41) is called the “kernel-LRIP” [5, 6] because it is a kernel-based analog to the LRIP. Being deterministic, the expected kernel is often easier to manipulate than the random feature map, and thus eases the proof of the kernel-LRIP. This is important because, if the kernel-LRIP (41) holds, then, using arguments based on the JL lemma, one can also establish [5, 6] the LRIP (33), as we show below.

The JL lemma [34] is a precursor of the restricted isometry property that is specialized to finite sets (5). It states that, given N arbitrary d -dimensional vectors $\mathbf{x}_i$, there exists an $m \times d$ matrix $\mathbf{A}$, with m on the order of $\log N$, such that $\|\mathbf{A}\mathbf{x}_i - \mathbf{A}\mathbf{x}_j\| \approx \|\mathbf{x}_i - \mathbf{x}_j\|$ for all i, j .

To illustrate how the JL lemma is useful in the context of compressive learning, let us momentarily restrict our attention to a learning problem where the collection Σ_Θ of parametric distributions is of *finite cardinality*, noting that we can extend this approach to continuous families of parametric distributions through discretization arguments involving the notion of covering numbers [4, Appendix C], as described below. As a concrete example, let us consider a discretized variant of compressive GMM using the RF feature map (3). As in [5, 6, 8], we assume that the d -dimensional frequency vectors $\mathbf{w}_j$ are drawn i.i.d. from the normal distribution $\mathcal{N}(\mathbf{0}, \sigma_w^2 \mathbf{I}_d)$, and we consider learning a mixture of Gaussian components $p_{\theta_\ell} = \mathcal{N}(\boldsymbol{\mu}_\ell, \mathbf{I})$ for $\ell = 1, \dots, k$, with equal weights $\alpha_\ell = 1/k$ for each ℓ , where the means $\boldsymbol{\mu}_\ell$ are assumed to be bounded, separated, and discretized on a regular grid. In this setting, there exists a finite number, N , of possible parametric mixture distributions $p_\theta \in \Sigma_\Theta$. As long as the MMD is a true metric (as defined in Box 7), the ratio $d(p_\theta, p_{\theta'}) / \text{MMD}(p_\theta, p_{\theta'})$ is finite for $p_\theta \neq p_{\theta'}$, hence whenever Σ_Θ is a finite collection, as in this concrete example, the kernel-LRIP (41) holds (for a sufficiently large C'_0). Moreover, specializing to an arbitrary pair of mixtures $p_\theta, p_{\theta'} \in \Sigma_\Theta$ and a given sketch dimension m , refinements of (40) (using *measure concentration*) enable one to show that the lower bound

$$\frac{1}{\sqrt{2}} \text{MMD}(p_\theta, p_{\theta'}) \leq \frac{1}{\sqrt{m}} \|\mathcal{A}(p_\theta) - \mathcal{A}(p_{\theta'})\| \quad (42)$$

holds with probability at least $1 - \exp(-c_0 m)$, where c_0 is a positive *concentration constant*. Since there are only N^2 pairs of parametric mixtures $p_\theta, p_{\theta'} \in \Sigma_\Theta$, the lower bound (42) is valid *uniformly* for all of them, except with probability at most $N^2 \exp(-c_0 m)$. This failure probability can be made smaller than any $\epsilon > 0$ by choosing m larger than $2c_0^{-1} \log(N/\epsilon)$. Combining (41) and (42) using a union bound, it can be shown [5, 6] that the LRIP (33) holds with high probability when C_0 is on the order of $C'_0 \sqrt{m}$ and the sketch dimension m grows logarithmically with N , the number of parametric distributions in the finite set Σ_Θ .

For infinite collections Σ_Θ , proving the kernel-LRIP (41), and eventually the LRIP (33), is more technical and can require some additional assumptions. As an example, for compressive clustering with RF features, it can be proven that there is a constant C'_0 such that (41) holds, provided that the centroids are sufficiently separated and bounded [5, 6]. The JL lemma can be extended by refining techniques used to establish the RIP (5). The main idea is to first prove the LRIP on some finite collection $\Sigma' \subset \Sigma_\Theta$ of N probability distributions, and then to extrapolate it (with slightly worse constants) to Σ_Θ . Technically, this involves the notion of covering numbers, and the cardinality $N = |\Sigma'|$ is typically exponential in the number of parameters needed to describe Σ_Θ . For example, for GMM, one needs kd parameters to describe the means $\boldsymbol{\mu}_\ell \in \mathbb{R}^d$, $1 \leq \ell \leq k$ of the mixture $p_\theta = \sum_{\ell=1}^k \frac{1}{k} \mathcal{N}(\boldsymbol{\mu}_\ell, \mathbf{I})$, and $\log N$ essentially depends linearly on kd .

The dimension m of the sketch for which the LRIP holds with high probability is thus on the order of kd/c_0 , up to some additional factors due to the proof technique [5, 6].

Empirical studies of compressive clustering [11, 12] and compressive GMM [8] suggest that a sketch dimension m on the order of kd (the number of parameters in these settings) is sufficient to yield accurate learning performance. The best known bounds on provably good sketch dimensions [5, 6] remain pessimistic compared to these empirically validated sketch dimensions. This is most likely related to sub-optimal bounds for the concentration constant c_0 and/or shortcomings in the techniques used to extend the LRIP from a finite collection Σ' to an infinite collection Σ_Θ .

5.6 THE CHALLENGE OF DESIGNING A FEATURE MAP GIVEN A LEARNING TASK

In the existing literature, the LRIP has been established [5, 6] for randomized feature maps $\Phi(\cdot)$ (e.g., random Fourier features, random quadratic features) that mimic related constructions from compressive sensing, developed either for sparse-vector recovery or low-rank matrix recovery.

When sketching with random Fourier features (e.g., for compressive clustering and compressive GMM), the main design choice for $\Phi(\cdot)$ is the distribution from which to draw the random frequencies w_j (i.e., the rows of $\mathbf{W}$ in (3)). In light of the connections to shift-invariant kernels (recall (35)), this design task is a particular instance of the difficult problem of kernel design [3, Sec. 4.4.5].

