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

Sketching Data Sets for Large-Scale Learning

Keeping only what you need

Big data can be a blessing: with very large training data sets 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 data sets 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 data set is massively compressed before learning. Here, a "sketch" is first constructed by computing carefully chosen nonlinear random features [e.g., random Fourier (RF) features] and averaging them over the whole data set. Parameters are then learned from the sketch, without access to the original data set. 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].

Digital Object Identifier 10.1109/MSP.2021.3092574
Date of current version: 27 August 2021

Introduction to compressive learning

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

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

Next, during the learning phase, an estimate $\tilde{\boldsymbol{\theta}}$ of some essential statistical parameters $\boldsymbol{\theta}$ of the data set are extracted from

the sketch $\tilde{z}$. These parameters are application dependent. For example, as illustrated in Figure 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 in the following. Python notebooks that illustrate examples from the article are available at https://gitlab.inria.fr/SketchedLearning/spm-notebook.

The transformation $\boldsymbol{\Phi}(\cdot)$, known as the *feature map*, is generally nonlinear and randomized. Although the use of $\boldsymbol{\Phi}(\cdot)$ is related to "kernel methods" in machine learning (e.g., [2] and [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 get 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 data set. This enables one to handle massive data sets while keeping computational resources under control.
3) Sketching can preserve privacy: the transformation $\boldsymbol{\Phi}(\cdot)$ 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 data sets is a simple mean of the sketches of those data sets.

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 Figure 2 and Table 1 for an illustration.

Principal component analysis

Principal component analysis (PCA) seeks to find the linear subspace of a fixed dimension $k < d$ that best fits the d-dimensional data $\{\boldsymbol{x}_i\}_{i=1}^n$ in the least-squares (LS) sense. In this case, we assume the data to be centered so that the target parameters $\boldsymbol{\theta}$ can be described by a k-dimensional orthonormal basis $\{\boldsymbol{u}_\ell\}_{\ell=1}^k$ that maximizes $\Sigma_{\ell=1}^k \Sigma_{i=1}^n \left| \boldsymbol{u}_\ell^\top \boldsymbol{x}_i \right|^2$.

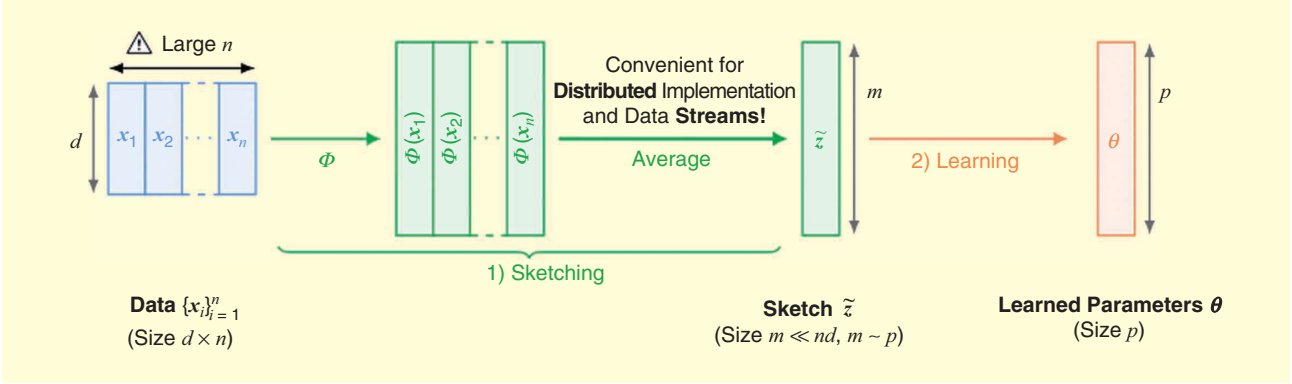

FIGURE 1. An overview of sketching and parameter learning.

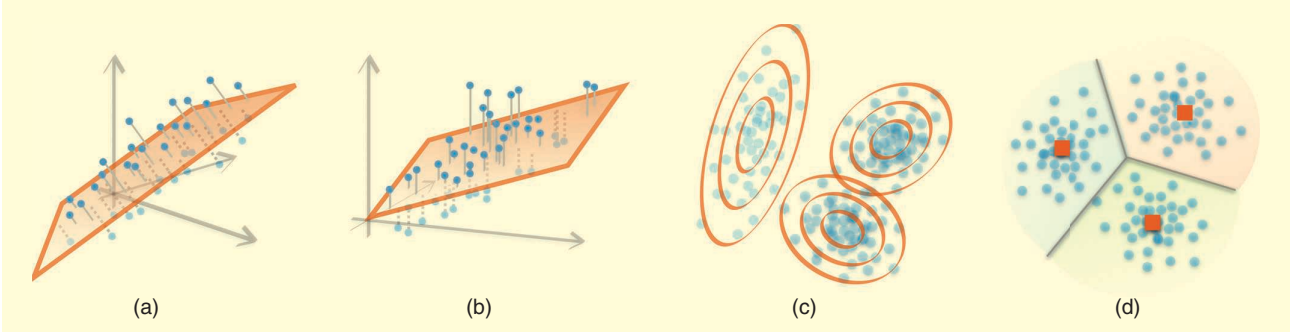

FIGURE 2. A schematic representation of the four running examples covered by this article. In the four figures, each $\boldsymbol{x}_i$ is associated with a blue colored disk; hence, the training collection corresponds to a point cloud, and the orange color geometrically represents the learned parameters.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

It is well known that one solution is given by the k principal eigenvectors of the empirical autocorrelation matrix $\hat{\boldsymbol{R}} = (1/n)\sum_{i=1}^{n} \boldsymbol{x}_i \boldsymbol{x}_i^\top$. This $\hat{\boldsymbol{R}}$ can be interpreted as a sketch of the form (1) that uses the feature map $\boldsymbol{\Phi}(\boldsymbol{x}) = \mathrm{vec}(\boldsymbol{x}\boldsymbol{x}^\top) = (x_1 x_1, x_2 x_1, \ldots, x_d x_d)^\top$ of dimension $m = d^2$. Here and in the sequel, the $\mathrm{vec}(\cdot)$ operator vectorizes a matrix by stacking its columns. As soon as the cardinality n of the data set exceeds the dimension d, the dimension of this sketch is smaller than nd, the total size of the data set. 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 $\boldsymbol{\Phi}(\boldsymbol{x}) = ((\boldsymbol{w}_1^\top \boldsymbol{x})^2, \ldots, (\boldsymbol{w}_m^\top \boldsymbol{x})^2)^\top$, where the d-dimensional vectors $\boldsymbol{w}_j$, $1 \le j \le m$ are the rows of a tall random matrix $\boldsymbol{W}$ 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 $\hat{\boldsymbol{R}}$ with provable accuracy [5], [51].

LS linear regression

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

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

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

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

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 (GMM) $p(\boldsymbol{x}|\boldsymbol{\theta}) = \sum_{\ell=1}^{k} \alpha_\ell \mathcal{N}(\boldsymbol{x}; \boldsymbol{\mu}_\ell, \boldsymbol{\Sigma}_\ell)$ to the data $\{\boldsymbol{x}_i\}_{i=1}^n$. 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 $\hat{\boldsymbol{\theta}} = \mathrm{argmax}_{\boldsymbol{\theta} \in \boldsymbol{\Theta}} \sum_{i=1}^{n} \log p(\boldsymbol{x}_i|\boldsymbol{\theta})$, but the EM algorithm processes all n data samples $\{\boldsymbol{x}_i\}_{i=1}^n$ at each iteration, which can be computationally burdensome when n is very large.

An alternative [6] is to compute a sketch $\tilde{\boldsymbol{z}}$ of the form (1) using RF features [7]:

$$\boldsymbol{\Phi}(\boldsymbol{x}) = \exp(-\mathrm{j}2\pi \boldsymbol{W}\boldsymbol{x}), \tag{3}$$

where $\mathrm{j} = \sqrt{-1}$ and $\boldsymbol{W} \in \mathbb{R}^{m \times d}$ 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 (for example, the connection between the RF map and the Gaussian kernel is discussed in the sidebar "Approximating Kernel Methods With Random Feature Maps"). 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 [6], where it was shown that 1,000 h of speech can be compressed down to a few kilobytes without loss of verification performance (see the sidebar "Compressive Clustering of Modified National Institute of Standards and Technology Digits").

k-means clustering

The goal of clustering is to group together "similar" data samples from $\{\boldsymbol{x}_i\}_{i=1}^n$. In the k-means approach to clustering, one aims to find the set of k centroids $\{\boldsymbol{c}_\ell\}_{\ell=1}^k$ that minimizes the average squared distance from each sample to its nearest centroid; i.e., $\sum_{i=1}^{n} \min_\ell \left\| \boldsymbol{x}_i - \boldsymbol{c}_\ell \right\|^2$. The famous Lloyd algorithm [8] 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 data set using (1) and extract the centroids from the sketch [9], [10]. For this purpose, one could use RF features (3) and (as explained in the sequel) solve for the centroids $\{\boldsymbol{c}_\ell\}_{\ell=1}^k$ and the (nonnegative, sum-to-one) weights $\{\alpha_\ell\}_{\ell=1}^k$ that minimize

Table 1. The notations used in this article.

Notation	
$\mathcal{X} = \{\boldsymbol{x}_i\}_{i=1}^n$	Data set of n training samples $\boldsymbol{x}_i \in \mathbb{R}^d$
$\tilde{\boldsymbol{z}}$	(Empirical) sketch of the data set, a vector of dimension m
$\boldsymbol{\Phi}(\cdot)$	Sketching feature map, a function mapping $\mathbb{R}^d$ to either $\mathbb{R}^m$ or $\mathbb{C}^m$
$\boldsymbol{\theta} \in \boldsymbol{\Theta}$	Target parameters to learn, of dimension p (e.g., principal component analysis matrix, centroids, and mixture model)
$\boldsymbol{W}$	Random matrix associated with certain feature maps
$\boldsymbol{\Phi}(\cdot)$	Usually drawn i.i.d. Gaussian
$\varrho(\cdot)$	Componentwise nonlinearity associated with certain feature maps $\boldsymbol{\Phi}(\cdot)$
w.h.p.	With high probability, i.e., with exponentially decaying failure probability
i.i.d.	Independent identically distributed
w.p.1	With probability 1
$\left\| \cdot \right\|$, $\langle \cdot, \cdot \rangle$	Euclidean norm of a vector, inner product between vectors
$\left\| \cdot \right\|_0$	ℓ^0 "norm" of a vector, i.e., its number of nonzero entries
$\left\| \cdot \right\|_F$	Frobenius norm of a matrix
$\left\| \cdot \right\|_{L^2}, \langle \cdot, \cdot \rangle_{L^2}$	L^2 norm of a function, L^2 inner product between functions.

$\left\| \tilde{z} - \Sigma_{\ell=1}^{k} \alpha_\ell \boldsymbol{\Phi}(\boldsymbol{c}_\ell) \right\|$. In this setting, the weights allow us to model unbalanced clusters, yet only the centroids need to be recovered. On large data sets, 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 $\{\boldsymbol{c}_\ell\}_{\ell=1}^{k}$ [9], [10]. For example, this method allows us to cluster the Modified National Institute of Standards and Technology (MNIST) digit data set, of dimension $d = 784$ and cardinality $n = 70{,}000$, using a complex-valued sketch of dimension $m = 400$ (see the sidebar "Compressive Clustering of MNIST Digits").

Although this article focuses on these examples, the general compressive learning framework extends beyond them. It has been applied to, for example, independent component analysis [11] (by sketching the cumulant tensor) and subspace clustering [12] (using a polynomial feature map). Extending the framework to a new task raises essential questions, such as: How do we choose the feature map $\boldsymbol{\Phi}(\cdot)$? How do we learn the essential parameters $\boldsymbol{\theta}$ 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 article sketches answers to these questions for the worked examples introduced in the preceding.

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) [13]. The goal of AQP is to build a short description of the content of a massive data set, 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 a 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 [13, Sec. 5.3.1] synopsis. In this context, the statistical parameters $\boldsymbol{\theta}$ learned from a sketch are interpreted as the result of a particular query on the data set.

Despite the apparent similarities between AQP and compressive learning, both data types 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.

Sketches for streaming and distributed methods

In AQP, a popular class of synopses is that of linear sketches [13, Ch. 5]. A linear sketch, which maps a data set to a vector, must satisfy the single condition that the sketch of the concatenation of two data sets equals the sum of their sketches. In mathematical terms, if we denote by $\tilde{z}(\mathcal{X})$ the sketch of a data set $\mathcal{X} = \{\boldsymbol{x}_i\}_{i=1}^{n}$, 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 $\tilde{z}(\mathcal{X}) = \Sigma_{i=1}^{n} \boldsymbol{\Phi}(\boldsymbol{x}_i)$ for some possibly nonlinear feature map $\boldsymbol{\Phi}(\cdot)$. 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,

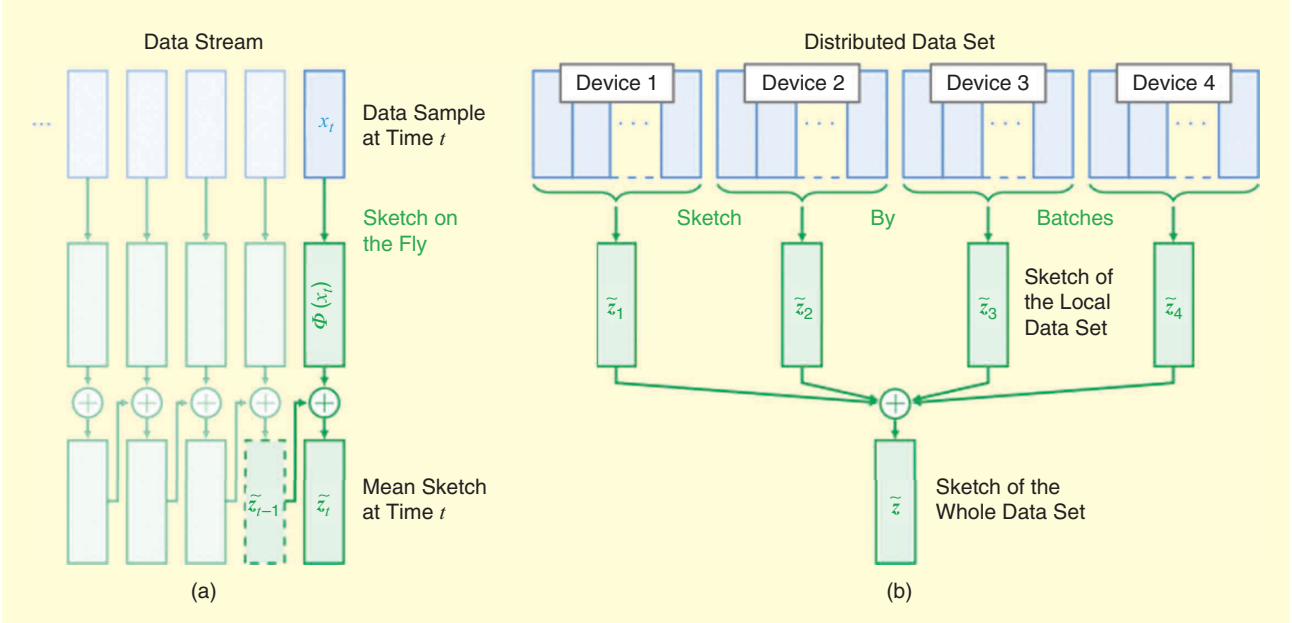

FIGURE 3. (a) A streaming scenario, where the data samples are sketched one by one and the mean sketch is updated at each time. (b) A distributed scenario, where each device computes a local sketch and a centralized entity further averages these local sketches.

inserting or deleting an element $\boldsymbol{x}_i$ from the data set corresponds to adding or subtracting $\boldsymbol{\Phi}(\boldsymbol{x}_i)$ from the sketch, respectively (see Figure 3). 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 [13, Ch. 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.