For example, when the rows of $\mathbf{W}$ are drawn i.i.d. $\mathcal{N}(\mathbf{0}, \sigma_w^2 \mathbf{I})$, the choice of the variance σ_w^2 determines the choice of the width $\sigma = 1/\sigma_w$ of the corresponding Gaussian kernel $\kappa_\sigma(\cdot, \cdot)$. Indeed, from a signal-processing standpoint, the corresponding mean kernel $k_\sigma(\cdot, \cdot)$ (recall (38)) acts to low-pass filter the underlying data distributions. To see why, observe that $\kappa_\sigma(\mathbf{x}, \mathbf{x}') = g_\sigma(\mathbf{x} - \mathbf{x}')$ with $g_\sigma(\mathbf{x}) := e^{-\|\mathbf{x}\|^2/2\sigma^2}$, and so

$$k_\sigma(p, q) = \mathbb{E}_{X \sim p, X' \sim q}[\kappa_\sigma(X, X')] = \iint e^{-\|X - X'\|^2/2\sigma^2} dp(X) dq(X') \quad (43)$$

$$= \langle g_\sigma \star p, q \rangle_{L^2} = \langle g_\sigma \star p, g_{\bar{\sigma}} \star q \rangle_{L^2}, \quad (44)$$

where $\star$ denotes convolution and $\bar{\sigma} = \sigma/\sqrt{2}$. Hence, the associated MMD (39) satisfies $\text{MMD}(p, q) = \|g_{\bar{\sigma}} \star p - g_{\bar{\sigma}} \star q\|_{L^2}$. Recall that learning-from-a-sketch is often performed by minimizing a “sketch matching” cost $\|\tilde{z} - \mathcal{A}(p_\theta)\|^2 = \|\mathcal{A}(\hat{p}_\mathcal{X}) - \mathcal{A}(p_\theta)\|^2$, as in (30), where $\hat{p}_\mathcal{X}$ denotes the empirical distribution of the data $\mathcal{X}$. In the limit of large sketch dimension m , this cost compares the smoothed versions of the probability distributions $\hat{p}_\mathcal{X}$ and p_θ , since $\frac{1}{m} \|\tilde{z} - \mathcal{A}(p_\theta)\|^2 \approx \|g_{\bar{\sigma}} \star \hat{p}_\mathcal{X} - g_{\bar{\sigma}} \star p_\theta\|_{L_2}^2$. Similarly, when using a greedy algorithm to learn a mixture model (or cluster centroids) from a sketch, the normalized inner product $\frac{1}{m} \langle \tilde{z}, \mathcal{A}(p_{\theta_i}) \rangle$ approximates the correlation $\langle g_{\bar{\sigma}} \star \hat{p}_\mathcal{X}, g_{\bar{\sigma}} \star p_{\theta_i} \rangle_{L_2}$ between the low-passed versions of the empirical data distribution $\hat{p}_\mathcal{X}$ and the candidate mixture component p_{θ_i} , respectively.

This latter idea is illustrated for compressive clustering in Figure 11. There, since the mixture component associated to a candidate centroid $\mathbf{c}$ is the Dirac $p_\mathbf{c}(\mathbf{x}) = \delta(\mathbf{x} - \mathbf{c})$ (recall the discussion before (19)), we have that $\langle \tilde{z}_\sigma, \mathcal{A}_\sigma(p_\mathbf{c}) \rangle = \langle \tilde{z}_\sigma, \Phi_\sigma(\mathbf{c}) \rangle$, where the dependence on the kernel width $\sigma = 1/\sigma_w$ has been made explicit. Meanwhile,

$$\langle g_{\bar{\sigma}} \star \hat{p}_\mathcal{X}, g_{\bar{\sigma}} \star p_\mathbf{c} \rangle_{L_2} = (g_\sigma \star \hat{p}_\mathcal{X})(\mathbf{c}) = \frac{1}{n} \sum_{i=1}^n g_\sigma(\mathbf{c} - \mathbf{x}_i). \quad (45)$$

In (45), we recognize a Parzen window density estimator [38], whose computation for a given $\mathbf{c}$ requires access to the n training samples $\mathbf{x}_i$. In contrast, its surrogate $\langle \tilde{z}_\sigma, \Phi_\sigma(\mathbf{c}) \rangle$ only requires access to the m -dimensional sketch $\tilde{z}$. In a large-scale setting, this can save huge amounts of memory and computation. As can be seen when comparing the top and bottom rows in Figure 11, $\langle \tilde{z}_\sigma, \Phi_\sigma(\mathbf{c}) \rangle$ well approximates $(g_\sigma \star \hat{p}_\mathcal{X})(\mathbf{c})$ for sufficiently large kernel width σ . By comparing the different columns of Figure 11, it can also be seen that the kernel width σ should be chosen compatible with the cluster width and separation. This choice involves a tradeoff between smoothing the unwanted gaps between data samples and over-smoothing the (desired) gaps between clusters. Existing theory [5, 6] identifies sufficient conditions on the choice of σ related to the number k of candidate clusters and their minimum separation.

Although i.i.d. Gaussian frequencies w_j were used with the RF map above, one may also consider the use of i.i.d. non-Gaussian frequencies. Such designs, as proposed in [8], can yield improved empirical behavior. As we will see in the next section, it is also possible to deviate from the RF map with the goal

Fig. 11: Criterion $\langle \tilde{z}_\sigma, \Phi_\sigma(c) \rangle$ used by greedy parameter estimation algorithms in compressive clustering with random Fourier features (top row) versus its expected value $(g_\sigma \star \hat{p}_\mathcal{X})(c)$ (bottom row) as a function of the centroid hypothesis location $c = [c_1, c_2]^\top$. The dataset (in blue) consists of $n = 100$ points drawn according to a mixture of $k = 3$ isotropic Gaussians. The frequencies w_j used to define the feature map $\Phi_\sigma(\cdot) = \exp(-j2\pi W \cdot)$ are drawn according to a standard Gaussian $\mathcal{N}(\mathbf{0}, \sigma_w^2 \mathbf{I})$ with $\sigma_w = 1/\sigma$. The sketch $\tilde{z}_\sigma$ is computed with the feature map $\Phi_\sigma(\cdot)$. With $\sigma^2 = \frac{1}{500}$ (left), there is insufficient smoothing, and the criterion displays many spurious local maxima; with $\sigma^2 = \frac{1}{10}$ (middle), there is appropriate smoothing, and local maxima are in good correspondence with the true cluster centers; with $\sigma^2 = 1$ (right), there is over-smoothing, and the criterion displays only a single maximum.

of improving computational efficiency. Such constructions also yield non-Gaussian expected kernels (37). Although they work well in practice, there is currently no proof that these latter kernels satisfy the kernel-LRIP.