In the sections 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

Compressive Mixture Modeling for Speaker Verification

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 Gaussian mixture model (GMM)–universal background model (UBM) [S1]. 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

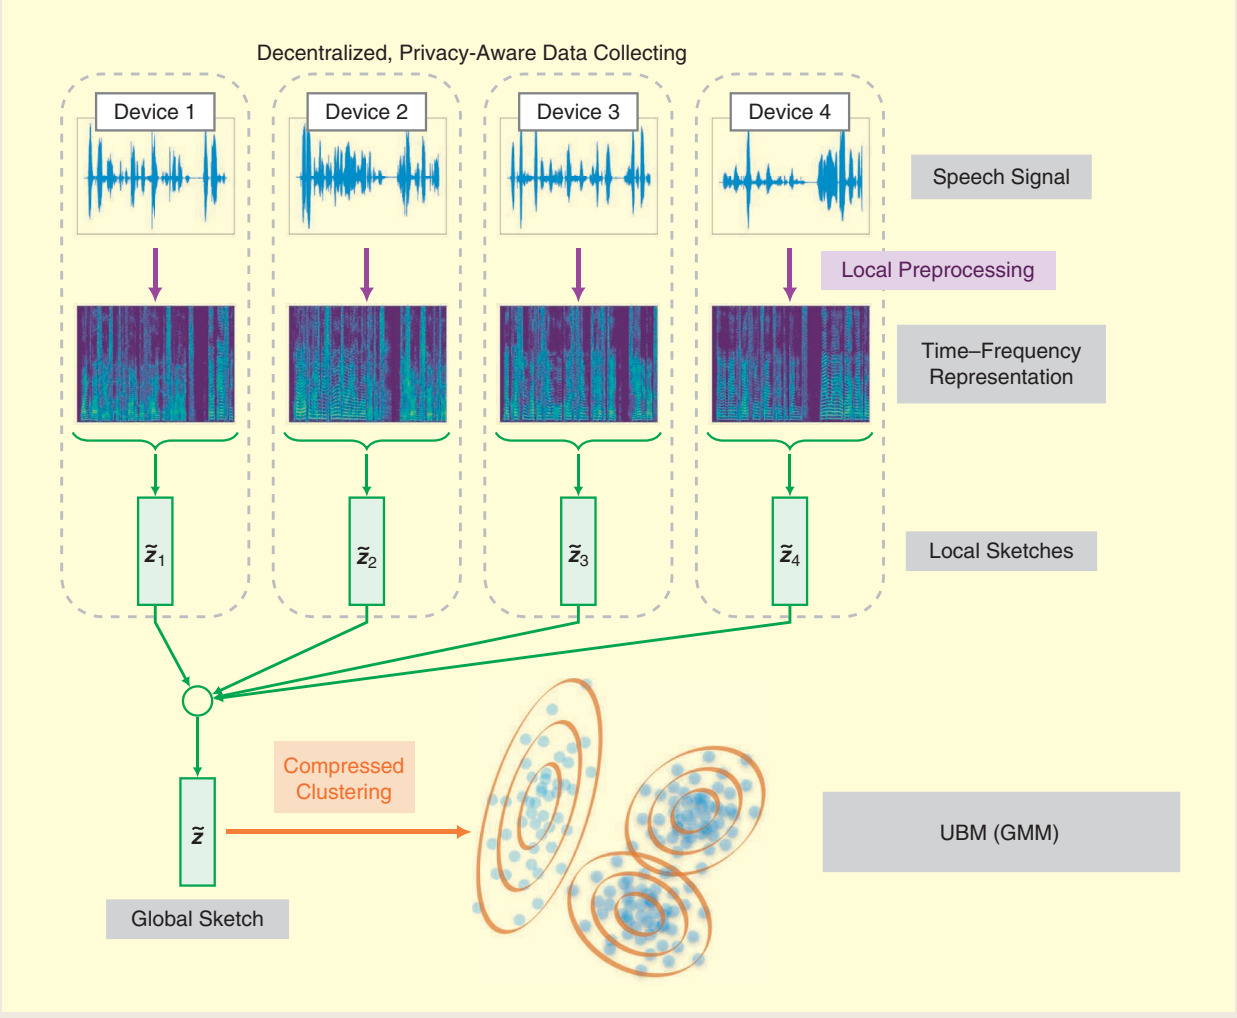

FIGURE S1. Speaker verification via compressive learning. Unlabeled training speech data are collected in a decentralized manner, preprocessed, and locally sketched. The local sketches are then merged into a global sketch, from which a UBM is learned.

attained by multiplying the data matrix $[\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n]^\top \in \mathbb{R}^{n \times d}$ by a random matrix on the left or the right [14]–[17]. *Sketching* may also refer to the use of sampling-based approaches [18], such as core sets [19] or the Nyström method [20]. 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 $\boldsymbol{\Phi}(\cdot)$, which is a key component of our sketch.

Sketches as (randomized) generalized moments

In signal processing and machine learning, the data samples $\boldsymbol{x}_i$ generally live in the vector space $\mathbb{R}^d$ and are often modeled as i.i.d. random vectors having a probability distribution with

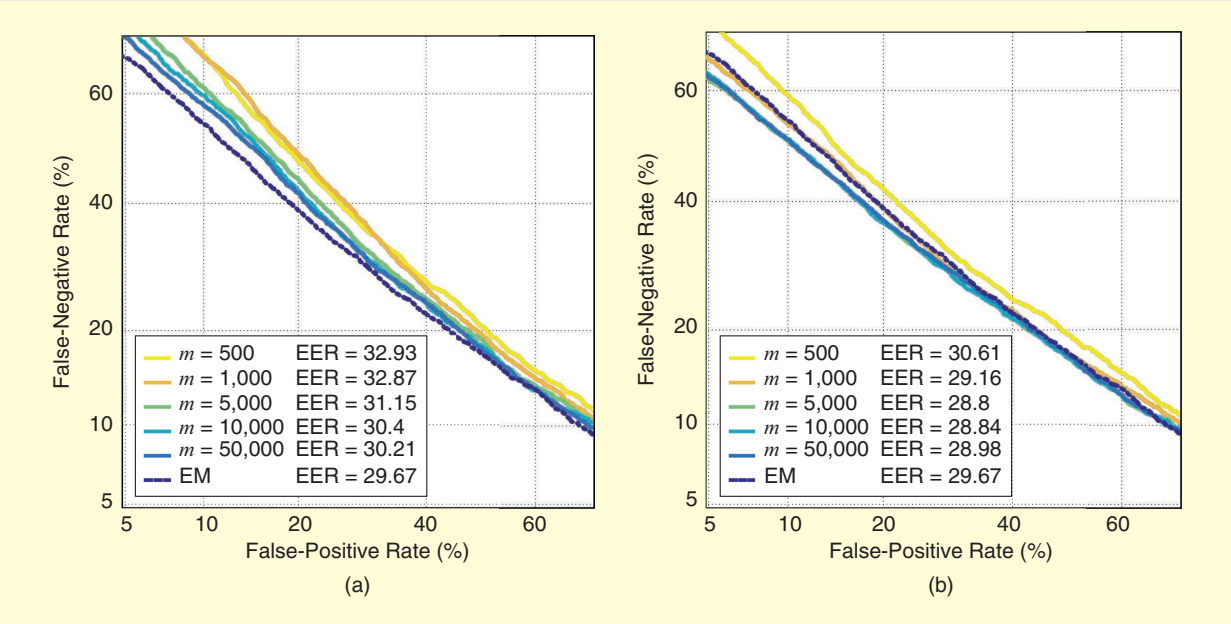

FIGURE S2. A detection error tradeoff curve (the false-positive/false-negative tradeoff obtained by varying the decision threshold) and an equal error rate (EER). On a single laptop, the EM was trained on 5 h of speech; the compressive GMM estimation used (a) 5 h and (b) 1,000 h for different sketch dimensions m.

decision is made if the ratio exceeds a certain threshold. In the GMM–UBM, the data are modeled using a GMM. That is, $p(\boldsymbol{x}|\boldsymbol{\theta})$ are multivariate Gaussian distributions applied to a suitable time–frequency transform of the raw audio signal [S1] (see Figure S1).

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 [6], using a random Fourier feature map $\boldsymbol{\Phi}(\cdot)$, the authors compressed 1,000 h 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 (15) introduced in the main text. They compared this compressive learning approach to the expectation–maximization (EM) algorithm, which, on the same laptop, could only be trained using 5 h of speech given the available random-access memory. 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 Figure S2).

Reference

[S1] D. A. Reynolds, T. F. Quatieri, and R. B. Dunn, "Speaker verification using adapted gaussian mixture models," *Dig. Sig. Proc.*, vol. 10, no. 1, pp. 1–3, 2000.

Compressive Clustering of MNIST Digits

Clustering is a classic machine learning task and a component of many machine learning pipelines. The most popular method is Lloyd's algorithm [8], 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 [9] applied this technique to handwritten digits from the Modified National Institute of Standards and Technology (MNIST) data set but used compressive k-means clustering in place of Lloyd's algorithm (see Figure S3). Using an $n = 10^7$ sample augmentation of the 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 four orders of magnitude, respectively (see Figure S4).

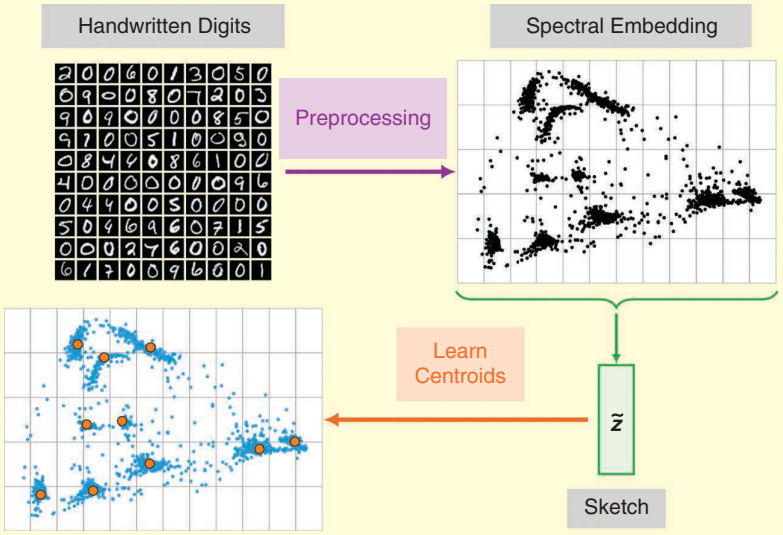

FIGURE S3. The clustering of handwritten digits via compressive k-means. First, scale-invariant feature transform descriptors are extracted from each image. Then, a similarity graph is constructed to obtain a so-called spectral embedding of the data set (using the first eigenvectors of the Laplacian of the graph). The spectral features are aggregated into a sketch, from which centroids are extracted.

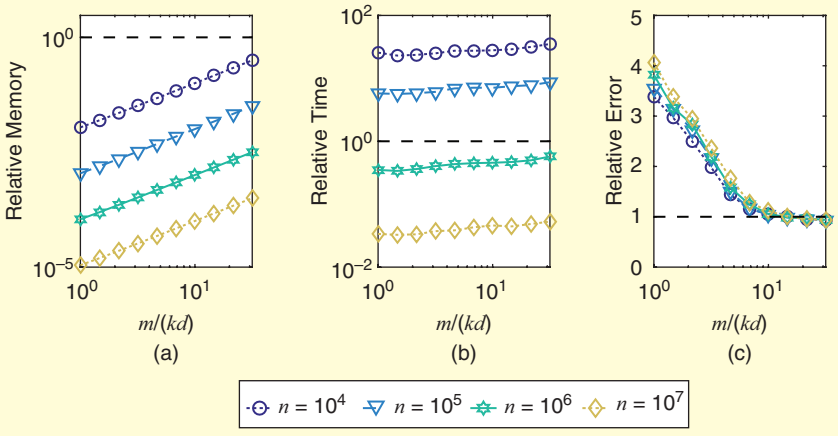

FIGURE S4. The (a) memory, (b) runtime, and (c) error of compressive k-means clustering relative to Lloyd's algorithm for various sample cardinalities n and sketch lengths m, using $k = 10$ centroids and $d = 10$ dimensional spectral features. Figure from [9].

Compressive Sensing and the Restricted Isometry Property

In compressive sensing, one observes a linear measurement $\boldsymbol{y} = \boldsymbol{A}\boldsymbol{x} \in \mathbb{R}^m$ of signal $\boldsymbol{x} \in \mathbb{R}^d$ 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 $\boldsymbol{x}$ in a given signal class, one desires that the distances between all signals in that class are preserved by the measurement operator $\boldsymbol{A}$. For the class of k-sparse signals Σ_k, i.e., signals with at most k nonzero entries, this property is satisfied up to a tolerance $\delta \in [0,1)$ when $\boldsymbol{A}$ obeys the restricted isometry property (RIP) [4], [22]

$$(1-\delta)\|\boldsymbol{x}-\boldsymbol{x}'\|^2 \le \|\boldsymbol{A}\boldsymbol{x}-\boldsymbol{A}\boldsymbol{x}'\|^2 \le (1+\delta)\|\boldsymbol{x}-\boldsymbol{x}'\|^2, \quad \forall \boldsymbol{x}, \boldsymbol{x}' \in \Sigma_k. \tag{S1}$$

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

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

The RIP also provides guarantees on robust recovery. Suppose that we have noisy measurements $\boldsymbol{y} = \boldsymbol{A}\boldsymbol{x} + \boldsymbol{e}$, with noise $\boldsymbol{e}$ of bounded norm $\|\boldsymbol{e}\| \le \varepsilon$. In addition, suppose that $\boldsymbol{x}$ is only approximately k-sparse, and use $\boldsymbol{x}_k$ to denote the best k-sparse approximation of $\boldsymbol{x}$. Finally, consider recovering an estimate $\hat{\boldsymbol{x}}$ of $\boldsymbol{x}$ by searching for the signal with the smallest ℓ_1 norm that agrees with the measurements up to a tolerance of ε, i.e., within the set $C_{\boldsymbol{y},\varepsilon} := \{\boldsymbol{u} : \|\boldsymbol{A}\boldsymbol{u} - \boldsymbol{y}\| \le \varepsilon\}$. Then, if the RIP (S1) holds, the estimation error $\hat{\boldsymbol{x}} - \boldsymbol{x}$ satisfies [4]

$$\|\hat{\boldsymbol{x}} - \boldsymbol{x}\| \le C\frac{\|\boldsymbol{x} - \boldsymbol{x}_k\|_1}{\sqrt{k}} + D\varepsilon, \tag{S2}$$

where constants $C, D > 0$ depend only on the value of δ that appears in (S1). Thus, the estimation error increases linearly with the noise level ε and the deviation $\|\boldsymbol{x} - \boldsymbol{x}_k\|_1$ from perfect sparsity. Please see Figure S5.

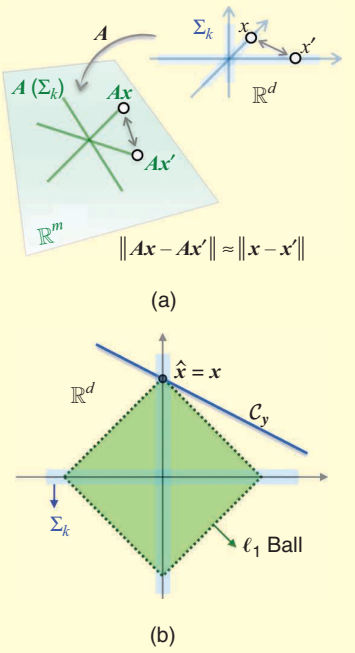

FIGURE S5. (a) A geometrical interpretation of the RIP. (b) A 2D illustration of the recoverability of $\boldsymbol{x}$ from $\boldsymbol{y} = \boldsymbol{A}\boldsymbol{x}$ by finding the vector $\hat{\boldsymbol{x}}$ in $C_{\boldsymbol{y}}$ with the smallest ℓ_1 norm.

density p_X. Consider what happens when the number of samples n goes to infinity. The strong law of large numbers says that

$$\lim_{n\to\infty} \frac{1}{n}\sum_{i=1}^{n} \boldsymbol{\Phi}(\boldsymbol{x}_i) \overset{\text{w.p.1}}{=} \mathbb{E}[\boldsymbol{\Phi}(X)] = \int p_X(\boldsymbol{x})\boldsymbol{\Phi}(\boldsymbol{x})\mathrm{d}\boldsymbol{x}, \tag{4}$$

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 $\boldsymbol{\Phi}(x) = x^k$ (so that $m = 1$), then $\mathbb{E}[\boldsymbol{\Phi}(X)]$ is the (uncentered) kth moment of the random variable X, a quantity that has a long history in statistics. By analogy, with a generic vector-valued feature map $\boldsymbol{\Phi}(\cdot)$ and in dimension $d > 1$, quantities of the form $\mathbb{E}[\boldsymbol{\Phi}(X)]$ 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) [21]. This method to "learn from a sketch" is very popular in, e.g., the field of econometrics [21, Ch. 1]. The GeMM can be seen as an alternative to ML 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 $\boldsymbol{\Phi}(\cdot)$ (see the sidebar "Compressive Mixture Modeling for Speaker Verification").

The GeMM differs from compressive learning in several aspects. In the GeMM, the feature map $\boldsymbol{\Phi}(\cdot)$ is typically constructed to make the parameter estimates computable in closed form. This narrows the range of learning tasks that the GeMM can handle. In compressive learning, $\boldsymbol{\Phi}(\cdot)$ 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, $\boldsymbol{\Phi}(\cdot)$ is typically randomized. This results in randomized generalized moments, which are rarely seen in the GeMM.

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 $\{\boldsymbol{x}_i\}_{i=1}^n$. This differs from the traditional approach in signal processing, where dimensionality reduction is performed on the features $\boldsymbol{x}_i \in \mathbb{R}^d$ 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], [22], a signal of interest is modeled as a high-dimensional vector $\boldsymbol{x} \in \mathbb{R}^d$, and a physical measurement of that signal is approximated by a linear transformation plus additive noise: $\boldsymbol{y} = \boldsymbol{A}\boldsymbol{x} + \boldsymbol{e} \in \mathbb{R}^m$ (see the sidebar "Compressive Sensing and the Restricted Isometry Property"). Here, the linear measurement operator is represented by the $m \times d$ matrix $\boldsymbol{A}$. 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 $\boldsymbol{x}$ from $\boldsymbol{y}$ is generally ill posed. Still, it is possible to accurately recover $\boldsymbol{x}$ from $\boldsymbol{y}$ when both $\boldsymbol{x}$ and $\boldsymbol{A}$ obey certain properties and the noise is small enough.

The celebrated Johnson–Lindenstrauss (JL) lemma and its variations specify, 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 CS: if $\boldsymbol{x}$ is sparse, in that relatively few of its coefficients deviate significantly from zero, and if $\boldsymbol{A}$ 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 $\boldsymbol{x}$ (see the sidebar "Compressive Sensing and the Restricted Isometry Property").

Now that CS has been described, we can clearly connect it to compressive learning. Recall that the sketch (1) $\tilde{\boldsymbol{z}}$ converges to the generalized moment $\boldsymbol{z} := \mathbb{E}[\boldsymbol{\Phi}(X)] = \int p_X(\boldsymbol{x})\boldsymbol{\Phi}(\boldsymbol{x})\,\mathrm{d}\boldsymbol{x}$ 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 $\boldsymbol{z}$ 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

$$\mathbb{E}_{X\sim p_X}[\boldsymbol{\Phi}(X)] = \alpha\mathbb{E}_{X_1\sim p_{X_1}}[\boldsymbol{\Phi}(X_1)] + (1-\alpha)\mathbb{E}_{X_2\sim p_{X_2}}[\boldsymbol{\Phi}(X_2)], \quad (5)$$

which implies that the generalized moment $\boldsymbol{z}$ is linear in p_X. With this understanding, we can write

$$\boldsymbol{z} = \mathcal{A}(p_X) := \mathbb{E}_{X\sim p_X}[\boldsymbol{\Phi}(X)], \quad (6)$$

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

Although $\mathcal{A}$ is a linear function of p_X, we emphasize that the feature map $\boldsymbol{\Phi}(\boldsymbol{x})$ is generally not a linear function of $\boldsymbol{x}$. 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 preceding arguments, we can handle the finite sample case. Consider the difference between the true generalized moment $\boldsymbol{z}$ and the empirical moment $\tilde{\boldsymbol{z}}$; i.e.,

$$\frac{1}{n}\sum_{i=1}^{n}\boldsymbol{\Phi}(\boldsymbol{x}_i) - \mathbb{E}[\boldsymbol{\Phi}(X)] =: \boldsymbol{e}. \quad (7)$$

As we discussed earlier, $\boldsymbol{e}$ converges to zero as n tends to infinity by the law of large numbers. Combining (1), (6), and (7), we obtain

$$\tilde{\boldsymbol{z}} = \mathcal{A}(p_X) + \boldsymbol{e}, \quad (8)$$

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 conditions the central limit theorem can be used to show that $\|\boldsymbol{e}\|$ decays as $1/\sqrt{n}$ with high probability [4, Ch. 8]. From the assumed independence of the data samples, this happens, for instance, if $\mathbb{P}[\|\boldsymbol{\Phi}(X)\| \geq t]$ decays exponentially fast when t increases. With the preceding interpretation of compressive learning, one recovers all of the traditional ingredients of CS:

- The measurement operator $\mathcal{A}$ is linear.
- The measurements $\tilde{\boldsymbol{z}}$ are drastically dimension reduced. In mathematical terms, probability distributions p_X belong to the infinite–dimensional vector space of so-called finite measures [23]. 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 $\boldsymbol{\Phi}(\cdot)$, such as one based on RF features, as in (3).
- The measurements $\tilde{\boldsymbol{z}}$ are noisy, as per (6).

The analogy between compressive learning and CS is illustrated in Figure 4. The CS analog to the sketching phase (1) is the signal sensing phase, where the signal $\boldsymbol{x}$ is linearly mapped to the observation vector $\boldsymbol{y}$. The analog to sparse recovery, where an estimate of $\boldsymbol{x}$ is computed from the observation $\boldsymbol{y}$ by solving an optimization problem, is to learn from a sketch, where an estimate of the data distribution p_X (or of distributional

parameters $\boldsymbol{\theta}$ of interest) is computed from the sketch $\tilde{\boldsymbol{z}}$. Its expression as an optimization problem will be further discussed in the "Learning From a Sketch" section.

RF sampling and superresolution 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 RF sampling and superresolution recovery.

We begin by considering the goal of recovering k centroids $\{\boldsymbol{c}_\ell\}_{\ell=1}^k$ in $\mathbb{R}^d$ with Euclidean norm $\leq r$ that are separated from each other in Euclidean distance by $\geq \varepsilon$. 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 $\hat{\boldsymbol{p}} \in \mathbb{R}_+^N$ 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 $\hat{\boldsymbol{p}}$ 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 [13] of the form $\tilde{\boldsymbol{z}} = \boldsymbol{A}\hat{\boldsymbol{p}}$, where matrix $\boldsymbol{A} \in \mathbb{R}^{m\times N}$ 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 nonnegative, k-sparse, sum-to-one vector using

$$\tilde{\boldsymbol{p}} = \operatorname*{argmin}_{\boldsymbol{p}\in\Sigma_k^+} \left\| \tilde{\boldsymbol{z}} - \boldsymbol{A}\boldsymbol{p} \right\|^2, \tag{9}$$

(or a convex or greedy relaxation of this problem) where Σ_k^+ here denotes the set of k-sparse, nonnegative, sum-to-one vectors. Then, one would identify the k grid locations in $\mathbb{R}^d$ corresponding to the nonzero indices of $\tilde{\boldsymbol{p}}$. CS theory [4] says that the support of $\tilde{\boldsymbol{p}}$ 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 $\boldsymbol{A}$ grows exponentially in d, making storage and multiplication by $\boldsymbol{A}$ impractical.

An alternative approach could be to construct $\boldsymbol{A}$ 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 $\boldsymbol{A}$ could be implemented by the fast Fourier transform (FFT) algorithm, and thus the matrix $\boldsymbol{A}$ 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., [24] and [25]. When the m frequencies are drawn uniformly at random, CS theory [4], [22] 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 $\boldsymbol{A}$ as an explicit matrix and allows efficient computation of the sketch, the cost of solving the optimization problem (9) 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 $\hat{\boldsymbol{p}}$) whose dimension grows exponentially in d. We can avoid this discretization (i.e., take $\epsilon \to 0$ and $r \to \infty$) by replacing the DFT with the continuous Fourier transform (CFT), in which case we are Fourier transforming the empirical distribution $\hat{p}_{\mathcal{X}}(\boldsymbol{x}) \coloneqq (1/n)\sum_{i=1}^n \delta(\boldsymbol{x} - \boldsymbol{x}_i)$ rather than its N-bin histogram $\hat{\boldsymbol{p}}$. The sketch $\tilde{\boldsymbol{z}}$ then has components

$$\tilde{z}_j = \int_{\mathbb{R}^d} \hat{p}_{\mathcal{X}}(\boldsymbol{x}) \exp(-\mathrm{j}2\pi \boldsymbol{w}_j^\top \boldsymbol{x})\, \mathrm{d}\boldsymbol{x} \tag{10}$$

$$= \frac{1}{n}\sum_{i=1}^n \exp(-\mathrm{j}2\pi \boldsymbol{w}_j^\top \boldsymbol{x}_i) \qquad j = 1, \ldots, m, \tag{11}$$

where $\{\boldsymbol{w}_j\}_{j=1}^m$ are d-dimensional frequencies that are drawn at random. Typically, $\boldsymbol{w}_j$ are drawn i.i.d. Gaussian, but other distributions can be used, as discussed in the sections "The Challenge of Designing a Feature Map Given a Learning Task" and "Sketching With Structured Random Matrices." Note that (11) corresponds precisely to the sketch (1) with the RF feature map $\boldsymbol{\Phi}(\cdot)$ from (3). Taking a statistical perspective, the components $\tilde{z}_j$ in (10) 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

$$\Psi_p(\boldsymbol{w}) \coloneqq \int_{\mathbb{R}^d} p(\boldsymbol{x}) \exp(\mathrm{j}2\pi \boldsymbol{w}^\top \boldsymbol{x})\, \mathrm{d}\boldsymbol{x}$$
$$= \mathbb{E}_{X\sim p}[\exp(\mathrm{j}2\pi \boldsymbol{w}^\top X)]. \tag{12}$$

Intuitively, when a probability distribution p has a "simple" structure, one can recover it (with high probability) from enough randomly chosen samples of its CFT. Centroid recovery from the sketch (10) is premised on the empirical distribution $\hat{p}_{\mathcal{X}}$ being well approximated by a mixture of k Diracs;

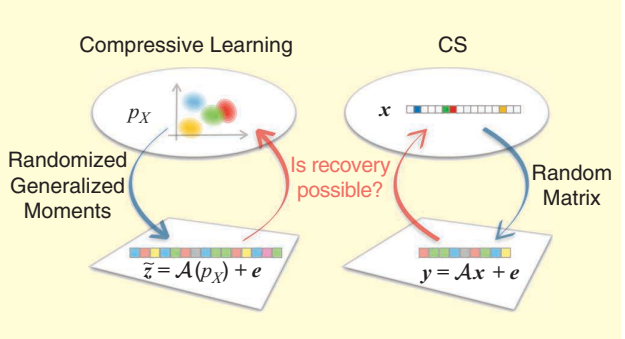

FIGURE 4. The analogy between compressive learning, which uses the dimensionality-reducing linear measurement $\mathcal{A}(p_X) = \mathbb{E}[\boldsymbol{\Phi}(X)]$ of a distribution p_X, and CS, which uses the dimensionality-reducing linear measurement $\boldsymbol{A}\boldsymbol{x}$ of a signal $\boldsymbol{x}$.

i.e., $\hat{p}_{\mathcal{X}}(\boldsymbol{x}) \approx \sum_{\ell=1}^{k} \alpha_\ell \delta(\boldsymbol{x} - \boldsymbol{c}_\ell)$. In this case, centroid recovery parallels the "superresolution" recovery problem (see the sidebar "Superresolution Recovery") through an optimization problem that is the continuous analog to (9) (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 $\boldsymbol{w}_j$ are randomly drawn. For example, when the centroids $\{\boldsymbol{c}_\ell\}_{\ell=1}^{k}$ 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], [51]. Similar guarantees hold when $\hat{p}_{\mathcal{X}}$ is approximately a sum of spatially localized components (e.g., in compressive GMM) [5], [51].

Learning from a sketch

Until now, we have primarily focused on the first stage of the compressive learning pipeline (see Figure 1), where the data set $\mathcal{X}$ is sketched down to $\tilde{\boldsymbol{z}}$, 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, $\boldsymbol{\theta}$, are recovered from the sketch $\tilde{\boldsymbol{z}}$. The close analogy between sketching and CS allows us to cast this "parameter learning" stage as an optimization problem; i.e.,

$$\tilde{\boldsymbol{\theta}} = \operatorname*{argmin}_{\boldsymbol{\theta}} C(\boldsymbol{\theta} \,|\, \tilde{\boldsymbol{z}}), \tag{13}$$

where the cost function $C(\cdot \,|\, \tilde{\boldsymbol{z}})$ is adapted to the considered learning task. As in CS, many candidate distributions p_X (and hence many candidate parameters $\boldsymbol{\theta}$) can yield the same sketch $\tilde{\boldsymbol{z}}$. 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 $\boldsymbol{x}$, 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 $\{p_{\boldsymbol{\theta}_\ell}\}_{\ell=1}^{k}$. Concretely, $p_X = \sum_{\ell=1}^{k} \alpha_\ell p_{\boldsymbol{\theta}_\ell}$, where the mixture weights $\{\alpha_\ell\}_{\ell=1}^{k}$ are nonnegative and sum to one. For example, with a GMM, we have that $p_{\boldsymbol{\theta}_\ell} = \mathcal{N}(\boldsymbol{\mu}_\ell, \boldsymbol{\Sigma}_\ell)$, where $\boldsymbol{\theta}_\ell = (\boldsymbol{\mu}_\ell, \boldsymbol{\Sigma}_\ell)$ contains the mean $\boldsymbol{\mu}_\ell$ and covariance $\boldsymbol{\Sigma}_\ell$. If p_X is well approximated by the mixture model $\sum_{\ell=1}^{k} \alpha_\ell p_{\boldsymbol{\theta}_\ell}$, then, according to (8), the sketch $\tilde{\boldsymbol{z}}$ is well approximated by the linear combination $\sum_{\ell=1}^{k} \alpha_\ell \mathcal{A}(p_{\boldsymbol{\theta}_\ell})$. Hence, one could try to extract the mixture parameters, $\boldsymbol{\theta} = \{\alpha_\ell, \boldsymbol{\theta}_\ell\}_{\ell=1}^{k}$, from the sketch $\tilde{\boldsymbol{z}}$ by solving the (nonconvex) optimization problem (13) with

$$C(\boldsymbol{\theta} \,|\, \tilde{\boldsymbol{z}}) \coloneqq \left\| \tilde{\boldsymbol{z}} - \sum_{\ell=1}^{k} \alpha_\ell \mathcal{A}(p_{\boldsymbol{\theta}_\ell}) \right\|^2. \tag{14}$$

The cost $C(\boldsymbol{\theta} \,|\, \tilde{\boldsymbol{z}})$ can be interpreted as the negative log likelihood (up to a shift and scale) of $\boldsymbol{\theta}$ given the sketch $\tilde{\boldsymbol{z}}$, under the classic modeling assumption of i.i.d. Gaussian measurement noise $\boldsymbol{e}$ in (8).

When the RF feature map $\boldsymbol{\Phi}(\cdot)$ is used to compute the sketch $\tilde{\boldsymbol{z}}$ and the component densities $p_{\boldsymbol{\theta}_\ell}$ are Gaussian or α stable, there exist analytic expressions for $\mathcal{A}(p_{\boldsymbol{\theta}_\ell})$ and for the gradient of $\mathcal{A}(p_{\boldsymbol{\theta}_\ell})$ with respect to the mixture parameters in $\boldsymbol{\theta}_\ell$ [6]. These expressions are convenient when numerically optimizing (14). For instance, greedy approaches, similar to the orthogonal matching pursuit (OMP) algorithm for CS (recall the sidebar

Superresolution Recovery

Superresolution is a general class of techniques to enhance the resolution of a sensing system, e.g., to observe subwavelength features in astronomy or medical imaging [23]. The problem addressed by superresolution is to recover a continuous-time (when dimension $d = 1$) or continuous-space (when dimension $d \geq 2$) sparse signal $s(\boldsymbol{t})$ from a few, possibly noisy, Fourier measurements $\{y_j\}_{j=1}^{m}$. This amounts to recovering a weighted sum $s(\boldsymbol{t}) = \sum_{\ell=1}^{k} \alpha_\ell \delta(\boldsymbol{t} - \boldsymbol{t}_\ell)$ of k Diracs with amplitudes α_ℓ and locations $\boldsymbol{t}_\ell \in \mathbb{R}^d$ from

$$y_j = \int_{\mathbb{R}^d} s(\boldsymbol{t}) \exp(-\mathrm{j}2\pi \boldsymbol{w}_j^{\mathsf{T}} \boldsymbol{t}) \, \mathrm{d}\boldsymbol{t} + e_j, \qquad j = 1, \ldots, m, \tag{S3}$$

with measurement noise e_j and frequency vectors $\{\boldsymbol{w}_j\}_{j=1}^{m}$ in $\mathbb{R}^d$. Recovery of $s(\cdot)$ can be posed as an infinite–dimensional convex problem on measures [23]. However, most reconstruction algorithms involve nonconvex steps [26]. When the frequencies $\boldsymbol{w}_j$ 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 $\boldsymbol{t}_\ell$ [S2] or positivity of the amplitudes α_ℓ [26].

The link between superresolution and compressive clustering follows from rewriting (10) as

$$\tilde{z}_j = \int_{\mathbb{R}^d} p_X(\boldsymbol{x}) \exp(-\mathrm{j}2\pi \boldsymbol{w}_j^{\mathsf{T}} \boldsymbol{x}) \, \mathrm{d}\boldsymbol{x} + e_j, \qquad j = 1, \ldots, m, \tag{S4}$$

where e_j captures the "noise" due to finite-sample effects [recall (7)]. Comparing (S4) to (S3), we see that they are mathematically equivalent when $p_X(\boldsymbol{x}) = \sum_{\ell=1}^{k} \alpha_\ell \delta(\boldsymbol{x} - \boldsymbol{c}_\ell)$ except for the fact that, in the case of compressive clustering, α_ℓ are nonnegative and sum to one.

Reference

[S2] C. Poon, N. Keriven, and G. Peyré, "The geometry of off-the-grid compressed sensing," 2020. arXiv: 1802.08464.

"Compressive Sensing and the Restricted Isometry Property"), can be used [6] to estimate the parameters $\{\alpha_\ell, \boldsymbol{\theta}_\ell\}_{\ell=1}^k$. These approaches sequentially estimate and subtract, from $\tilde{\boldsymbol{z}}$, each of the k components $\alpha_\ell \mathcal{A}(p_{\boldsymbol{\theta}_\ell})$ that best align with it, where "best" is measured via the correlation between $\tilde{\boldsymbol{z}}$ and $\mathcal{A}(p_{\boldsymbol{\theta}_\ell})$. As another example, an iterative approach [10] was proposed that exploits the log likelihood interpretation of $C(\boldsymbol{\theta} \,|\, \tilde{\boldsymbol{z}})$ in (14) and the i.i.d. random nature of the linear transform $\boldsymbol{W}$ in the RF map (3). In [6], applications of compressive mixture modeling are demonstrated on speaker verification (see the sidebar "Compressive Mixture Modeling for Speaker Verification") and source separation.