While existing theory focuses on proving the LRIP for a given random feature map and learning task, an important open question is: How should one *design the feature map* to best match a given learning task? In particular, can we design a random feature map that satisfies the LRIP (33) for a given learning task defined by a loss function $L(\theta | x_i)$ and embodied by a task-driven distance (26)? A promising yet still challenging avenue would be to first identify a positive definite kernel $\kappa_0(x, x')$ for which the corresponding MMD satisfies the kernel-LRIP (41), and then use Bochner’s theorem or Mercer’s theorem (see Box 8) to design a random feature map $\Phi(\cdot)$ whose expected kernel (37) is precisely κ_0 .

[Box 8] Approximating kernel methods with random feature maps

While each random feature map $\Phi(\cdot)$ implicitly defines a positive definite kernel by taking the expectation (37), the converse is also true.

For shift-invariant kernels, *i.e.*, kernels for which $\kappa(x, x') = \kappa(x - x', 0)$ only depends on the difference $x - x'$ (such as the Gaussian kernel), this is a consequence of Bochner’s theorem [9], which states that if κ is a positive definite kernel such that $\kappa(x, x) = 1$ for all x , then its Fourier transform yields a probability distribution $p_W(w) = \int \kappa(x, 0) \exp(-j2\pi w^\top x) dx$. Conversely, the kernel can be obtained by the inverse Fourier transform, which can also be phrased as an expectation:

$$\kappa(x, x') = \int \exp(j2\pi w^\top (x - x')) p_W(w) dw = \mathbb{E}_{W \sim p_W} \exp(j2\pi W^\top (x - x')).$$

Hence, drawing i.i.d. frequency vectors w_j according to p_W yields an RF feature map (3) whose expected kernel is precisely κ . For instance, a Laplace kernel can be approximated if the rows of W are drawn i.i.d. from the Cauchy distribution [39].

More generally, under mild assumptions on a positive definite kernel κ , one can invoke Mercer’s

theorem [40] to similarly show the existence of a random feature map whose expected kernel, in the sense of (37), matches κ .

6 SKETCHING WITH REDUCED COMPUTATIONAL RESOURCES

The computational cost of sketching via (1) is heavily dependent on the feature map $\Phi(\cdot)$. The computational cost of parameter estimation via (17) is also heavily dependent on $\Phi(\cdot)$, since it often involves iterative application of $\Phi(\cdot)$.

Often, the feature map is constructed as a randomized linear operation followed by a componentwise non-linear operation, *i.e.*,

$$\Phi(\mathbf{x}) = \varrho(\mathbf{W}\mathbf{x}), \quad (46)$$

where $\mathbf{W}$ is a (randomly drawn) matrix of size $m \times d$ and $\varrho(\cdot)$ applies a scalar non-linear function identically to each element of the vector $\mathbf{W}\mathbf{x}$. For example, the feature maps described earlier for compressive PCA, GMM, and clustering all have this form. Reducing the computational cost of each stage has been the goal of several studies. For example, using a fast transform for $\mathbf{W}$ drastically reduces the memory and computational complexity demands relative to an explicit matrix. Also, quantized versions of the nonlinearity $\varrho(\cdot)$ are much more easily implemented in hardware than, say, the complex exponential nonlinearity $\varrho_{\text{RF}}(\cdot) := \exp(-j2\pi \cdot)$ used in the RF map (3). We discuss such constructions of $\mathbf{W}$ and $\varrho(\cdot)$ below.

6.1 SKETCHING WITH STRUCTURED RANDOM MATRICES

In most of our previous examples, we constructed $\mathbf{W}$ by drawing its m rows i.i.d. from the normal distribution $\mathcal{N}(\mathbf{0}, \sigma_w^2 \mathbf{I}_d)$ with some variance $\sigma_w^2 > 0$. Recall that, with the RF map $\varrho_{\text{RF}}(\cdot)$, the rows of $\mathbf{W}$ correspond to the (d -dimensional) frequencies used when sampling the Fourier transform of the (empirical) data distribution. In any case, when $\mathbf{W}$ is an explicit matrix, the computational complexity of computing $\Phi(\mathbf{x})$ of the form (46) is dominated by the matrix-vector product $\mathbf{W}\mathbf{x}$. Thus, it is on the order of md , which is also the order of the memory needed to store $\mathbf{W}$.

As an alternative to these approaches, it has been suggested to construct $\mathbf{W}$ as a structured random matrix with a fast implementation that mimics an i.i.d. Gaussian matrix. Multiple ways of accomplishing this goal have been proposed in the literature. We focus on the approach suggested in [41], which was successfully applied to compressive learning in [42]. There, the idea is to construct $\mathbf{W}$ as a vertical concatenation of $b = \lceil m/d \rceil$ blocks $\{\mathbf{B}_j\}_{j=1}^b$, each of size $d \times d$. These blocks have the form $\mathbf{B}_j = \mathbf{D}_j^{(0)} \mathbf{H} \mathbf{D}_j^{(1)} \mathbf{H} \mathbf{D}_j^{(2)} \mathbf{H} \mathbf{D}_j^{(3)}$, where $\mathbf{H}$ is the Walsh-Hadamard matrix and $\mathbf{D}_j^{(k)}$ are random diagonal matrices. In particular, the diagonal elements of $\mathbf{D}_j^{(1)}$, $\mathbf{D}_j^{(2)}$, and $\mathbf{D}_j^{(3)}$ are drawn i.i.d. from the uniform distribution over $\{-1, 1\}$, and the diagonal elements of $\mathbf{D}_j^{(0)}$ are drawn i.i.d. from the χ distribution with d degrees-of-freedom, which is the distribution of the norm of a d -variate Gaussian vector. This construction is depicted in Fig. 12. The fast Walsh-Hadamard transform offers an order $d \log d$ -complexity implementation of the matrix-vector multiplication $\mathbf{H}\mathbf{x}$ and prevents the need to explicitly store $\mathbf{H}$. With this structured and fast incarnation of $\mathbf{W}$, the sketching complexity shrinks from order md to order $m \log d$. Moreover, since only the diagonal matrices need to be stored, the storage cost shrinks from order md to order m .