As a second example, we consider compressive k-means clustering. Here, the goal is to recover, from the sketch $\tilde{\boldsymbol{z}}$ the centroids $\boldsymbol{\theta} = \{\boldsymbol{c}_\ell\}_{\ell=1}^k$ that minimize the average squared Euclidean distance from each sample to its nearest centroid; i.e., $(1/n)\sum_{i=1}^n \min_\ell \|\boldsymbol{x}_i - \boldsymbol{c}_\ell\|^2$. 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 $\boldsymbol{c}_\ell$, and covariance matrices $\boldsymbol{\Sigma}_\ell$. Then, in the special case that $\alpha_\ell = 1/k$ and $\boldsymbol{\Sigma}_\ell = \sigma^2 \boldsymbol{I}$ for all components $1 \le \ell \le k$, we can write the likelihood as $\prod_{i=1}^n p(\boldsymbol{x}_i|\boldsymbol{\theta}) \propto \prod_{i=1}^n \sum_{\ell=1}^k \exp(-(1/2\sigma^2)\|\boldsymbol{x}_i - \boldsymbol{c}_l\|^2)$, and so the negative log likelihood becomes $-\sum_{i=1}^n \log \sum_{\ell=1}^k \exp(-(1/2\sigma^2)\|\boldsymbol{x}_i - \boldsymbol{c}_l\|^2)$ up to an additive constant. We can then use the log–sum–exp approximation $\log \sum_\ell \exp(f_\ell(\boldsymbol{x})) \approx \max_\ell f_\ell(\boldsymbol{x})$ to approximate this latter expression as $(1/2\sigma^2)\sum_{i=1}^n \min_\ell \|\boldsymbol{x}_i - \boldsymbol{c}_\ell\|^2$, which agrees with the k-means cost up to a scaling. If we furthermore consider the case of a vanishing variance $\sigma^2 \to 0$, then the component density $p_{\boldsymbol{\theta}_\ell}$ reduces to a point mass; i.e., $p_{\boldsymbol{\theta}_\ell}(\boldsymbol{x}) \to \delta(\boldsymbol{x} - \boldsymbol{c}_\ell)$. In this limiting case, the linear measurement of the point mass $p_{\boldsymbol{\theta}_\ell}$ is $\mathcal{A}(p_{\boldsymbol{\theta}_\ell}) = \mathbb{E}_{X\sim p_{\boldsymbol{\theta}_\ell}}[\boldsymbol{\Phi}(X)] = \int \boldsymbol{\Phi}(\boldsymbol{x}) p_{\boldsymbol{\theta}_\ell}(\boldsymbol{x})\, \mathrm{d}\boldsymbol{x} = \boldsymbol{\Phi}(\boldsymbol{c}_\ell)$, according to (6).

Thus, with these justifications, the cost function (14) for the compressive GMM would change to

$$C(\boldsymbol{\theta} \,|\, \tilde{\boldsymbol{z}}) := \left\| \tilde{\boldsymbol{z}} - \frac{1}{k} \sum_{\ell=1}^k \boldsymbol{\Phi}(\boldsymbol{c}_\ell) \right\|^2 \tag{15}$$

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 [9]:

$$C(\boldsymbol{\theta} \,|\, \tilde{\boldsymbol{z}}) := \min_{\boldsymbol{\alpha}} \left\| \tilde{\boldsymbol{z}} - \sum_{\ell=1}^k \alpha_\ell \boldsymbol{\Phi}(\boldsymbol{c}_\ell) \right\|^2. \tag{16}$$

Similar to the compressive GMM problem described earlier, minimization of the cost function (16) for compressive k-means can be tackled by greedy approaches, as described in [9]. Despite the fact that (16) does not directly minimize the k-means cost, it has been shown empirically [9] that the centroids estimated by such greedy algorithms nearly minimize this cost; see, e.g., the sidebar "Compressive Clustering of MNIST Digits" for an example on MNIST data. This claim is also supported by theoretical results guaranteeing that the minimizer of (16) is endowed with statistical learning guarantees with respect to the original k-means cost [5], [51]. Such guarantees will be discussed shortly.

Depending on the choice of parameters $\boldsymbol{\theta}$ and the feature map $\boldsymbol{\Phi}(\cdot)$, the form of the optimization problem posed to learn from a sketch can differ considerably from that for GMMs in (14) and that for the k-means in (15) and (16). 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 $\{\boldsymbol{x}_i\}_{i=1}^n$ in an LS sense. It is well known that this subspace is spanned by k principal eigenvectors of the empirical autocorrelation matrix $\hat{\boldsymbol{R}} = (1/n)\sum_{i=1}^n \boldsymbol{x}_i \boldsymbol{x}_i^\top$, or, equivalently, the column space of the (positive semidefinite symmetric) matrix $\hat{\boldsymbol{R}}_k$ that is closest to $\hat{\boldsymbol{R}}$ in the Frobenius norm:

$$\hat{\boldsymbol{R}}_k := \operatorname*{argmin}_{\boldsymbol{R}:\operatorname{rank}(\boldsymbol{R})\le k} \left\| \hat{\boldsymbol{R}} - \boldsymbol{R} \right\|_F = \operatorname*{argmin}_{\boldsymbol{R}:\operatorname{rank}(\boldsymbol{R})\le k} \left\| \operatorname{vec}(\hat{\boldsymbol{R}}) - \operatorname{vec}(\boldsymbol{R}) \right\|^2. \tag{17}$$

When sketching using quadratic features $\boldsymbol{\Phi}(\boldsymbol{x}) = ((\boldsymbol{w}_1^\top \boldsymbol{x})^2, \ldots, (\boldsymbol{w}_m^\top \boldsymbol{x})^2)$ with random $\boldsymbol{w}_j \in \mathbb{R}^d$, the jth component of the sketch becomes $\tilde{z}_j = (1/n)\sum_{i=1}^n \boldsymbol{w}_j^\top \boldsymbol{x}_i \boldsymbol{x}_i^\top \boldsymbol{w}_j = \boldsymbol{w}_j^\top \hat{\boldsymbol{R}} \boldsymbol{w}_j$. Importantly, this $\tilde{z}_j$ is a linear function of $\hat{\boldsymbol{R}}$, and so there exists an $m \times d^2$ matrix $\boldsymbol{A}$ such that $\tilde{\boldsymbol{z}} = \boldsymbol{A} \operatorname{vec}(\hat{\boldsymbol{R}})$. Thus, by analogy with (17), one could first fit a positive semidefinite symmetric low-rank matrix to the sketch $\tilde{\boldsymbol{z}}$ via

$$\tilde{\boldsymbol{R}} = \operatorname*{argmin}_{\boldsymbol{R}:\operatorname{rank}(\boldsymbol{R})\le k, \boldsymbol{R}^\top = \boldsymbol{R}, \boldsymbol{R} \ge 0} \left\| \tilde{\boldsymbol{z}} - \boldsymbol{A} \operatorname{vec}(\boldsymbol{R}) \right\|^2 \tag{18}$$

and then set the parameter estimate $\tilde{\boldsymbol{\theta}}$ equal to the column space of $\tilde{\boldsymbol{R}}$. While the optimum of (17) is automatically symmetric and positive semidefinite, this property needs to be enforced explicitly in (18).

The low-rank matrix recovery problem (18) has been thoroughly investigated by the signal processing and machine learning communities (e.g., [27]). It arises, for example, in applications such as collaborative filtering for recommender systems and signal reconstruction from phaseless measurements. Early approaches to solving the nonconvex problem (18) involved convex relaxation via nuclear norm regularization [4, Ch. 4]. More recent approaches exploit the nonconvex geometry of (18) [28].

Compressive learning with theoretical guarantees

In the previous section, learning from a sketch was posed as the optimization problem (13), repeated here for convenience:

$$\tilde{\boldsymbol{\theta}} = \operatorname*{argmin}_{\boldsymbol{\theta}} C(\boldsymbol{\theta} \,|\, \tilde{\boldsymbol{z}}). \tag{19}$$

The quantity $C(\cdot \,|\, \tilde{\boldsymbol{z}})$ is a real-valued cost function whose minimizer $\tilde{\boldsymbol{\theta}}$ is the "best" (in some sense) estimate of $\boldsymbol{\theta}$ from the sketch $\tilde{\boldsymbol{z}}$. This approach contrasts with traditional statistical learning [29], [30], which computes the parameter estimate

$$\hat{\boldsymbol{\theta}} := \operatorname*{argmin}_{\boldsymbol{\theta}} R(\boldsymbol{\theta} \,|\, \mathcal{X}) \quad \text{with} \quad R(\boldsymbol{\theta} \,|\, \mathcal{X}) := \frac{1}{n} \sum_{i=1}^n L(\boldsymbol{\theta} \,|\, \boldsymbol{x}_i), \tag{20}$$

where $R(\cdot\,|\,\mathcal{X})$ is the "empirical risk" and $L(\cdot\,|\,\boldsymbol{x}_i)$ is the "loss function" for the ith sample. Table 2 presents typical loss functions for our four running examples. In most cases, performing the minimization in (20) involves querying the full training data set $\mathcal{X}$ many times, e.g., for stochastic gradient descent. In compressive learning, the cost $C(\cdot\,|\,\tilde{\boldsymbol{z}})$, which depends only on the low-dimensional sketch $\tilde{\boldsymbol{z}}$, 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 data set $\mathcal{X}$, the minimization in (13) can be made much more efficient than that in (20).

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

$$\boldsymbol{\theta}^\star := \operatorname*{argmin}_{\boldsymbol{\theta}} R^\star(\boldsymbol{\theta}) \quad \text{with} \quad R^\star(\boldsymbol{\theta}) := \mathbb{E}\,[L(\boldsymbol{\theta}\,|\,X)], \tag{21}$$

where $R^\star(\boldsymbol{\theta})$ is known as the "true risk." Indeed, the feature map $\boldsymbol{\Phi}(\cdot)$ and cost $C(\cdot\,|\,\tilde{\boldsymbol{z}})$ are designed precisely to ensure that the surrogate minimization (19) approximates the true risk minimization (21).

Excess risk guarantees

To establish a guarantee on the goodness of an estimate $\boldsymbol{\theta}$, one must prove that the "excess risk" $R^\star(\boldsymbol{\theta}) - R^\star(\boldsymbol{\theta}^\star) \geq 0$ 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 $R^\star(\cdot)$ is the primary metric by which one judges the quality of an arbitrary estimate $\boldsymbol{\theta}$. Note that controlling the excess risk is different from proving that $\boldsymbol{\theta}$ is close to the ideal estimate $\boldsymbol{\theta}^\star$ in, e.g., the Euclidean distance $\|\boldsymbol{\theta} - \boldsymbol{\theta}^\star\|$. Consider, for example, the problem of fitting a 1D linear subspace to a data set (i.e., PCA with $k = 1$). When the two largest eigenvalues of the autocorrelation matrix $\boldsymbol{R}$ are equal, any nontrivial linear combination of the corresponding eigenvectors generates a 1D subspace $\boldsymbol{\theta}$ 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.

Table 2. The loss functions for our four running examples.

Running Example	Parameters $\boldsymbol{\theta}$	Loss Function $L(\boldsymbol{\theta}\,|\,\boldsymbol{x}_i)$
PCA	Orthonormal basis $\boldsymbol{\theta} = \{\boldsymbol{u}_\ell\}_{\ell=1}^k$	$\sum_{\ell=1}^k \lvert \boldsymbol{x}_i^\top \boldsymbol{u}_\ell \rvert^2$
LS linear regression	A weight matrix $\boldsymbol{\theta}$	$\lVert \boldsymbol{x}_{1i} - \boldsymbol{\theta}\boldsymbol{x}_{2i} \rVert^2$, $\boldsymbol{x}_i := (\boldsymbol{x}_{1i}^\top, \boldsymbol{x}_{2i}^\top)^\top$
Gaussian mixture modeling	GMM parameters $\boldsymbol{\theta} = \{\alpha_\ell, \boldsymbol{\mu}_\ell, \boldsymbol{\Sigma}_\ell\}_{\ell=1}^k$	$-\log \sum_{\ell=1}^k \alpha_\ell \mathcal{N}(\boldsymbol{x}_i; \boldsymbol{\mu}_\ell, \boldsymbol{\Sigma}_\ell)$
k-means clustering	A set of centroids $\boldsymbol{\theta} = \{\boldsymbol{c}_\ell\}_{\ell=1}^k$	$\min_\ell \lVert \boldsymbol{x}_i - \boldsymbol{c}_\ell \rVert^2$

Despite the fact that the estimates $\tilde{\boldsymbol{\theta}}, \hat{\boldsymbol{\theta}}$, and $\boldsymbol{\theta}^\star$ minimize different objective functions, they often yield similar true risks; i.e., the excess risks of $\tilde{\boldsymbol{\theta}}$ and $\hat{\boldsymbol{\theta}}$ are provably small. For $\hat{\boldsymbol{\theta}}$, the proofs use classic results from statistical learning [29], [30], under assumptions that we will briefly discuss in the sequel. For $\tilde{\boldsymbol{\theta}}$, with an appropriately designed feature map $\boldsymbol{\Phi}(\cdot)$ and cost function $C(\cdot\,|\,\tilde{\boldsymbol{z}})$, the proofs informally follow from the fact that $C(\cdot\,|\,\tilde{\boldsymbol{z}})$ and $R^\star(\cdot)$ have a similar shape. This fact is illustrated in the sidebar "Traditional Statistical Learning Versus Compressive Learning" for a toy example of compressive clustering. There it can be seen that, with a properly designed feature map $\boldsymbol{\Phi}(\cdot)$, a well-chosen cost function $C(\cdot\,|\,\tilde{\boldsymbol{z}})$, and a sufficiently large sketch dimension m, the minimizer $\tilde{\boldsymbol{\theta}}$ of the cost yields a nearly minimal true risk $R^\star(\tilde{\boldsymbol{\theta}})$ and lives in a large basin of attraction in $C(\cdot\,|\,\tilde{\boldsymbol{z}})$. If the sketch dimension m is chosen too small, however, then the excess risk $R^\star(\tilde{\boldsymbol{\theta}}) - R^\star(\boldsymbol{\theta}^\star)$ increases.

A theory of compressive learning [5], [51] has been developed to better understand how (random) feature maps $\boldsymbol{\Phi}(\cdot)$ and cost functions $C(\cdot\,|\,\tilde{\boldsymbol{z}})$ can be designed to ensure that the sketch $\tilde{\boldsymbol{z}}$ captures sufficient information to control the excess risk on $\tilde{\boldsymbol{\theta}}$. In the following four sections, we aim to summarize the key aspects of this theory using broadly accessible language. For those interested, the full technical details can be found in [5], [51].

In a nutshell, the theory relies on interpreting $\tilde{\boldsymbol{\theta}}$ 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 the compressive GMM, $\tilde{p}$ is a Gaussian mixture).
- $\tilde{p}$ minimizes $\|\tilde{\boldsymbol{z}} - \mathcal{A}(p)\|$ among all simple distributions p.

For example, in the compressive GMM, this holds when $\tilde{p} = p_{\tilde{\boldsymbol{\theta}}}$, i.e., the Gaussian mixture distribution parameterized by $\tilde{\boldsymbol{\theta}}$. 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 the section "Task-Driven Distances and Excess Risk Bounds").
- Given a family of "simple" probability distributions (see the section "Exploiting Simplified Models to Learn From a Sketch"), this distance can indeed be controlled provided a certain "lower RIP" (LRIP) holds (see the section "The LRIP and Excess Risk Control"). This is analogous to traditional CS, where the RIP allows one to control the performance of sparse signal recovery.
- To establish the LRIP for random feature maps $\boldsymbol{\Phi}(\cdot)$, one can build on connections between kernel methods and the JL lemma (see the section "Establishing the LRIP via Kernels and the JL Lemma").
- Finally, to design task-specific feature maps $\boldsymbol{\Phi}(\cdot)$, one can appeal to the problem of kernel design (see the section "The Challenge of Designing a Feature Map Given a Learning Task").

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

$$d(p'_X, p_X) = \sup_{\boldsymbol{\theta}} \left| R^\star(\boldsymbol{\theta} \,|\, p'_X) - R^\star(\boldsymbol{\theta} \,|\, p_X) \right|, \tag{22}$$

where $R^\star(\boldsymbol{\theta} \,|\, p) := \mathbb{E}_{X\sim p}[L(\boldsymbol{\theta} \,|\, X)]$ 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

$$\begin{aligned} R^\star(\boldsymbol{\theta}^\star \,|\, p_X) \le R^\star(\boldsymbol{\theta}^{\star\prime} \,|\, p_X) &\le R^\star(\boldsymbol{\theta}^{\star\prime} \,|\, p'_X) + d(p'_X, p_X) \\ &\le R^\star(\boldsymbol{\theta}^\star \,|\, p'_X) + d(p'_X, p_X) \\ &\le R^\star(\boldsymbol{\theta}^\star \,|\, p_X) + 2d(p'_X, p_X), \end{aligned} \tag{23}$$

where $\boldsymbol{\theta}^\star := \operatorname{argmin}_{\boldsymbol{\theta}} R^\star(\boldsymbol{\theta} \,|\, p_X)$ and $\boldsymbol{\theta}^{\star\prime} := \operatorname{argmin}_{\boldsymbol{\theta}} R^\star(\boldsymbol{\theta} \,|\, p'_X)$. In particular, these inequalities yield the following bounds on the excess risk of $\boldsymbol{\theta}^{\star\prime}$:

$$0 \le R^\star(\boldsymbol{\theta}^{\star\prime} \,|\, p_X) - R^\star(\boldsymbol{\theta}^\star \,|\, p_X) \le 2d(p'_X, p_X). \tag{24}$$

For example, if p'_X equals the empirical distribution $\hat{p}_{\mathcal{X}}(\boldsymbol{x}) := (1/n)\Sigma_{i=1}^n \delta(\boldsymbol{x} - \boldsymbol{x}_i)$, then (24) bounds the "generalization error" of empirical risk minimization, i.e., the error incurred when training with $\mathcal{X} = \{\boldsymbol{x}_i\}_{i=1}^n$ 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, $d(\hat{p}_{\mathcal{X}}, p_X) = O(1/\sqrt{n})$ with high probability on the draw of i.i.d. training samples $\mathcal{X} = \{\boldsymbol{x}_i\}_{i=1}^n$. The term *uniform convergence* stems from the analogy between the distance (22) and the ℓ_∞ norm.

Traditional Statistical Learning Versus Compressive Learning

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

To gain intuition into how these quantities manifest in compressive learning, Figure S6 shows a simple example: compressive k-means clustering of 1D data $\mathcal{X} = \{\boldsymbol{x}_i\}_{i=1}^n$ with two centroids, $\boldsymbol{\theta}_1$ and $\boldsymbol{\theta}_2$. The true risk $R^\star(\cdot)$ and the cost $C(\cdot|\tilde{\boldsymbol{z}})$ are plotted versus $\boldsymbol{\theta} = (\boldsymbol{\theta}_1, \boldsymbol{\theta}_2)$ and versus a 1D slice in the $\boldsymbol{\theta}$ plane. As the sketch dimension m increases, it can be seen that the excess risk $R^\star(\tilde{\boldsymbol{\theta}}) - R^\star(\boldsymbol{\theta}^\star)$ decreases. Moreover, although the cost function $C(\cdot|\tilde{\boldsymbol{z}})$ 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.

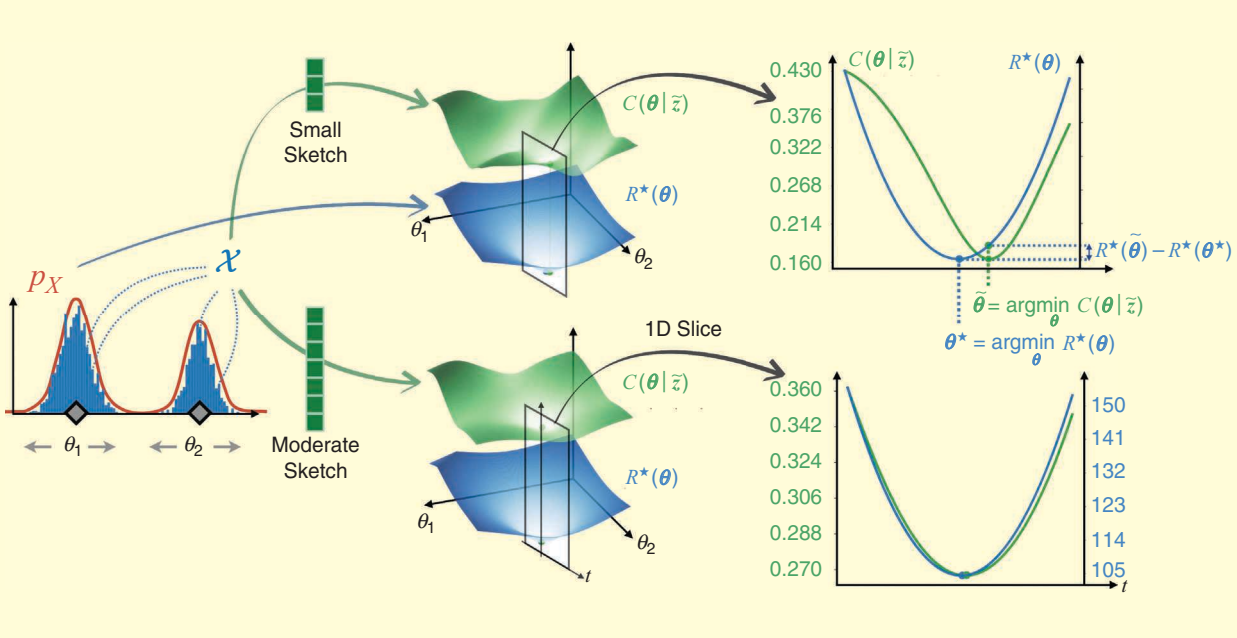

FIGURE S6. The k-means clustering of 1D data $\mathcal{X}$, showing cost function $C(\boldsymbol{\theta}|\tilde{\boldsymbol{z}})$, risk $R^\star(\boldsymbol{\theta})$, ideal estimate $\boldsymbol{\theta}^\star$, compressive learning estimate $\tilde{\boldsymbol{\theta}}$, and excess risk $R^\star(\tilde{\boldsymbol{\theta}}) - R^\star(\boldsymbol{\theta}^\star)$, for a sketch of a small-dimension m (top) and for a sketch of a moderate-size m (bottom).

Similarly, excess risk bounds for compressive learning are achieved by interpreting the minimizer $\tilde{\boldsymbol{\theta}}$ of (19) as the minimizer of the risk $R^*(\boldsymbol{\theta} \,|\, \tilde{p})$ for some "simple" probability distribution $\tilde{p}$ that minimizes $\| \tilde{\boldsymbol{z}} - \mathcal{A}(\tilde{p}) \|$ 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 Figure 4.

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 $\boldsymbol{x}_0 \in \mathbb{R}^d$ given the noisy linear measurement $\boldsymbol{y} = \boldsymbol{A}\boldsymbol{x}_0 + \boldsymbol{e} \in \mathbb{R}^m$, where $m \ll d$. Ideally, recovery aims to find the k-sparse vector whose (noiseless) measurement is closest to the observed $\boldsymbol{y}$; i.e.,

$$\hat{\boldsymbol{x}} = \operatorname*{argmin}_{\boldsymbol{x} \in \Sigma_k} \| \boldsymbol{y} - \boldsymbol{A}\boldsymbol{x} \|^2. \tag{25}$$

In practice, convex or greedy approximations of this combinatorial approach are often used. Recovery guarantees can be established using the RIP in (S1), which says that the Euclidean distance $\| \boldsymbol{A}\boldsymbol{x} - \boldsymbol{A}\boldsymbol{x}' \|$ between the (noiseless) measurements of two k-sparse vectors is almost the same as the Euclidean distance $\| \boldsymbol{x} - \boldsymbol{x}' \|$ between the vectors themselves (see the sidebar "Compressive Sensing and the Restricted Isometry Property"). 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 $p_{\boldsymbol{\theta}} \in \Sigma_{\boldsymbol{\Theta}}$ from the sketch $\tilde{\boldsymbol{z}}$ rather than recovering the model parameters $\boldsymbol{\theta} \in \boldsymbol{\Theta}$ themselves. Here, $\Sigma_{\boldsymbol{\Theta}} := \{p_{\boldsymbol{\theta}} : \boldsymbol{\theta} \in \boldsymbol{\Theta}\}$ denotes some set of admissible distributions $p_{\boldsymbol{\theta}}$. For example, in the compressive GMM, the goal becomes that of recovering a k-term GMM $p_{\boldsymbol{\theta}} = \Sigma_{\ell=1}^{k} \alpha_\ell \mathcal{N}(\boldsymbol{\mu}_\ell, \boldsymbol{\Sigma}_\ell)$ rather than the GMM parameters $\boldsymbol{\theta} = \{\alpha_\ell, \boldsymbol{\mu}_\ell, \boldsymbol{\Sigma}_\ell\}_{\ell=1}^{k}$. Likewise, in compressive PCA, the goal becomes that of recovering a distribution $p_{\boldsymbol{\theta}}$ whose autocorrelation matrix has a rank of at most k, rather than a k-dimensional orthonormal basis $\boldsymbol{\theta}$ for the column space of that autocorrelation matrix. Then, mirroring the CS recovery approach (25), compressive learning recovery aims to find a parametric distribution $p_{\boldsymbol{\theta}} \in \Sigma_{\boldsymbol{\Theta}}$ whose (noiseless) sketch is closest to the observed sketch $\tilde{\boldsymbol{z}}$; i.e.,

$$\tilde{p} = p_{\tilde{\boldsymbol{\theta}}} := \operatorname*{argmin}_{p_{\boldsymbol{\theta}} \in \Sigma_{\boldsymbol{\Theta}}} \| \tilde{\boldsymbol{z}} - \mathcal{A}(p_{\boldsymbol{\theta}}) \|^2. \tag{26}$$

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

$$\tilde{\boldsymbol{\theta}} = \operatorname*{argmin}_{\boldsymbol{\theta} \in \boldsymbol{\Theta}} C(\boldsymbol{\theta} \,|\, \tilde{\boldsymbol{z}}) \quad \text{with} \quad C(\boldsymbol{\theta} \,|\, \tilde{\boldsymbol{z}}) = \| \tilde{\boldsymbol{z}} - \mathcal{A}(p_{\boldsymbol{\theta}}) \|^2. \tag{27}$$

Note the similarity between (27) and (13) and (14). Moreover, $\tilde{\boldsymbol{\theta}}$ also minimizes the risk $R^*(\boldsymbol{\theta} \,|\, \tilde{p})$ [5], [51].

From the preceding description, 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., $\boldsymbol{x}_0$ 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 $\boldsymbol{x}_0$ rather than $\boldsymbol{x}_0$ itself. In compressive learning, this is accomplished by seeking to recover the best model distribution $p_{\boldsymbol{\theta}} \in \Sigma_{\boldsymbol{\Theta}}$ rather than the true distribution p_X. In both cases, if the linear operator [i.e., $\boldsymbol{A}$ in CS or $\mathcal{A}(\cdot)$ in compressive learning] is well designed, then the model parameters (in Σ_k for CS or in $\Sigma_{\boldsymbol{\Theta}}$ for compressive learning) can be accurately recovered.

The LRIP and excess risk control

Recovery guarantees can be established using a tool analogous to the RIP (S1) in CS.

Take, for example, the case of compressive PCA. As discussed around (18), one could use a sketch of the form $\tilde{\boldsymbol{z}} = \boldsymbol{A}\mathrm{vec}(\hat{\boldsymbol{R}}) \in \mathbb{R}^m$, where $\hat{\boldsymbol{R}} = (1/n)\Sigma_{i=1}^{n} \boldsymbol{x}_i \boldsymbol{x}_i^\top$ is the $d \times d$ empirical autocorrelation matrix and $m \ll d^2$, and then search for $\tilde{\boldsymbol{R}}$, the symmetric matrix with a rank of at most k that minimizes the cost $C(\boldsymbol{R} \,|\, \tilde{\boldsymbol{z}}) = \| \tilde{\boldsymbol{z}} - \boldsymbol{A}\mathrm{vec}(\boldsymbol{R}) \|^2$. Recovery guarantees can be established using a RIP of the form

$$\begin{aligned}(1-\delta)\| \boldsymbol{R} - \boldsymbol{R}' \|_F^2 \le \| \boldsymbol{A}\mathrm{vec}(\boldsymbol{R}) - \boldsymbol{A}\mathrm{vec}(\boldsymbol{R}') \|^2 \\ \le (1+\delta)\| \boldsymbol{R} - \boldsymbol{R}' \|_F^2, \quad \forall \boldsymbol{R}, \boldsymbol{R}' \in \Sigma_k,\end{aligned} \tag{28}$$

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, Ch. 9] that the i.i.d. Gaussian $\boldsymbol{A}$ satisfies the RIP (28) with high probability when m is on the order of kd times a constant that depends on the tolerance δ. Furthermore, the cost-minimizing $\tilde{\boldsymbol{R}}$ is near optimal in the sense of minimizing the excess risk, even when a convex relaxation of the cost is used [5], [51].

These compressive PCA guarantees hold even under model mismatch: although Σ_k constrains $\tilde{\boldsymbol{R}}$ to have a rank of at most k, the empirical autocorrelation matrix $\hat{\boldsymbol{R}}$ used to construct $\tilde{\boldsymbol{z}}$ tends to have a full-rank d in practice. To explore this idea in more detail, suppose for the moment that the data $\mathcal{X}$ lie in a rank-k subspace. In this case, $\hat{\boldsymbol{R}}$ would have a rank of at most k, and there exists a symmetric $\boldsymbol{R}$ with a rank of at most k that drives the cost $C(\boldsymbol{R} \,|\, \tilde{\boldsymbol{z}}) = \| \boldsymbol{A}\mathrm{vec}(\hat{\boldsymbol{R}}) - \boldsymbol{A}\mathrm{vec}(\boldsymbol{R}) \|^2$ to zero. Furthermore, when $\boldsymbol{A}$ satisfies the RIP (28), the left inequality in (28) implies that this cost-minimizing $\tilde{\boldsymbol{R}}$ must equal $\hat{\boldsymbol{R}}$. When the data $\mathcal{X}$ do not lie in a rank-k subspace, the minimal cost will be nonzero. In this case, an upper bound on the error $\| \hat{\boldsymbol{R}} - \tilde{\boldsymbol{R}} \|_F$ of the cost-minimizing estimate $\tilde{\boldsymbol{R}}$ as well as an upper bound on the excess risk of the principal subspace $\tilde{\boldsymbol{\theta}}$ of $\tilde{\boldsymbol{R}}$ can both be obtained as a consequence of the RIP (28).

In place of the RIP, compressive learning uses the so-called LRIP,

$$d(p_{\boldsymbol{\theta}}, p_{\boldsymbol{\theta}'}) \le C_0 \| \mathcal{A}(p_{\boldsymbol{\theta}}) - \mathcal{A}(p_{\boldsymbol{\theta}'}) \|, \tag{29}$$

assumed to be valid for each pair $p_{\boldsymbol{\theta}}, p_{\boldsymbol{\theta}'} \in \Sigma_{\boldsymbol{\Theta}}$, where C_0 is a positive constant. The LRIP says that the Euclidean distance $\| \mathcal{A}(p_{\boldsymbol{\theta}}) - \mathcal{A}(p_{\boldsymbol{\theta}'}) \|$ between the (noiseless) sketches of two

distributions is—up to a scaling—controlling the distance $d(p_{\boldsymbol{\theta}}, p_{\boldsymbol{\theta}'})$ between the distributions themselves (22). Note that the lower bound $d(\cdot,\cdot)$ in (29) is task specific and not Euclidean as in the CS case (S1).

The first main theoretical guarantee about compressive learning [5], [51] is that when the LRIP (29) holds, the estimate $\tilde{\boldsymbol{\theta}}$ obtained by minimizing (27) automatically has controlled excess risk. In particular, using $\tilde{\boldsymbol{\theta}} = \operatorname{argmin}_{\boldsymbol{\theta}} R^*(\theta \,|\, p_{\tilde{\boldsymbol{\theta}}})$ [5], [51], the excess risk bound (24) can be combined with the LRIP (29) and the definition of $\tilde{\boldsymbol{\theta}}$ (27) as follows:

$$\begin{aligned} R^*(\tilde{\boldsymbol{\theta}} \,|\, p_X) - R^*(\boldsymbol{\theta}^* \,|\, p_X) &\le 2d(p_{\tilde{\boldsymbol{\theta}}}, p_X) \\ &\le 2[d(p_{\tilde{\boldsymbol{\theta}}}, p_{\boldsymbol{\theta}'}) + d(p_{\boldsymbol{\theta}'}, p_X)] \\ &\le 2[C_0 \| \mathcal{A}(p_{\tilde{\boldsymbol{\theta}}}) - \mathcal{A}(p_{\boldsymbol{\theta}'}) \| + d(p_{\boldsymbol{\theta}'}, p_X)] \\ &\le 2[C_0 \| \mathcal{A}(p_{\tilde{\boldsymbol{\theta}}}) - \tilde{\boldsymbol{z}} \| + C_0 \| \tilde{\boldsymbol{z}} - \mathcal{A}(p_{\boldsymbol{\theta}'}) \| + d(p_{\boldsymbol{\theta}'}, p_X)] \\ &\le 2[2C_0 \| \mathcal{A}(p_{\boldsymbol{\theta}'}) - \tilde{\boldsymbol{z}} \| + d(p_{\boldsymbol{\theta}'}, p_X)] \\ &\le 2[2C_0 \| \mathcal{A}(p_X) - \tilde{\boldsymbol{z}} \| + 2C_0 \| \mathcal{A}(p_{\boldsymbol{\theta}'}) - \mathcal{A}(p_X) \| \\ &\qquad + d(p_{\boldsymbol{\theta}'}, p_X)], \end{aligned} \tag{30}$$

for any distribution $p_{\boldsymbol{\theta}'}$ in $\Sigma_{\boldsymbol{\Theta}}$ [because $d(\cdot,\cdot)$ is a valid distance metric; i.e., $d(p_1, p_2) \le d(p_1, p_3) + d(p_3, p_2)\ \forall p_3$]. One option is to choose $p_{\boldsymbol{\theta}'} = p_{\boldsymbol{\theta}^*}$, in which case the term $2C_0 \| \mathcal{A}(p_{\boldsymbol{\theta}^*}) - \mathcal{A}(p_X) \| + d(p_{\boldsymbol{\theta}^*}, p_X)$ in the upper bound reflects the excess risk due to modeling and $2C_0 \| \mathcal{A}(p_X) - \tilde{\boldsymbol{z}} \|$ reflects the excess risk due to sketching from a finite data set $\mathcal{X}$. For a tighter bound, we could choose $p_{\boldsymbol{\theta}'}$ as the distribution in $\Sigma_{\boldsymbol{\Theta}}$ that minimizes the right side of (30). These approaches and more refined variants have been applied to analyze various learning tasks in, e.g., [5], [11], [12], and [51].

It should be emphasized that recovery guarantees based on the RIP (S1) or LRIP (29) hold even in the presence of measurement noise $\boldsymbol{e}$ and/or modeling errors. In CS, measurement noise arises due to, e.g., thermal noise or interference, while modeling errors arise when $\boldsymbol{x}_0$ is not truly k-sparse. Many sparse recovery techniques are provably robust to such noise and modeling errors; see (S2) in the sidebar "Compressive Sensing and the Restricted Isometry Property." In compressive learning, measurement noise arises due to the finite cardinality of the data set $\mathcal{X}$ [recall (7)], while modeling errors arise when $p_X \notin \Sigma_{\boldsymbol{\Theta}}$. For example, GMM recovery guarantees can be established even when p_X is not truly a GMM.