6.2 SKETCHING WITH QUANTIZED CONTRIBUTIONS

With the RF map (3), which is commonly used in compressive clustering and GMM, the non-linear operation $\varrho(\cdot)$ in (46) becomes $\varrho_{\text{RF}}(\cdot) := \exp(-j2\pi \cdot)$. Since implementing $\varrho_{\text{RF}}(\cdot)$ with high accuracy is somewhat costly and not amenable to easy hardware acceleration, one might consider quantizing it. For example, dropping the imaginary part of $\varrho_{\text{RF}}(\cdot)$ and quantizing the real part to one bit of precision yields the 2π -periodic function $\varrho_q(\cdot) := \text{sign}(\cos(2\pi \cdot))$ which is simply a “square wave”.

To alleviate the effects of quantization, one can apply dithering [10] to the input. In this case, the feature map becomes $\Phi_q(\mathbf{x}) = \varrho_q(\mathbf{W}\mathbf{x} + \boldsymbol{\xi}) \in \{-1, 1\}^m$, with i.i.d. dither components ξ_j drawn uniformly over $[0, 1)$. The effect of dithering is to make the quantized Φ_q behave similarly to non-quantized Φ_{RF} on average. For instance, it was shown in [43] that for each $\mathbf{W}$, $\mathbf{x}$, $\mathbf{x}'$, and $\boldsymbol{\xi}$,

$$\langle \Phi_q(\mathbf{x}), \Phi_{\text{RF}}(\mathbf{x}') \rangle \approx \mathbb{E}_{\boldsymbol{\xi}} \langle \Phi_q(\mathbf{x}), \Phi_{\text{RF}}(\mathbf{x}') \rangle = c \langle \Phi_{\text{RF}}(\mathbf{x}), \Phi_{\text{RF}}(\mathbf{x}') \rangle, \quad (47)$$

Fig. 12: Structured random matrix design from [41]. Each block is a composition of several Hadamard and diagonal matrices. For convenience, we draw W^T instead of W .

Fig. 13: Criterion $\langle \tilde{z}, \Phi(c) \rangle$ versus centroid hypothesis $c = [c_1, c_2]^T$ used by greedy parameter estimation algorithms in compressive clustering with $k = 3$ and a 2D dataset (in blue), with and without quantization of the sketch $\tilde{z}$ and/or feature map Φ . The surface plot on the left highlights the irregularity of $c \mapsto \langle \tilde{z}_q, \Phi_q(c) \rangle$.

where c is a constant. The approximation above is accurate for a typical draw of the m -dimensional dither vector ξ when the sketch dimension m is large enough [44]. Note that, when this dithered quantizer is used to compute a sketch $\tilde{z}_q := \frac{1}{n} \sum_{i=1}^n \Phi_q(x_i) = \frac{1}{n} \sum_{i=1}^n \varrho_q(Wx_i + \xi)$, it is important to use the same dither realization ξ for all samples i .

The question then arises: when estimating parameters θ via (31), how should we account for quantization in the sketch? Simply replacing the $\mathcal{A}(p_{\theta_\ell})$ term with $\mathcal{A}_q(p_{\theta_\ell}) := \mathbb{E}_{X \sim p_{\theta_\ell}}[\Phi_q(X)]$ may sound appealing. For example, with compressive GMM (18) or compressive clustering (20), this would mean minimizing $\|\tilde{z}_q - \sum_{\ell=1}^k \alpha_\ell \mathcal{A}_q(p_{\theta_\ell})\|^2$. However, the optimization problem (31) would become more challenging, as suggested by comparing the left two panels to the right two panels in Fig. 13, which plots the centroid-selection criterion $\langle \tilde{z}, \Phi(c) \rangle = \langle \tilde{z}, \mathcal{A}(p_c) \rangle$ used by greedy algorithms. Also, this approach would not inherit the theoretical guarantees that were carefully established using the LRIP (33), which does not easily translate to the quantized case.

Instead, we suggest to use $\mathcal{A}_{RF}(p_{\theta_\ell}) := \mathbb{E}_{X \sim p_{\theta_\ell}}[\Phi_{RF}(X)]$ for the $\mathcal{A}(p_{\theta_\ell})$ term in (18), and to re-scale $\tilde{z}'_q = \tilde{z}_q/c$ with the constant c from (47). Indeed, thanks to (47), the resulting cost function $C(\theta|\tilde{z}'_q) = \|\tilde{z}'_q - \sum_{\ell=1}^k \alpha_\ell \mathcal{A}_{RF}(p_{\theta_\ell})\|^2$ is, in expectation over ξ , *exactly the same* (up a constant additive bias term) as the non-quantized cost function $C(\theta|\tilde{z}_{RF})$ [43]. The similarity between $C(\theta|\tilde{z}'_q)$ and $C(\theta|\tilde{z}_{RF})$, even for a single realization of the m -dimensional dither vector ξ , is suggested by comparing the right two panels in Fig. 13.

Empirical results [43] suggest that, when the sketch dimension is inflated by about 25%, this quantized compressive learning procedure yields the same performance as the non-quantized procedure. Moreover, accurate probabilistic bounds for approximation (47), established in [44], allow one to extend the theoretical compressive learning guarantees in (34) to this new cost function.

[Box 9] Potential of Sketching for Privacy-Aware Learning and Alternative Approaches

Given a dataset, privacy preservation can be achieved by asking a trusted dataholder to corrupt all queries of the dataset [45, 46] in a controlled manner. There are, however, challenges to this so-called *interactive* approach. For example, because the privacy-preserving effects of this corruption can often be diminished through the mining of multiple query responses (especially if the queries are adaptive), the per-query corruption levels must be designed with the type and total number of queries in mind. These corruption levels are often designed using a so-called “privacy budget,” which is expended over multiple queries in order to meet an overall privacy level. Once the entire privacy budget has been used up, the data can no longer be accessed by a given datauser. Also, the dataholder must ensure that responses to *different* datausers cannot be combined in a way that circumvents the intended privacy preservation.

In contrast, the *non-interactive* approach [45, 46] is to publish an intermediate privacy-preserving synopsis of the dataset, to which the public is allowed unlimited access. For example, with a low-dimensional dataset, one could publish a privacy-preserving histogram of the data [47], from which aggregate statistics could be subsequently extracted. The non-interactive approach is attractive for several reasons. For example, there is no need to formulate nor allocate a privacy budget; it is sufficient to set an overall privacy level. Also, there is no need to worry about datausers sharing/combining data.

By adding noise to a sketch of the form (1), one can easily generate a privacy preserving synopsis of a dataset. By construction, such sketches capture the global statistics of the dataset $\mathcal{X} = \{x_i\}_{i=1}^n$ while being relatively insensitive to each individual data sample x_i , especially when the sample cardinality n is large. Also, when the original data is distributed across multiple devices, a privacy preserving global sketch can be constructed by first locally sketching at each device and then averaging those local sketches at a fusion center, as illustrated in Box 1. In this scenario, the local sketches will themselves be privacy preserving, which alleviates concerns about privacy leaks during data fusion.

7 PRIVACY PRESERVATION

In addition to its efficient use of computational resources, sketching is a promising tool for privacy-preserving machine learning. In numerous applications, such as when working with medical records, online surveys, or measurements coming from personal devices, data samples contain sensitive personal information and data providers ask that individuals’ contributions to the dataset remain private, *i.e.*, not publicly discoverable. Learning from such data collections while protecting the privacy of individual contributors has become a crucial challenge [45, 46, 48].

A common way to preserve privacy is to have a trusted dataholder (or “curator”) corrupt the response to each query of the dataset [45] in a controlled manner. A query may ask for something as simple as counting the number of times a given event occurred, or it may ask for more sophisticated information that requires the dataholder to run an inference algorithm. As the corruption becomes more significant, the privacy guarantee gets stronger, but the quality of the response to the query (called the “utility”) degrades. This can be conceptualized by a *privacy-utility tradeoff* [45, 48].

When sketching a dataset via (1), there is a very simple way to preserve privacy in any subsequent learning task: simply add i.i.d. noise (with appropriate distribution) to the sketch. The privacy level can be adjusted by changing the variance of that noise, as described below. This one-time approach to privacy preservation (more generally known as *privacy-preserving data publishing* [45]) has several benefits over the query-based approach to privacy preservation discussed in the previous paragraph. (See also Box 9.)

7.1 SKETCHING WITH DIFFERENTIAL PRIVACY GUARANTEES

Differential privacy [46] is a standard framework for privacy preservation that has a precise mathematical definition and is well-known in machine learning and signal processing (*e.g.*, [48]). When a given (randomized) learning pipeline is differentially private, its output depends negligibly on the presence or absence of any individual sample in the dataset. Differential privacy is robust to many forms of attack, such as when the adversary can access side information that nullifies privacy guarantees based on anonymization or mutual information measures (*e.g.*, when the adversary can control some of the data vectors x_i , or can access additional databases that are correlated with the primary database).

Fig. 14: When sketching with differential privacy, the output log-density of the sketch s remains close when changing one sample in the dataset (since (49) is equivalent to $|\log(p_{\tilde{s}(\mathcal{X})}(s)) - \log(p_{\tilde{s}(\mathcal{X}')} (s))| \leq \epsilon$ for all possible s). An adversary with knowledge of $\Phi(\cdot)$ and s —as symbolized by the red arrows—could then hardly decide whether a given sample x_j was used to compute the sketch or not.

For compressive learning methods, enforcing differential privacy guarantees is as simple as adding well-calibrated noise v to the usual sketch $\tilde{z}$, *i.e.*, constructing

$$\tilde{s}(\mathcal{X}) := \tilde{z}(\mathcal{X}) + v, \quad (48)$$

where we find it helpful to explicitly denote the dependence of the dataset $\mathcal{X}$. We will assume that the realization $\Phi(\cdot)$ of the random feature map is fixed and publicly known, in contrast to other approaches like [49, 50] that use linear mixing matrices as encryption keys to ensure privacy preservation. As a result, when we treat $\tilde{s}(\mathcal{X})$ as random, this is due to the randomness in v , not the randomness in $\Phi(\cdot)$ or $\mathcal{X}$.

Formally, the sketching mechanism $\tilde{s}(\cdot)$ is said to be ϵ -differentially private [46] if, for any dataset $\mathcal{X} = \{x_i\}_{i=1}^n$ and “neighboring” dataset $\mathcal{X}' = \mathcal{X} \setminus \{x_j\} \cup \{x'_j\}$ that replaces the individual sample x_j by another sample x'_j , and for any possible sketch outcome s , we have that

$$\exp(-\epsilon) \leq \frac{p_{\tilde{s}(\mathcal{X})}(s)}{p_{\tilde{s}(\mathcal{X}')} (s)} \leq \exp(\epsilon). \quad (49)$$

Here, $\epsilon > 0$ plays the role of a privacy level: smaller ϵ implies a stronger privacy guarantee. In words, (49) says that, when ϵ is small, the densities of $\tilde{s}(\mathcal{X})$ and $\tilde{s}(\mathcal{X}')$ are almost indistinguishable, as depicted in Figure 14.

The condition (49) can be interpreted as bounding a “likelihood ratio” [51], a familiar quantity in signal processing. Consider two hypotheses: one that the dataset equals $\mathcal{X}$ (*i.e.*, includes x_j), and the other that the dataset equals $\mathcal{X}'$ (*i.e.*, includes x'_j instead). Then $p_{\tilde{s}(\mathcal{X})}(s)$ would be the likelihood of observing private sketch s under the first hypothesis, while $p_{\tilde{s}(\mathcal{X}')} (s)$ would be the same for the second hypothesis. Say an adversary wanted to detect whether or not the dataset contains x_j . By appropriately thresholding the likelihood ratio $p_{\tilde{s}(\mathcal{X})}(s)/p_{\tilde{s}(\mathcal{X}')} (s)$, one can obtain hypothesis tests that are optimal from various perspectives (*e.g.*, Bayes, minimax, Neyman-Pearson) [51, Ch. 2]. Thus, when (49) holds with small ϵ , it is fundamentally difficult for an adversary to determine whether x_j or x'_j was present in the sketch. Even if the adversary had non-trivial prior knowledge of the true hypothesis (as in so-called “linkage attacks,” which make use of a second public dataset to which the target user contributed), (49) implies that—for any method—the probability of recovering the true hypothesis from the sketch is only slightly higher than that which is achievable *without* observing the sketch.

To ensure that the noisy sketch (48) is differentially private, it is sufficient to draw the noise v as i.i.d. Laplacian with appropriate variance. The variance needed to achieve a given privacy level ϵ can be determined by analyzing the so-called *sensitivity* of the noiseless sketch, *i.e.*, the biggest possible change that

can result from removing one sample. When using the random Fourier feature map (3), which generates complex-valued $\tilde{z}(\mathcal{X})$, it has been established [52, 53] that it is sufficient for the real and imaginary components of $\mathbf{v}$ to be i.i.d. Laplacian with standard deviation $\sigma_v \propto \frac{m}{n\epsilon}$.