Establishing the LRIP via kernels and the JL lemma

In light of the fact that the LRIP (29) yields statistical learning guarantees (30) for compressive learning, a key question becomes: How can we choose the sketch dimension m so that the LRIP (29) holds? Similar to how the RIP is proven in CS (see, e.g., [4, Lemma 9.33]), refinements of a mathematical tool called the *JL lemma* [31] can be used to obtain a value of m

Kernel Methods and Kernel Embeddings of Probability Distributions

Sketching shares connections with kernel methods [3], a family of machine learning techniques that produces decisions or insights using a kernel function, $\kappa(\boldsymbol{x}, \boldsymbol{x}') \in \mathbb{R}$, which measures the "similarity" between $\boldsymbol{x}$ and $\boldsymbol{x}'$. A kernel is said to be "positive definite" if the $n \times n$ matrix $\boldsymbol{K}$, constructed with entries $\kappa(\boldsymbol{x}_i, \boldsymbol{x}_j)$ for $1 \le i, j \le n$, is positive semidefinite for every possible $\{\boldsymbol{x}_i\}_{i=1}^n$. 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, $\kappa(\boldsymbol{x}, \boldsymbol{x}') = \langle \bar{\boldsymbol{\Phi}}(\boldsymbol{x}), \bar{\boldsymbol{\Phi}}(\boldsymbol{x}') \rangle$ for some (not necessarily explicitly known) mapping $\boldsymbol{\Phi}(\cdot)$ 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 [29], and dictionary learning [S3]—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 $\kappa(\boldsymbol{x}, \boldsymbol{x}')$ operating on signals $\boldsymbol{x}$ and $\boldsymbol{x}'$ 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

$$k(p, q) := \mathbb{E}_{X \sim p, X' \sim q}[\kappa(X, X')]. \tag{S5}$$

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 (6). The maximum mean discrepancy (MMD) [S4], [S5]

$$\mathrm{MMD}(p, q) := \sqrt{k(p, p) + k(q, q) - 2k(p, q)} \tag{S6}$$

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 [S4], is now well known in machine learning. When the MMD behaves as a true metric, i.e., when $\mathrm{MMD}(p, q) = 0 \Leftrightarrow p = q$, 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 [S5].

References

[S3] R. Rubinstein, A. M. Bruckstein, and M. Elad, "Dictionaries for sparse representation modeling," *Proc. IEEE*, vol. 98, no. 6, pp. 1045–1057, 2010. doi: 10.1109/JPROC.2010.2040551.

[S4] A. Gretton, K. M. Borgwardt, M. J. Rasch, B. Schölkopf, and A. J. Smola, "A Kernel method for the two-sample problem," in *Proc. Adv. Neural Inform. Process. Syst. (NIPS)*, 2007.

[S5] B. K. Sriperumbudur, A. Gretton, K. Fukumizu, B. Schölkopf, and G. R. Lanckriet, "Hilbert space embeddings and metrics on probability measures," *J. Mach. Learning Res.*, vol. 11, no. 10, pp. 1517–1561, Mar. 2010.

sufficient for the LRIP to hold with high probability on the draw of the random feature map $\boldsymbol{\Phi}(\cdot)$. We now summarize this approach.

To begin, we establish connections to kernel methods (see the sidebar "Kernel Methods and Kernel Embeddings of Probability Distributions" 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 $\{\boldsymbol{w}_j\}_{j=1}^m$, we have, for all $\boldsymbol{x}, \boldsymbol{x}'$,

$$\left\langle \frac{1}{\sqrt{m}}\boldsymbol{\Phi}(\boldsymbol{x}), \frac{1}{\sqrt{m}}\boldsymbol{\Phi}(\boldsymbol{x}') \right\rangle = \frac{1}{m}\sum_{j=1}^{m} \exp(-\mathrm{j}2\pi \boldsymbol{w}_j^{\top}(\boldsymbol{x}-\boldsymbol{x}')), \tag{31}$$

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 the sidebar "Kernel Methods and Kernel Embeddings of Probability Distributions." This kernel is also "shift-invariant" in that it depends only on the difference $\boldsymbol{x}-\boldsymbol{x}'$. Next, imagine that the d-dimensional frequency vectors $\boldsymbol{w}_j$, which constitute the m rows of the random matrix $\boldsymbol{W}$, 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 (31) says

$$\left\langle \frac{1}{\sqrt{m}}\boldsymbol{\Phi}(\boldsymbol{x}), \frac{1}{\sqrt{m}}\boldsymbol{\Phi}(\boldsymbol{x}') \right\rangle \overset{\text{w.p.1}}{\to} \mathbb{E}_W \exp(-\mathrm{j}2\pi W^{\top}(\boldsymbol{x}-\boldsymbol{x}')). \tag{32}$$

The right-hand side of (32) is the CFT of the probability density p_W evaluated at $\boldsymbol{x}-\boldsymbol{x}'$, i.e., the characteristic function $\boldsymbol{\Psi}_{p_W}(\boldsymbol{x}'-\boldsymbol{x})$ using the notation from (12). Because $\boldsymbol{\Psi}_{p_W}(\cdot)$ is the CFT of a nonnegative function, it is a so-called positive definite function, which means that for any $\{\boldsymbol{x}_i\}_{i=1}^n$, the $n \times n$ matrix $\boldsymbol{\Psi}$ defined with elements $\boldsymbol{\Psi}_{ij} = \boldsymbol{\Psi}_{p_W}(\boldsymbol{x}_i - \boldsymbol{x}_j)$ for $1 \le i, j \le n$ will be positive semidefinite. Thus, if we construct a kernel as $\kappa(\boldsymbol{x}, \boldsymbol{x}') := \boldsymbol{\Psi}_{p_W}(\boldsymbol{x}'-\boldsymbol{x})$, then it will be a positive definite kernel. When $p_W = \mathcal{N}(\boldsymbol{0}, \sigma_w^2 \boldsymbol{I}_d)$, this approach yields the familiar Gaussian kernel (a particular type of "radial basis function"); i.e., $\kappa_\sigma(\boldsymbol{x}, \boldsymbol{x}') := \exp(-\|\boldsymbol{x}-\boldsymbol{x}'\|^2/2\sigma^2)$, here of width $\sigma = 1/\sigma_w$.

More generally, by considering any parametric feature map of the form $\boldsymbol{\Phi}(\boldsymbol{x}|\boldsymbol{W})$, where the parameter $\boldsymbol{W}$ is drawn at random according to some probability distribution, one can define [in many papers, the $1/\sqrt{m}$ scaling in (31)–(33) is subsumed in the feature map $\boldsymbol{\Phi}(\cdot)$] the "expected kernel"

$$\kappa(\boldsymbol{x},\boldsymbol{x}') := \mathbb{E}_W \left\langle \frac{1}{\sqrt{m}}\boldsymbol{\Phi}(\boldsymbol{x}|\boldsymbol{W}), \frac{1}{\sqrt{m}}\boldsymbol{\Phi}(\boldsymbol{x}'|\boldsymbol{W}) \right\rangle, \tag{33}$$

not to be confused with the "mean kernel" defined in (S5) (see the sidebar "Kernel Methods and Kernel Embeddings of Probability Distributions"). This setting includes RF features (3) with i.i.d. frequencies $\boldsymbol{w}_j$ (as in the preceding) or with a frequency matrix $\boldsymbol{W}$ that includes structured blocks of rows, as we will soon discuss.

Now that the kernel connections have been established, we return to our original objective of understanding when the LRIP (29) holds. Importantly, the law of large numbers shows that, in the limit of large sketch dimension m, the right side of (29) is related to the maximum mean discrepancy (MMD) from (S6), which is a kernel-based distance between distributions (see the sidebar "Kernel Methods and Kernel Embeddings of Probability Distributions"). Indeed, for arbitrary probability distributions p and q, we have

$$\lim_{m\to\infty} \frac{1}{\sqrt{m}} \left\| \mathcal{A}(p) - \mathcal{A}(q) \right\| \overset{\text{w.p.1}}{=} \mathrm{MMD}(p,q), \tag{34}$$

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 (29) holds, it must also be true that, for a sufficiently large sketch dimension m,

$$d(p_{\boldsymbol{\theta}}, p_{\boldsymbol{\theta}'}) \le C_0' \mathrm{MMD}(p_{\boldsymbol{\theta}}, p_{\boldsymbol{\theta}'}), \quad \text{for all } p_{\boldsymbol{\theta}}, p_{\boldsymbol{\theta}'} \in \Sigma_{\boldsymbol{\Theta}}, \tag{35}$$

where C_0' is a positive constant. Property (35) connects two different metrics on probability distributions: the left-hand side of (35) is defined by the learning task [recall (22)], while the right-hand side of (35) is defined by the expected kernel κ, from (33), associated with the randomized feature map. Note that (35) is a deterministic property; it does not depend on the draw of the randomized feature map, unlike the LRIP (29). In the literature, (35) is called the *kernel LRIP* [5], [51] 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 it eases the proof of the kernel LRIP. This is important because, if the kernel LRIP (35) holds, then, using arguments based on the JL lemma, one can also establish [5], [51] the LRIP (29), as we show in the following.

The JL lemma [31] is a precursor of the RIP that is specialized to finite sets (S1). It states that given N arbitrary d-dimensional vectors $\boldsymbol{x}_i$, there exists an $m \times d$ matrix $\boldsymbol{A}$, with m on the order of $\log N$, such that $\|\boldsymbol{A}\boldsymbol{x}_i - \boldsymbol{A}\boldsymbol{x}_j\| \approx \|\boldsymbol{x}_i - \boldsymbol{x}_j\|$ 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 $\Sigma_{\boldsymbol{\Theta}}$ 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 in the following. As a concrete example, let us consider a discretized variant of the compressive GMM using the RF feature map (3). As in [5], [6], and [51], we assume that the d-dimensional frequency vectors $\boldsymbol{w}_j$ are drawn i.i.d. from the normal distribution $\mathcal{N}(\boldsymbol{0}, \sigma_w^2 \boldsymbol{I}_d)$, and we consider learning a mixture of Gaussian components $p_{\boldsymbol{\theta}_\ell} = \mathcal{N}(\boldsymbol{\mu}_\ell, \boldsymbol{I})$ for $\ell = 1, \ldots, k$, 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 $p_{\boldsymbol{\theta}} \in \Sigma_{\boldsymbol{\Theta}}$. As long as the MMD

is a true metric (as defined in the sidebar "Kernel Methods and Kernel Embeddings of Probability Distributions"), the ratio $d(p_{\boldsymbol{\theta}},p_{\boldsymbol{\theta}'})/\mathrm{MMD}(p_{\boldsymbol{\theta}},p_{\boldsymbol{\theta}'})$ is finite for $p_{\boldsymbol{\theta}} \neq p_{\boldsymbol{\theta}'}$; hence, whenever $\Sigma_{\boldsymbol{\Theta}}$ is a finite collection, as in this concrete example, the kernel LRIP (35) holds (for a sufficiently large C_0'). Moreover, specializing to an arbitrary pair of mixtures $p_{\boldsymbol{\theta}},p_{\boldsymbol{\theta}'} \in \Sigma_{\boldsymbol{\Theta}}$ and a given sketch dimension m, refinements of (34) (using measure concentration) enable one to show that the lower bound

$$\frac{1}{\sqrt{2}}\mathrm{MMD}(p_{\boldsymbol{\theta}},p_{\boldsymbol{\theta}'}) \leq \frac{1}{\sqrt{m}}\left\| \mathcal{A}(p_{\boldsymbol{\theta}}) - \mathcal{A}(p_{\boldsymbol{\theta}'}) \right\| \tag{36}$$

holds with a probability of 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 $p_{\boldsymbol{\theta}},p_{\boldsymbol{\theta}'} \in \Sigma_{\boldsymbol{\Theta}}$, the lower bound (36) is valid uniformly for all of them, except with a probability of at most $N^2\exp(-c_0 m)$. This failure probability can be made smaller than any $\epsilon > 0$ by choosing an m larger than $2c_0^{-1}\log(N/\sqrt{\epsilon})$. Combining (35) and (36) using a union bound, it can be shown [5], [51] that the LRIP (29) 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 $\Sigma_{\boldsymbol{\Theta}}$.

For infinite collections $\Sigma_{\boldsymbol{\Theta}}$, proving the kernel LRIP (35), and eventually the LRIP (29), 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 (35) holds, provided that the centroids are sufficiently separated and bounded [5], [51]. The JL lemma can be extended by refining techniques used to establish the RIP (S1). The main idea is to first prove the LRIP on some finite collection $\Sigma' \subset \Sigma_{\boldsymbol{\Theta}}$ of N probability distributions and then to extrapolate it (with slightly worse constants) to $\Sigma_{\boldsymbol{\Theta}}$. Technically, this involves the notion of covering numbers, and the cardinality $N = |\Sigma'|$ is typically exponential in the number of parameters needed to describe $\Sigma_{\boldsymbol{\Theta}}$. For example, for the GMM, one needs kd parameters to describe the means $\boldsymbol{\mu}_\ell \in \mathbb{R}^d, 1 \leq \ell \leq k$ of the mixture $p_{\boldsymbol{\theta}} = \sum_{\ell=1}^{k}(1/k)\mathcal{N}(\boldsymbol{\mu}_\ell,\boldsymbol{I})$, 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], [51].

Empirical studies of compressive clustering [9], [10] and the compressive GMM [6] 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], [51] remain pessimistic compared to these empirically validated sketch dimensions. This is most likely related to suboptimal 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 $\Sigma_{\boldsymbol{\Theta}}$.

The challenge of designing a feature map given a learning task

In the existing literature, the LRIP has been established [5], [51] for randomized feature maps $\boldsymbol{\Phi}(\cdot)$ (e.g., RF features and random quadratic features) that mimic related constructions from CS, developed either for sparse vector recovery or low-rank matrix recovery.

When sketching with RF features (e.g., for compressive clustering and the compressive GMM), the main design choice for $\boldsymbol{\Phi}(\cdot)$ is the distribution from which to draw the random frequencies $\boldsymbol{w}_j$ [i.e., the rows of $\boldsymbol{W}$ in (3)]. In light of the connections to shift-invariant kernels [recall (31)], 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 $\boldsymbol{W}$ are drawn i.i.d. $\mathcal{N}(\boldsymbol{0}, \sigma_w^2\boldsymbol{I})$, 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 (S5)] acts to low-pass-filter the underlying data distributions. To see why, observe that $\kappa_\sigma(\boldsymbol{x},\boldsymbol{x}') = g_\sigma(\boldsymbol{x},\boldsymbol{x}')$ with $g_\sigma(\boldsymbol{x}) \coloneqq e^{-\|\boldsymbol{x}\|^2/2\sigma^2}$, and so

$$\begin{aligned} k_\sigma(p,q) &= \mathbb{E}_{X\sim p, X'\sim q}[\kappa_\sigma(X,X')] \\ &= \iint e^{-\|X-X'\|^2/2\sigma^2}\,\mathrm{d}p(X)\,\mathrm{d}q(X') \qquad (37) \\ &= \langle\, g_\sigma \star p, q \,\rangle_{L^2} = \langle\, g_{\bar{\sigma}} \star p,\, g_{\bar{\sigma}} \star q \,\rangle_{L^2}, \qquad (38) \end{aligned}$$

where $\star$ denotes convolution and $\bar{\sigma} = \sigma/\sqrt{2}$. Hence, the associated MMD (S6) satisfies $\mathrm{MMD}(p,q) = \| g_{\bar{\sigma}} \star p - g_{\bar{\sigma}} \star q \|_{L^2}$. Recall that learning from a sketch is often performed by minimizing a "sketch matching" cost $\|\tilde{\boldsymbol{z}} - \mathcal{A}(p_{\boldsymbol{\theta}})\|^2 = \|\mathcal{A}(\hat{p}_{\mathcal{X}}) - \mathcal{A}(p_{\boldsymbol{\theta}})\|^2$, as in (26), 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 $p_{\boldsymbol{\theta}}$ since $(1/m)\|\tilde{\boldsymbol{z}} - \mathcal{A}(p_{\boldsymbol{\theta}})\|^2 \approx \|g_{\bar{\sigma}} \star \hat{p}_{\mathcal{X}} - g_{\bar{\sigma}} \star p_{\boldsymbol{\theta}}\|_{L^2}^2$. Similarly, when using a greedy algorithm to learn a mixture model (or cluster centroids) from a sketch, the normalized inner product $(1/m)\langle\, \tilde{\boldsymbol{z}}, \mathcal{A}(p_{\boldsymbol{\theta}_\ell}) \,\rangle$ approximates the correlation $\langle\, g_{\bar{\sigma}} \star \hat{p}_{\mathcal{X}}, g_{\bar{\sigma}} \star p_{\boldsymbol{\theta}_\ell} \,\rangle_{L^2}$ between the low-passed versions of the empirical data distribution $\hat{p}_{\mathcal{X}}$ and the candidate mixture component $p_{\boldsymbol{\theta}_\ell}$, respectively.