A weaker form of privacy, known as *approximate differential privacy* or (ϵ, δ) -*differential privacy* [46], can be attained by adding Gaussian noise $\mathbf{v}$ with smaller variance. For example, with RF features, it is sufficient for the real and imaginary components of $\mathbf{v}$ to be i.i.d. Gaussian with standard deviation $\sigma_v \propto \frac{\sqrt{m}}{n\epsilon}$.

The privacy-utility tradeoff facilitates the comparison of different privacy-preserving learning strategies. For example, given two strategies, one could match the privacy levels ϵ and compare utilities, or one could match utilities and compare privacy levels. For compressive learning, the utility of interest is the risk (recall (25)).

In this section, we focused on differential privacy. Other definitions of privacy exist in the literature, such as information-theoretic ones [67]. Likewise, cryptographic methods, such as fully homomorphic encryption [68], can be used to transmit and manipulate data in a secure and private manner, although this notion of “privacy” is quite different from the former ones. Additional work is needed to understand whether compressive learning is “private” according to definitions other than differential privacy.

8 PERSPECTIVES AND OPEN CHALLENGES

By averaging well-chosen randomized feature transformations over large training collections, sketching significantly compresses data in a way that facilitates provably accurate yet scalable learning from huge and/or streaming datasets, while simultaneously preserving privacy.

In this article, we described several approaches to accelerate the sketching process, including feature quantization and the use of randomized fast transforms. Another approach is to randomly mask each feature vector $\Phi(\mathbf{x}_i)$ prior to averaging, *i.e.*, set a random subset of its components to zero. It has been established [53] that such random masking does not increase nor decrease the differential privacy level ϵ . But it does reduce the need to compute all entries of each feature vector, and thus reduces sketching complexity. Another promising approach consists of mixed analog-digital sketches, where, *e.g.*, optical processing units are used to significantly improve the energy efficiency of the linear stage [54].

When discussing methods *to learn* from a sketch, we focused on optimization-based approaches. Although heuristics based on orthogonal matching pursuit [8] and approximate message passing [12] have been proposed that yield promising empirical results, performance guarantees for these approaches have yet to be established. Alternatives, such as total-variation minimization over the space of signed measures [25, 28, 55], principled greedy methods [56], and gradient flows on systems of particles [57] could be leveraged to make progress on this front, and black-box optimization could be used to learn from sketches computed by optical processing units.

As for *applications* of compressive learning, most of the current literature, and hence most of our article, has focused on unsupervised learning tasks. Further work is needed to develop compressive learning methods for supervised tasks like regression and classification [33]. For example, one approach to compressive classification was proposed in [58]: for each class $\ell = 1, \dots, k$, one computes a sketch $\tilde{z}_\ell$ using only the training examples with label ℓ (*i.e.*, we sketch the k *conditional distributions* of the data). From those sketches, one could estimate the conditional densities of each class (using a mixture model, for example), from which a maximum likelihood (ML) or maximum a posteriori (MAP) classifier could be derived. Another approach is to perform classification directly in the compressed domain: to an unseen example $\mathbf{x}'$, we would assign the class ℓ that maximizes the correlation $\langle \tilde{z}_\ell, \Phi(\mathbf{x}') \rangle$. This strategy can be interpreted [58] as compressively evaluating a Parzen-window classifier [38]. Further work is also needed on unsupervised matrix-factorization tasks like dictionary learning, low-rank matrix completion, and non-negative matrix factorization (NMF) [59].

ACKNOWLEDGMENTS

The authors are grateful to Luc Giffon and the anonymous IEEE-SPM reviewers for their comments on the manuscript. The authors are funded in part by the Fonds de la Recherche Scientifique (FNRS), the FNRD PDR grant T.0136.20 (Learn2Sense), the Agence Nationale de la Recherche (ANR) under grant ANR-19-CHIA-0009 (AllegroAssai), and the National Science Foundation (NSF) under grant CCF-1955587.

AUTHORS

Rémi Gribonval (M'02, SM'08, FM'14) is a Senior Researcher with Inria at ENS de Lyon, France. His research, at the interface of mathematical signal processing, machine learning, and approximation theory, focuses on the versatile notion of sparsity and its applications in these fields with an emphasis on dimension reduction, inverse problems and high-dimensional statistics. He founded the series of international workshops SPARS on Signal Processing with Adaptive/Sparse Representations. He received the 2011 Blaise Pascal Award in Applied Mathematics and Scientific Engineering from the SMAI by the French National Academy of Sciences. He is a fellow of the IEEE and of EURASIP. ORCID: 0000-0002-9450-8125, remi.gribonval@inria.fr

Antoine Chatalic received the B.S. degree in 2013, the M.S. degree in computer science in 2017, and the Ph.D. degree in signal processing in 2020, all from the University of Rennes 1, France. He is now working as a postdoctoral researcher at the machine learning Genoa center (Genoa, Italy) under the supervision of Lorenzo Rosasco. His research interests include sketching and dimensionality reduction techniques, kernel methods, compressive sensing, online algorithms and privacy-aware learning. ORCID: <https://orcid.org/0000-0003-2574-2417>, antoine.chatalic@dibris.unige.it

Nicolas Keriven is a Researcher at CNRS and GIPSA-lab (Grenoble). He received his Master's degree from École polytechnique and École Normale Supérieure de Cachan in 2014, and his Ph.D. from Université de Rennes 1 in 2017, under the supervision of Rémi Gribonval. He was previously a Postdoctoral Researcher at École Normale Supérieure de Paris, funded by the CFM-ENS Laplace chair *Modèles et Sciences des données*. He received the Best Student Paper award at the SPARS workshop in 2017 for his work on *Random Moments for Sketched Mixture Learning*. ORCID: 0000-0002-3846-8763, Nicolas.Keriven@cnrs.fr

Vincent Schellekens is a FNRS Research Fellow, doing his PhD in signal processing at the ICTEAM institute in UCLouvain (Louvain-la-Neuve, Belgium). He received his B.Sc. (2015) and M.Sc. (2017) in electrical engineering from UCLouvain. His research, under the supervision of Laurent Jacques, covers compressive sensing, general random embeddings, and privacy-preservation methods for large-scale learning. ORCID: 0000-0001-5874-5732, vincent.schellekens@uclouvain.be

Laurent Jacques (B.Sc., M.Sc., & PhD in Math. Physics, '96, '98, & '04) is currently a Professor and a FNRS Research Associate since '12 with the Image and Signal Processing Group of ICTEAM/UCLouvain in Louvain-La-Neuve, Belgium. His research focuses on sparse signal representations, quantized compressive sensing theory, and computational imaging. Since '99, Laurent Jacques has co-authored more than 40 papers in international journals, 80 conference proceedings and presentations in signal and image processing conferences, and 4 book chapters. ORCID: 0000-0002-6261-0328, laurent.jacques@uclouvain.be

Philip Schniter (S'92-M'93-SM'05-F'14) received the Ph.D. degree in electrical engineering from Cornell University, Ithaca, NY, USA, in 2000. He is currently a Professor in the Department of Electrical and Computer Engineering at The Ohio State University, Columbus, OH, USA. He received the IEEE Signal Processing Society (SPS) Best Paper Award in 2016, and has served on several IEEE SPS Technical Committees: Signal Processing for Communications and Networking, Sensor Array and Multichannel, and Computational Imaging. ORCID: 0000-0003-0939-7545, schniter.1@osu.edu

REFERENCES

- [1] R. Gribonval, A. Chatalic, N. Keriven, V. Schellekens, L. Jacques, and P. Schniter. "Sketching Datasets for Large-Scale Learning (long version)". In: *arXiv:2008.01839 [cs, math, stat]* (June 2021). arXiv: 2008.01839.
- [2] F. Pérez-Cruz and O. Bousquet. "Kernel methods and their potential use in signal processing". In: *IEEE Signal Process. Mag.* 21.3 (2004).
- [3] B. Schölkopf and A. Smola. *Learning with Kernels*. Adaptive Computation and Machine Learning. Cambridge: MIT Press, 2002.
- [4] S. Foucart and H. Rauhut. *A Mathematical Introduction to Compressive Sensing*. Springer, May 2012.
- [5] R. Gribonval, G. Blanchard, N. Keriven, and Y. Traonmilin. "Compressive Statistical Learning with Random Feature Moments". en. In: *Mathematical Statistics and Learning* (2021).
- [6] R. Gribonval, G. Blanchard, N. Keriven, and Y. Traonmilin. "Statistical learning guarantees for compressive clustering and compressive mixture modeling". In: *Mathematical Statistics and Learning* 3.2 (Aug. 2021).
- [7] D. A. Reynolds, T. F. Quatieri, and R. B. Dunn. "Speaker Verification Using Adapted Gaussian Mixture Models". In: *Dig. Sig. Proc.* 10.1-3 (2000).
- [8] N. Keriven, A. Bourrier, R. Gribonval, and P. Pérez. "Sketching for large-scale learning of mixture models". In: *Inform. Inference* 7.3 (2017).
- [9] A. Rahimi and B. Recht. "Random Features for Large Scale Kernel Machines". In: *Proc. Neural Inform. Process. Syst. Conf.* (2007).