This latter idea is illustrated for compressive clustering in Figure 5. There, since the mixture component associated to a candidate centroid $\boldsymbol{c}$ is the Dirac $p_{\boldsymbol{c}}(\boldsymbol{x}) = \delta(\boldsymbol{x}-\boldsymbol{c})$ [recall the discussion before (15)], we have that $\langle \tilde{\boldsymbol{z}}_\sigma, \mathcal{A}_\sigma(p_{\boldsymbol{c}}) \rangle = \langle \tilde{\boldsymbol{z}}_\sigma, \boldsymbol{\Phi}_\sigma(\boldsymbol{c}) \rangle$, where the dependence on the kernel width $\sigma = 1/\sigma_w$ has been made explicit. Meanwhile,

$$\langle\, g_{\bar{\sigma}} \star \hat{p}_{\mathcal{X}}, g_{\bar{\sigma}} \star p_{\boldsymbol{c}} \,\rangle_{L^2} = (g_\sigma \star \hat{p}_{\mathcal{X}})(\boldsymbol{c}) = \frac{1}{n}\sum_{i=1}^{n} g_\sigma(\boldsymbol{c}-\boldsymbol{x}_i). \tag{39}$$

In (39), we recognize a Parzen window density estimator [32], whose computation for a given $\boldsymbol{c}$ requires access to the n training samples $\boldsymbol{x}_i$. In contrast, its surrogate $\langle \tilde{\boldsymbol{z}}_\sigma, \boldsymbol{\Phi}_\sigma(\boldsymbol{c}) \rangle$ only requires access to the m-dimensional sketch $\tilde{\boldsymbol{z}}$. 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 5, $\langle \tilde{\boldsymbol{z}}_\sigma, \boldsymbol{\Phi}_\sigma(\boldsymbol{c}) \rangle$ well approximates $(g_\sigma \star \hat{p}_{\mathcal{X}})(\boldsymbol{c})$ for a sufficiently large kernel width σ. By comparing the different columns of Figure 5, 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 oversmoothing the (desired) gaps between clusters. Existing theory [5], [51] 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 $\boldsymbol{w}_j$ were used with the RF map in the preceding, one may also consider the use of i.i.d. non-Gaussian frequencies. Such designs, as proposed in [7], 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 of improving computational efficiency. Such constructions also yield non-Gaussian expected kernels (33). 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 (29) for a given learning task defined by a loss function $L(\boldsymbol{\theta}\,|\,\boldsymbol{x}_i)$ and embodied by a task-driven distance (22)? A promising yet still challenging avenue would be to first identify a positive definite kernel $\kappa_0(\boldsymbol{x},\boldsymbol{x}')$ for which the corresponding MMD satisfies the kernel LRIP (35) and then use Bochner's theorem or Mercer's theorem (see the sidebar "Approximating Kernel Methods With Random Feature Maps") to design a random feature map $\boldsymbol{\Phi}(\cdot)$ whose expected kernel (33) is precisely κ_0.

Sketching with reduced computational resources

The computational cost of sketching via (1) is heavily dependent on the feature map $\boldsymbol{\Phi}(\cdot)$. The computational cost of parameter estimation via (13) is also heavily dependent on $\boldsymbol{\Phi}(\cdot)$, since it often involves iterative application of $\boldsymbol{\Phi}(\cdot)$.

Often, the feature map is constructed as a randomized linear operation followed by a componentwise nonlinear operation; i.e.,

$$\boldsymbol{\Phi}(\boldsymbol{x}) = \varrho(\boldsymbol{W}\boldsymbol{x}), \tag{40}$$

where $\boldsymbol{W}$ is a (randomly drawn) matrix of size $m \times d$ and $\varrho(\cdot)$ applies a scalar nonlinear function identically to each element of the vector $\boldsymbol{W}\boldsymbol{x}$. For example, the feature maps described earlier for compressive PCA, the 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 $\boldsymbol{W}$ 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 $\varrho_{\mathrm{RF}}(\cdot) := \exp(-\mathrm{j}2\pi\cdot)$ used in the RF map (3). We discuss such constructions of $\boldsymbol{W}$ and $\varrho(\cdot)$ in the following.

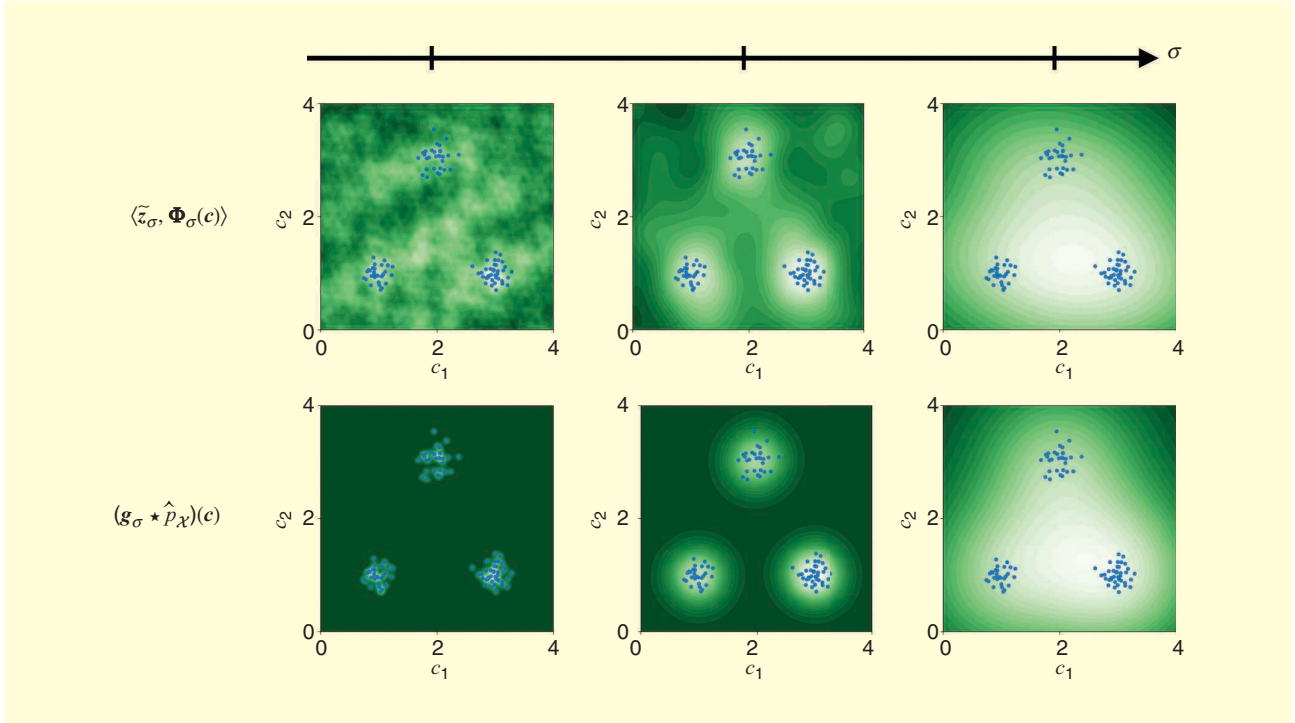

FIGURE 5. The criterion $\langle \tilde{\boldsymbol{z}}_\sigma, \boldsymbol{\Phi}_\sigma(\boldsymbol{c})\rangle$ used by greedy parameter estimation algorithms in compressive clustering with RF features (top row) versus its expected value $(g_\sigma \star \hat{p}_{\mathcal{X}})(\boldsymbol{c})$ (bottom row) as a function of the centroid hypothesis location $\boldsymbol{c} = [c_1, c_2]^\top$. The data set (in blue) consists of $n = 100$ points drawn according to a mixture of $k = 3$ isotropic Gaussians. The frequencies $\boldsymbol{w}_j$ used to define the feature map $\boldsymbol{\Phi}_\sigma(\cdot) = \exp(-\mathrm{j}2\pi\boldsymbol{W}\cdot)$ are drawn according to a standard Gaussian $\mathcal{N}(\boldsymbol{0}, \sigma_w^2\boldsymbol{I})$ with $\sigma_w = 1/\sigma$. The sketch $\tilde{\boldsymbol{z}}_\sigma$ is computed with the feature map $\boldsymbol{\Phi}_\sigma(\cdot)$. With $\sigma^2 = 1/500$ (left), there is insufficient smoothing, and the criterion displays many spurious local maxima; with $\sigma^2 = 1/10$ (middle), there is appropriate smoothing, and local maxima are in good correspondence with the true cluster centers; with $\sigma^2 = 1$ (right), there is oversmoothing, and the criterion displays only a single maximum.

Sketching with structured random matrices

In most of our previous examples, we constructed $\boldsymbol{W}$ by drawing its m rows i.i.d. from the normal distribution $\mathcal{N}(\boldsymbol{0}, \sigma_w^2 \boldsymbol{I}_d)$ with some variance $\sigma_w^2 > 0$. Recall that, with the RF map $\varrho_{\mathrm{RF}}(\cdot)$, the rows of $\boldsymbol{W}$ correspond to the (d-dimensional) frequencies used when sampling the CFT of the (empirical) data distribution. In any case, when $\boldsymbol{W}$ is an explicit matrix, the computational complexity of computing $\boldsymbol{\Phi}(\boldsymbol{x})$ of the form (40) is dominated by the matrix–vector product $\boldsymbol{Wx}$. Thus, it is on the order of md, which is also the order of the memory needed to store $\boldsymbol{W}$.

As an alternative to these approaches, it has been suggested to construct $\boldsymbol{W}$ 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 [33], which was successfully applied to compressive learning in [34]. There, the idea is to construct $\boldsymbol{W}$ as a vertical concatenation of $b = \lceil m/d \rceil$ blocks $\{\boldsymbol{B}_j\}_{j=1}^{b}$, each of size $d \times d$. These blocks have the form $\boldsymbol{B}_j = \boldsymbol{D}_j^{(0)} \boldsymbol{H} \boldsymbol{D}_j^{(1)} \boldsymbol{H} \boldsymbol{D}_j^{(2)} \boldsymbol{H} \boldsymbol{D}_j^{(3)}$, where $\boldsymbol{H}$ is the Walsh–Hadamard matrix and $\boldsymbol{D}_j^{(k)}$ are random diagonal matrices. In particular, the diagonal elements of $\boldsymbol{D}_j^{(1)}$, $\boldsymbol{D}_j^{(2)}$, and $\boldsymbol{D}_j^{(3)}$ are drawn i.i.d. from the uniform distribution over $\{-1, 1\}$ and the diagonal elements of $\boldsymbol{D}_j^{(0)}$ are drawn i.i.d. from the $\mathcal{X}$ distribution with d degrees of freedom, which is the distribution of the norm of a d-variate Gaussian vector. This construction is depicted in Figure 6. The fast Walsh–Hadamard transform offers an order $d \log d$-complexity implementation of the matrix–vector multiplication $\boldsymbol{Hx}$ and prevents the need to explicitly store $\boldsymbol{H}$. With this structured and fast incarnation of $\boldsymbol{W}$, 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.

Sketching with quantized contributions

With the RF map (3), which is commonly used in compressive clustering and the GMM, the nonlinear operation $\varrho(\cdot)$ in (40) becomes $\varrho_{\mathrm{RF}}(\cdot) := \exp(-\mathrm{j}2\pi \cdot)$. Since implementing

Approximating Kernel Methods With Random Feature Maps

While each random feature map $\boldsymbol{\Phi}(\cdot)$ implicitly defines a positive definite kernel by taking the expectation (33), the converse is also true. For shift-invariant kernels, i.e., kernels for which $\kappa(\boldsymbol{x}, \boldsymbol{x}') = \kappa(\boldsymbol{x} - \boldsymbol{x}', 0)$ only depends on the difference $\boldsymbol{x} - \boldsymbol{x}'$ (such as the Gaussian kernel), this is a consequence of Bochner's theorem [7], which states that if κ is a positive definite kernel such that $\kappa(\boldsymbol{x}, \boldsymbol{x}) = 1$ for all $\boldsymbol{x}$, then its CFT yields a probability distribution $p_w(\boldsymbol{w}) = \int \kappa(\boldsymbol{x}, 0) \exp(-\mathrm{j}2\pi \boldsymbol{w}^\top \boldsymbol{x}) \mathrm{d}\boldsymbol{x}$. Conversely, the kernel can be obtained by the inverse CFT, which can also be phrased as an expectation:

$$\kappa(\boldsymbol{x}, \boldsymbol{x}') = \int \exp(\mathrm{j}2\pi \boldsymbol{w}^\top (\boldsymbol{x} - \boldsymbol{x}')) p_w(\boldsymbol{w}) \mathrm{d}\boldsymbol{w} = \mathbb{E}_{w \sim p_w} \exp(\mathrm{j}2\pi W^\top (\boldsymbol{x} - \boldsymbol{x}')).$$

Hence, drawing i.i.d. frequency vectors $\boldsymbol{w}_j$ according to p_w yields a random Fourier feature map (3) whose expected kernel is precisely κ. For instance, a Laplace kernel can be approximated if the rows of $\boldsymbol{W}$ are drawn i.i.d. from the Cauchy distribution [S6].

More generally, under mild assumptions on a positive definite kernel κ, one can invoke Mercer's theorem [S7] to similarly show the existence of a random feature map whose expected kernel, in the sense of (33), matches κ.

References

[S6] P. T. Boufounos, S. Rane, and H. Mansour, "Representation and coding of signal geometry," *Inform. Inference*, vol. 6, p. 4, Dec. 2017.

[S7] F. Bach, "On the equivalence between Kernel quadrature rules and random feature expansions," *J. Mach. Learning Res.*, vol. 18, no. 1, pp. 714–751, 2017.

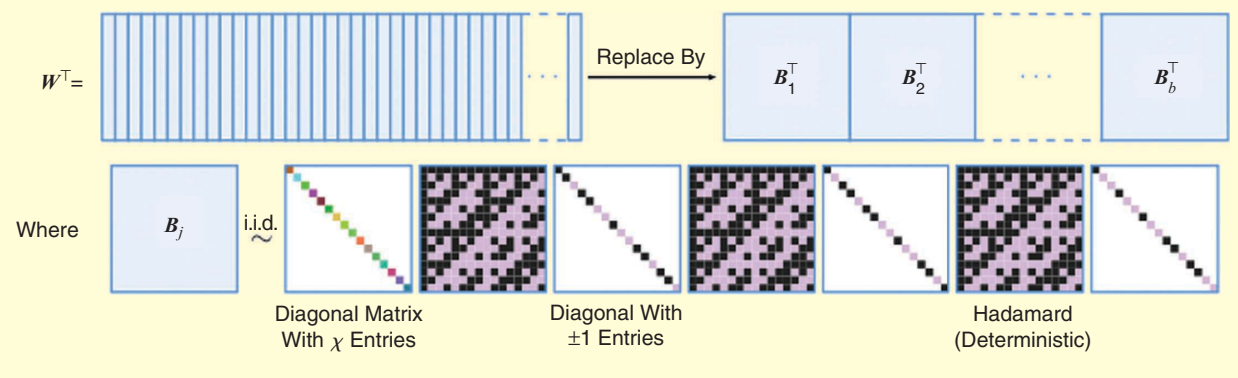

FIGURE 6. The structured random matrix design from [40]. Each block is a composition of several Hadamard and diagonal matrices. For convenience, we draw $\boldsymbol{W}^\top$ instead of $\boldsymbol{W}$.

$\varrho_{\mathrm{RF}}(\cdot)$ 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 $\varrho_{\mathrm{RF}}(\cdot)$ and quantizing the real part to 1 bit of precision yields the 2π-periodic function $\varrho_q(\cdot) := \mathrm{sign}(\cos(2\pi\cdot))$ which is simply a "square wave."

To alleviate the effects of quantization, one can apply dithering [8] to the input. In this case, the feature map becomes $\boldsymbol{\Phi}_q(\boldsymbol{x}) = \varrho_q(\boldsymbol{W}\boldsymbol{x} + \boldsymbol{\xi}) \in \{-1, +1\}^m$, with i.i.d. dither components ξ_j drawn uniformly over $[0, 1)$. The effect of dithering is to make the quantized $\boldsymbol{\Phi}_q$ behave similarly to the nonquantized $\boldsymbol{\Phi}_{\mathrm{RF}}$, on average. For instance, it was shown in [35] that for each $\boldsymbol{W}$, $\boldsymbol{x}$, $\boldsymbol{x}'$, and $\boldsymbol{\xi}$,

$$\begin{aligned}\langle \boldsymbol{\Phi}_q(\boldsymbol{x}), \boldsymbol{\Phi}_{\mathrm{RF}}(\boldsymbol{x}') \rangle &\approx \mathbb{E}_{\boldsymbol{\xi}} \langle \boldsymbol{\Phi}_q(\boldsymbol{x}), \boldsymbol{\Phi}_{\mathrm{RF}}(\boldsymbol{x}') \rangle \\ &= c \langle \boldsymbol{\Phi}_{\mathrm{RF}}(\boldsymbol{x}), \boldsymbol{\Phi}_{\mathrm{RF}}(\boldsymbol{x}') \rangle, \end{aligned} \tag{41}$$

where c is a constant. The approximation in the preceding is accurate for a typical draw of the m-dimensional dither vector $\boldsymbol{\xi}$ when the sketch dimension m is large enough [36]. Note that, when this dithered quantizer is used to compute a sketch $\tilde{\boldsymbol{z}}_q := (1/n)\sum_{i=1}^n \boldsymbol{\Phi}_{\boldsymbol{q}}(\boldsymbol{x}_i) = (1/n)\sum_{i=1}^n \varrho_q(\boldsymbol{W}\boldsymbol{x}_i + \boldsymbol{\xi})$, it is important to use the same dither realization $\boldsymbol{\xi}$ for all samples i.

The question then arises: When estimating parameters $\boldsymbol{\theta}$ via (27), how should we account for quantization in the sketch? Simply replacing the $\mathcal{A}(p_{\boldsymbol{\theta}_\ell})$ term with $\mathcal{A}_q(p_{\boldsymbol{\theta}_\ell}) := \mathbb{E}_{X\sim p_{\boldsymbol{\theta}_\ell}}[\boldsymbol{\Phi}_q(X)]$ may sound appealing. For example, with the compressive GMM (14) or compressive clustering (16), this would mean minimizing $\left\| \tilde{\boldsymbol{z}}_q - \sum_{\ell=1}^k \alpha_\ell \mathcal{A}_q(p_{\boldsymbol{\theta}_\ell}) \right\|^2$. However, the optimization problem (27) would become more challenging, as suggested by comparing the left two panels to the right two panels in Figure 7, which plots the centroid selection criterion $\langle \tilde{\boldsymbol{z}}, \boldsymbol{\Phi}(\boldsymbol{c}) \rangle = \langle \tilde{\boldsymbol{z}}, \mathcal{A}(p_{\boldsymbol{c}}) \rangle$ used by greedy algorithms. Also, this approach would not inherit the theoretical guarantees that were carefully established using the LRIP (29), which does not easily translate to the quantized case.