- [10] R. M. Gray and D. L. Neuhoff. “Quantization”. In: *IEEE Trans. Inform. Theory* 44.6 (1998).
- [11] N. Keriven, N. Tremblay, Y. Traonmilin, and R. Gribonval. “Compressive K-means”. In: *Proc. IEEE Int. Conf. Acoust. Speech & Signal Process.* 2017.
- [12] E. Byrne, A. Chatalic, R. Gribonval, and P. Schniter. “Sketched clustering via hybrid approximate message passing”. In: *IEEE Trans. Signal Process.* 67.17 (2019).
- [13] M. P. Sheehan, M. S. Kotzagiannidis, and M. E. Davies. “Compressive Independent Component Analysis”. In: *2019 27th European Signal Processing Conference (EUSIPCO)*. IEEE, 2019.
- [14] M. P. Sheehan, A. Gonon, and M. E. Davies. “Compressive Learning for Semi-Parametric Models”. In: *arXiv:1910.10024* (2019).
- [15] G. Cormode, M. Garofalakis, P. J. Haas, and C. Jermaine. “Synopsis for Massive Data”. In: *Foundations and Trends in Databases* 4.1 (2011).
- [16] D. Achlioptas. “Database-friendly random projections: Johnson-Lindenstrauss with binary coins”. In: *Journal of Computer and System Sciences* 66.4 (2003).
- [17] M. W. Mahoney. “Randomized algorithms for matrices and data”. In: *Foundations and Trends in Machine Learning* 3 (2010).
- [18] C. Boutsidis, A. Zouzias, and P. Drineas. “Random Projections for k-means Clustering”. In: *Advances in Neural Information Processing Systems*. 2010.
- [19] D. P. Woodruff. “Sketching as a tool for numerical linear algebra”. In: *Foundations and Trends in Theoretical Computer Science* 10 (2014).
- [20] D. Achlioptas, F. McSherry, and B. Schölkopf. “Sampling techniques for kernel methods”. In: *Advances in Neural Information Processing Systems*. 2002.
- [21] D. Feldman, M. Monemizadeh, C. Sohler, and D. P. Woodruff. “Coresets and sketches for high dimensional subspace approximation problems”. In: *Proceedings of the Annual ACM-SIAM Symposium on Discrete Algorithms* 1 (2010).
- [22] C. Williams and M. W. Seeger. “Using the Nystrom Method to Speed Up Kernel Machines”. In: *Advances in Neural Information Processing Systems*. Vol. 13. 2001.
- [23] A. R. Hall. *Generalized Method of Moments*. Oxford University Press, 2005.
- [24] E. J. Candès and M. B. Wakin. “An introduction to compressive sampling”. In: *IEEE Signal Process. Mag.* 25.2 (2008).
- [25] E. J. Candès and C. Fernandez-Granda. “Towards a mathematical theory of super-resolution”. In: *Commun. Pure & Appl. Math.* 67.6 (2014).
- [26] M. Lustig, D. L. Donoho, J. M. Santos, and J. M. Pauly. “Compressed sensing MRI”. In: *IEEE Signal Process. Mag.* 25.2 (2008).
- [27] J. A. Fessler. “Optimization Methods for Magnetic Resonance Image Reconstruction: Key Models and Optimization Algorithms”. In: *IEEE Signal Processing Magazine* 37.1 (2020).
- [28] Q. Denoyelle, V. Duval, G. Peyré, and E. Soubies. “The sliding Frank-Wolfe algorithm and its application to super-resolution microscopy”. English. In: *Inverse problems* 36.1 (Jan. 2020).
- [29] C. Poon, N. Keriven, and G. Peyré. “The geometry of off-the-grid compressed sensing”. In: *arXiv:1802.08464* (2020). arXiv: 1802.08464.
- [30] Y. Chen and Y. Chi. “Harnessing Structures in Big Data via Guaranteed Low-Rank Matrix Estimation: Recent Theory and Fast Algorithms via Convex and Nonconvex Optimization”. In: *IEEE Signal Processing Magazine* 35.4 (2018).
- [31] R. Ge, J. D. Lee, and T. Ma. “Matrix Completion has No Spurious Local Minimum”. In: *Advances in Neural Information Processing Systems* 29. Ed. by D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett. Curran Associates, Inc., 2016.
- [32] V. Vapnik. *The Nature of Statistical Learning Theory*. Springer, 2013.
- [33] S. Shalev-Shwartz and S. Ben-David. *Understanding Machine Learning*. Cambridge: Cambridge University Press, 2009.
- [34] W. B. Johnson and J. Lindenstrauss. “Extensions of Lipschitz mappings into a Hilbert space”. In: *Contemporary mathematics* 26.189-206 (1984).
- [35] R. Rubinfeld, A. M. Bruckstein, and M. Elad. “Dictionaries for Sparse Representation Modeling”. In: *Proceedings of the IEEE* 98.6 (2010).
- [36] A. Gretton, K. M. Borgwardt, M. J. Rasch, B. Schölkopf, and A. J. Smola. “A Kernel Method for the Two-Sample Problem”. In: *Advances in Neural Information Processing Systems (NIPS)*. 2007. eprint: 0805.2368.
- [37] B. K. Sriperumbudur, A. Gretton, K. Fukumizu, B. Schölkopf, and G. R. Lanckriet. “Hilbert space embeddings and metrics on probability measures”. In: *The Journal of Machine Learning Research* 11 (2010).
- [38] Z.-W. Pan, D.-H. Xiang, Q.-W. Xiao, and D.-X. Zhou. “Parzen windows for multi-class classification”. In: *Journal of complexity* 24.5-6 (2008).
- [39] P. T. Boufounos, S. Rane, and H. Mansour. “Representation and coding of signal geometry”. In: *Inform. Inference* 6.4 (2017).
- [40] F. Bach. “On the Equivalence between Kernel Quadrature Rules and Random Feature Expansions”. In: *Journal of Machine Learning Research* (2017).
- [41] F. X. Yu, A. T. Suresh, K. M. Choromanski, D. N. Holtmann-Rice, and S. Kumar. “Orthogonal Random Features”. In: *Proc. Neural Inform. Process. Syst. Conf.* 2016.
- [42] A. Chatalic, R. Gribonval, and N. Keriven. “Large-Scale High-Dimensional Clustering with Fast Sketching”. In: *Proc. IEEE Int. Conf. Acoust. Speech & Signal Process.* 2018.
- [43] V. Schellekens and L. Jacques. “Quantized Compressive K-Means”. In: *IEEE Signal Process. Lett.* 25.8 (Aug. 2018).
- [44] V. Schellekens and L. Jacques. “Breaking the waves: asymmetric random periodic features for low-bitrate kernel machines”. In: *Information and Inference: A Journal of the IMA* (May 2021). arXiv: 2004.06560.
- [45] B. C. M. Fung, K. Wang, R. Chen, and P. S. Yu. “Privacy-preserving Data Publishing: A Survey of Recent Developments”. In: *ACM Comput. Surv.* 42.4 (June 2010).
- [46] C. Dwork and A. Roth. “The Algorithmic Foundations of Differential Privacy”. In: *Theoretical Computer Science* 9.3-4 (2014).
- [47] W. Qardaji, W. Yang, and N. Li. “Differentially private grids for geospatial data”. In: *2013 IEEE 29th international conference on data engineering (ICDE)*. IEEE, 2013.
- [48] A. D. Sarwate and K. Chaudhuri. “Signal processing and machine learning with differential privacy: Algorithms and challenges for continuous data”. In: *IEEE Signal Process. Mag.* 30.5 (2013).
- [49] M. Testa, D. Valsesia, T. Bianchi, and E. Magli. *Compressed Sensing for Privacy-Preserving Data Processing*. Springer, 2019.

- [50] S. Rane and P. T. Boufounos. "Privacy-preserving nearest neighbor methods: Comparing signals without revealing them". In: *IEEE Signal Process. Mag.* 30.2 (2013).
- [51] H. V. Poor. *An Introduction to Signal Detection and Estimation*. Springer Science & Business Media, 2013.
- [52] V. Schellekens, A. Chatalic, F. Houssiau, Y.-A. de Montjoye, L. Jacques, and R. Gribonval. "Differentially Private Compressive k-means". In: *Proc. IEEE Int. Conf. Acoust. Speech & Signal Process.* 2019.
- [53] A. Chatalic, V. Schellekens, F. Houssiau, Y. A. de Montjoye, L. Jacques, and R. Gribonval. "Compressive learning with privacy guarantees". In: *Information and Inference: A Journal of the IMA* iaab005 (May 2021).
- [54] A. Saade, F. Caltagirone, I. Carron, L. Daudet, A. Drémeau, S. Gigan, and F. Krzakala. "Random projections through multiple optical scattering". In: *Proc. IEEE Int. Conf. Acoust. Speech & Signal Process.* 2016.
- [55] N. Boyd, G. Schiebinger, and B. Recht. "The Alternating Descent Conditional Gradient Method for Sparse Inverse Problems". In: *SIAM J. Optim.* 27.2 (2017).
- [56] Y. Traonmilin and J.-F. Aujol. "The basins of attraction of the global minimizers of the non-convex sparse spike estimation problem". en. In: *Inverse Problems* 36.4 (Feb. 2020). Publisher: IOP Publishing.
- [57] L. Chizat and F. Bach. "On the Global Convergence of Gradient Descent for Over-parameterized Models using Optimal Transport". In: *Advances in Neural Information Processing System*. 2018.
- [58] V. Schellekens and L. Jacques. "Compressive Classification (Machine Learning without learning)". In: *arXiv:1812.01410* (2018).
- [59] M. Udell, C. Horn, R. Zadeh, and S. Boyd. "Generalized Low Rank Models". In: *Found. Trends Mach. Learn.* 9.1 (2016).