Instead, we suggest to use $\mathcal{A}_{\mathrm{RF}}(p_{\boldsymbol{\theta}_\ell}) := \mathbb{E}_{X\sim p_{\boldsymbol{\theta}_\ell}}[\boldsymbol{\Phi}_{\mathrm{RF}}(X)]$ for the $\mathcal{A}(p_{\boldsymbol{\theta}_\ell})$ term in (14) and to rescale $\tilde{\boldsymbol{z}}'_q = \tilde{\boldsymbol{z}}_q / c$ with the constant c from (41). Indeed, thanks to (41), the resulting cost function $C(\boldsymbol{\theta} \,|\, \tilde{\boldsymbol{z}}'_q) = \left\| \tilde{\boldsymbol{z}}'_q - \sum_{\ell=1}^k \alpha_\ell \mathcal{A}_{\mathrm{RF}}(p_{\boldsymbol{\theta}_\ell}) \right\|^2$ is, in expectation over $\boldsymbol{\xi}$, exactly the same (up a constant additive bias term) as the nonquantized cost function $C(\boldsymbol{\theta} \,|\, \tilde{\boldsymbol{z}}_{\mathrm{RF}})$ [35]. The similarity between $C(\boldsymbol{\theta} \,|\, \tilde{\boldsymbol{z}}'_q)$ and $C(\boldsymbol{\theta} \,|\, \tilde{\boldsymbol{z}}_{\mathrm{RF}})$, even for a single realization of the m-dimensional dither vector $\boldsymbol{\xi}$, is suggested by comparing the right two panels in Figure 7.

A common way to preserve privacy is to have a trusted data holder (or "curator") corrupt the response to each query of the data set in a controlled manner.

Empirical results [35] suggest that, when the sketch dimension is inflated by about 25%, this quantized compressive learning procedure yields the same performance as the nonquantized procedure. Moreover, accurate probabilistic bounds for approximation (41), established in [36], allow one to extend the theoretical compressive learning guarantees in (30) to this new cost function.

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 data set remain private, i.e., not publicly discoverable. Learning from such data collections while protecting the privacy of individual contributors has become a crucial challenge [37], [38], [39].

A common way to preserve privacy is to have a trusted data holder (or "curator") corrupt the response to each query of the data set [37] 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 data holder 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 [37], [39].

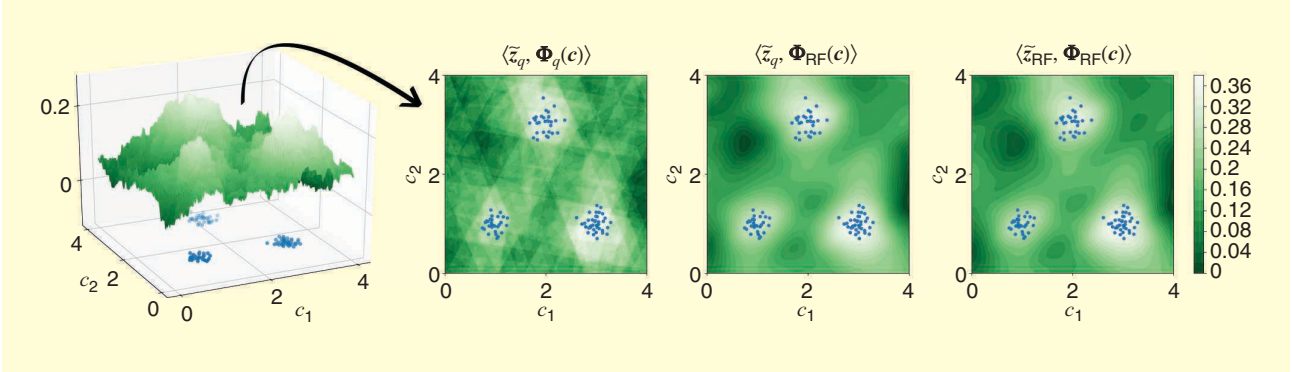

FIGURE 7. The criterion $\langle \tilde{\boldsymbol{z}}, \boldsymbol{\Phi}(\boldsymbol{c}) \rangle$ versus the centroid hypothesis $\boldsymbol{c} = [c_1, c_2]^\top$ used by greedy parameter estimation algorithms in compressive clustering with $k = 3$ and a 2D data set (in blue), with and without quantization of the sketch $\tilde{\boldsymbol{z}}$ and/or feature map $\boldsymbol{\Phi}$. The surface plot on the left highlights the irregularity of $\boldsymbol{c} \mapsto \langle \tilde{\boldsymbol{z}}_q, \boldsymbol{\Phi}_q(\boldsymbol{c}) \rangle$.

When sketching a data set 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 in the following. This one-time approach to privacy preservation (more generally known as *privacy-preserving data publishing* [37]) has several benefits over the query-based approach to privacy preservation discussed in the previous paragraph (see also the sidebar "Potential of Sketching for Privacy-Aware Learning and Alternative Approaches").

Sketching with differential privacy guarantees

Differential privacy [38] 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., [39]). When a given (randomized) learning pipeline is differentially private, its output depends negligibly on the presence or absence of any individual sample in the data set. 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 $\boldsymbol{x}_i$ or can access additional databases that are correlated with the primary database).

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

$$\tilde{\boldsymbol{s}}(\mathcal{X}) := \tilde{\boldsymbol{z}}(\mathcal{X}) + \boldsymbol{v}, \tag{42}$$

where we find it helpful to explicitly denote the dependence of the data set $\mathcal{X}$. We will assume that the realization $\boldsymbol{\Phi}(\cdot)$ of the random feature map is fixed and publicly known, in contrast to other approaches, like [40] and [41], that use linear mixing matrices as encryption keys to ensure privacy preservation. As a result, when we treat $\tilde{\boldsymbol{s}}(\mathcal{X})$ as random, this is due to the randomness in v, not the randomness in $\boldsymbol{\Phi}(\cdot)$ or $\mathcal{X}$.

Formally, the sketching mechanism $\tilde{\boldsymbol{s}}(\cdot)$ is said to be ϵ -differentially private [38] if, for any data set $\mathcal{X} = \{\boldsymbol{x}_i\}_{i=1}^n$ and "neighboring" data set $\mathcal{X}' = \mathcal{X} \setminus \{\boldsymbol{x}_j\} \cup \{\boldsymbol{x}'_j\}$ that replaces the individual sample $\boldsymbol{x}_j$ by another sample $\boldsymbol{x}'_j$, and for any possible sketch outcome $\boldsymbol{s}$, we have that

$$\exp(-\epsilon) \leq \frac{p_{\tilde{\boldsymbol{s}}(\mathcal{X})}(\boldsymbol{s})}{p_{\tilde{\boldsymbol{s}}(\mathcal{X}')}(\boldsymbol{s})} \leq \exp(\epsilon). \tag{43}$$

Here, $\epsilon > 0$ plays the role of a privacy level: a smaller ϵ implies a stronger privacy guarantee. In words, (43) says that, when ϵ is small, the densities of $\tilde{\boldsymbol{s}}(\mathcal{X})$ and $\tilde{\boldsymbol{s}}(\mathcal{X}')$ are almost indistinguishable, as depicted in Figure 8.

Potential of Sketching for Privacy-Aware Learning and Alternative Approaches

Given a data set, privacy preservation can be achieved by asking a trusted data holder to corrupt all queries of the data set [37], [38] 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 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 data user. Also, the data holder must ensure that responses to different data users cannot be combined in a way that circumvents the intended privacy preservation.

In contrast, the noninteractive approach [37], [38] is to publish an intermediate privacy-preserving synopsis of the data set, to which the public is allowed unlimited access. For example, with a low-dimensional data set, one could publish a privacy-preserving histogram of the data [S8], from which aggregate statistics could be subsequently extracted. The noninteractive approach is attractive for several reasons. For example, there is no need to formulate or allocate a privacy budget; it is sufficient to set an overall privacy level. Also, there is no need to worry about data users sharing/combining data.

By adding noise to a sketch of the form (1), one can easily generate a privacy-preserving synopsis of a data set. By construction, such sketches capture the global statistics of the data set $\mathcal{X} = \{\boldsymbol{x}_i\}_{i=1}^n$ while being relatively insensitive to each individual data sample $\boldsymbol{x}_i$, especially when the sample cardinality n is large. Also, when the original data are 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 the sidebar "Compressive Mixture Modeling for Speaker Verification." In this scenario, the local sketches will themselves be privacy preserving, which alleviates concerns about privacy leaks during data fusion.

Reference

[S8] W. Qardaji, W. Yang, and N. Li, "Differentially private grids for geospatial data," in *Proc. IEEE 29th Int. Conf. Data Eng. (ICDE)*, 2013, pp. 757–768.

The condition (43) can be interpreted as bounding a "likelihood ratio" [42], a familiar quantity in signal processing. Consider two hypotheses: one that the data set equals $\mathcal{X}$ (i.e., it includes $\boldsymbol{x}_j$) and the other that the data set equals $\mathcal{X}'$ (i.e., it includes $\boldsymbol{x}'_j$ instead). Then, $p_{\tilde{\boldsymbol{s}}(\mathcal{X})}(\boldsymbol{s})$ would be the likelihood of observing private sketch $\boldsymbol{s}$ under the first hypothesis, while $p_{\tilde{\boldsymbol{s}}(\mathcal{X}')}(\boldsymbol{s})$ would be the same for the second hypothesis. Say an adversary wanted to detect whether or not the data set contains $\boldsymbol{x}_j$. By appropriately thresholding the likelihood ratio $p_{\tilde{\boldsymbol{s}}(\mathcal{X})}(\boldsymbol{s})/p_{\tilde{\boldsymbol{s}}(\mathcal{X}')}(\boldsymbol{s})$, one can obtain hypothesis tests that are optimal from various perspectives (e.g., Bayes, minimax, and Neyman–Pearson) [42, Ch. 2]. Thus, when (43) holds with a small ϵ, it is fundamentally difficult for an adversary to determine whether $\boldsymbol{x}_j$ or $\boldsymbol{x}'_j$ was present in the sketch. Even if the adversary had nontrivial prior knowledge of the true hypothesis (as in so-called linkage attacks, which make use of a second public data set to which the target user contributed), (43) 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 (42) is differentially private, it is sufficient to draw the noise $\boldsymbol{v}$ 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 RF feature map (3), which generates a complex-valued $\tilde{\boldsymbol{z}}(\mathcal{X})$, it has been established [43], [44] that it is sufficient for the real and imaginary components of $\boldsymbol{v}$ to be i.i.d. Laplacian with a standard deviation of $\sigma_v \propto \frac{m}{n\epsilon}$.

Sketching significantly compresses data in a way that facilitates provably accurate yet scalable learning from huge and/or streaming data sets.

A weaker form of privacy, known as *approximate differential privacy* or (ϵ, δ)*-differential privacy* [38], can be attained by adding Gaussian noise $\boldsymbol{v}$ with a smaller variance. For example, with RF features, it is sufficient for the real and imaginary components of $\boldsymbol{v}$ to be i.i.d. Gaussian with a standard deviation of $\sigma_v \propto \frac{\sqrt{m}}{n\epsilon}$.

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 (21)].

In this section, we focused on differential privacy. Other definitions of privacy exist in the literature, such as information-theoretic ones [49]. Likewise, cryptographic methods, such as fully homomorphic encryption [50], 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.

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 data sets, while simultaneously preserving privacy.

In this article, we described several approaches to accelerate the sketching process, including feature quantization and the use

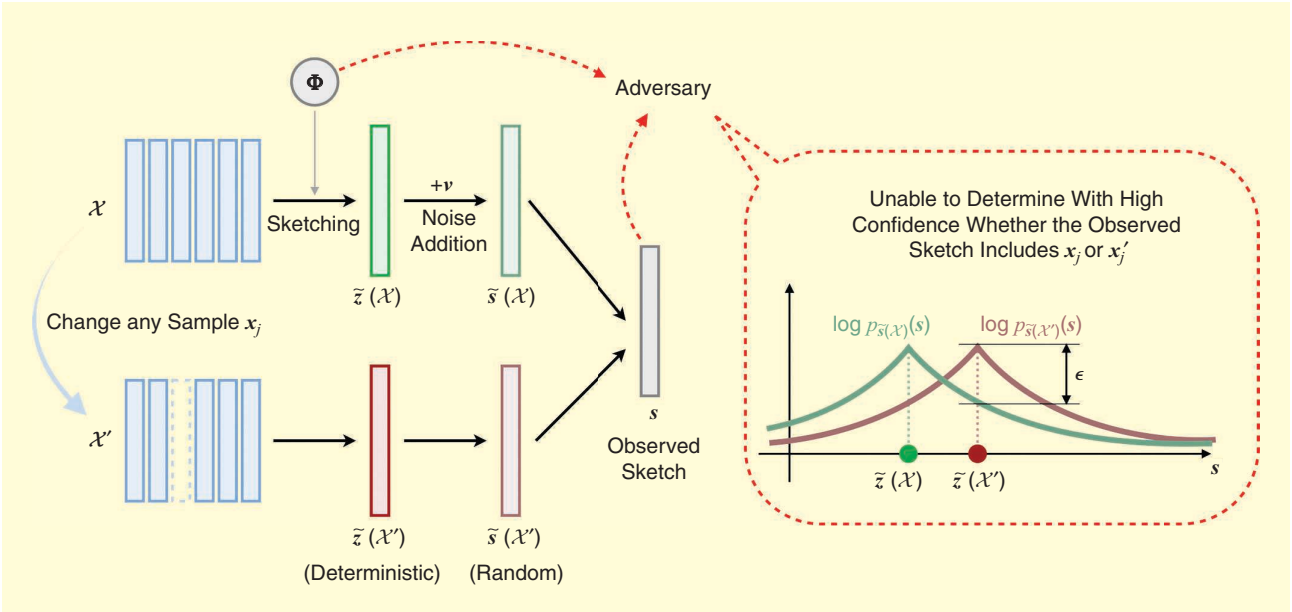

FIGURE 8. When sketching with differential privacy, the output log density of the sketch $\boldsymbol{s}$ remains close when changing one sample in the data set [since (49) is equivalent to $|\log(p_{\tilde{s}(\mathcal{X})}(\boldsymbol{s})) - \log(p_{\tilde{s}(\mathcal{X}')}(\boldsymbol{s}))| \leq \epsilon$ for all possible $\boldsymbol{s}$]. An adversary with knowledge of $\boldsymbol{\Phi}(\cdot)$ and $\boldsymbol{s}$, as symbolized by the red arrows, could then hardly decide whether a given sample $\boldsymbol{x}_j$ was used to compute the sketch or not.

of randomized fast transforms. Another approach is to randomly mask each feature vector $\boldsymbol{\Phi}(\boldsymbol{x}_i)$ prior to averaging, i.e., setting a random subset of its components to zero. It has been established [44] that such random masking does not increase or decrease the differential privacy level ϵ. But it does reduce the need to compute all entries of each feature vector, and thus it 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 [45].

Further work is needed to develop compressive learning methods for supervised tasks like regression and classification.

When discussing methods to learn from a sketch, we focused on optimization-based approaches. Although heuristics based on the OMP [6] and approximate message passing [10] 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 [23], [26], [46], principled greedy methods [47], and gradient flows on systems of particles [48], 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 [30]. For example, one approach to compressive classification was proposed in [49]: for each class $\ell = 1, \ldots, k$, 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 an ML or maximum a posteriori classifier could be derived. Another approach is to perform classification directly in the compressed domain: to an unseen example $\boldsymbol{x}'$, we would assign the class ℓ that maximizes the correlation $\langle \tilde{z}_\ell, \boldsymbol{\Phi}(\boldsymbol{x}') \rangle$. This strategy can be interpreted [49] as compressively evaluating a Parzen window classifier [32]. Further work is also needed on unsupervised matrix factorization tasks like dictionary learning, low-rank matrix completion, and nonnegative matrix factorization [50].

Acknowledgments

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

Authors

Rémi Gribonval (remi.gribonval@inria.fr) received his Habilitation to Conduct Research in applied mathematics from University of Rennes 1, France. He is a senior researcher with Inria at École Normale Supérieure de Lyon, Lyon, 69007, 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, with an emphasis on dimension reduction, inverse problems, and high-dimensional statistics. He founded the series of international workshops Signal Processing With Adaptive/Sparse Representations and received the 2011 Blaise Pascal Award in Applied Mathematics and Scientific Engineering from the Société de Mathématiques Appliquées et Industrielles through the French National Academy of Sciences. He is a Fellow of IEEE and the European Association of Signal Processing.

Antoine Chatalic (antoine.chatalic@dibris.unige.it) received his Ph.D. degree in signal processing in 2020 from the University of Rennes 1, France. He is a postdoctoral researcher at the Machine Learning Genoa Center, Genoa, 16146, 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.

Nicolas Keriven (nicolas.keriven@cnrs.fr) received his Ph.D. degree from the Université de Rennes 1 in 2017, under the supervision of Rémi Gribonval. He is a researcher at the Centre National de la Recherche Scientifique, Paris, 75000, France, and the Grenoble Images Parole Signal Automatique lab, Grenoble, 38402, France. He received the 2017 Best Student Paper award at the Signal Processing With Adaptive/Sparse Representations workshop, for "Random Moments for Sketched Mixture Learning."

Vincent Schellekens (vincent.schellekens@uclouvain.be) received his Ph.D. degree in signal processing in 2021 from the UCLouvain, Louvainla-Neuve, B1348, Belgium. He is a postdoctoral researcher funded by the Fonds de la Recherche Scientique, under the supervision of Laurent Jacques, at the Information and Communication Technologies, Electronics, and Applied Mathematics, UCLouvain, Louvain-la-Neuve, B1348, Belgium. His research focuses on compressive sensing, general random embeddings, and privacy-preservation methods for large-scale learning.

Laurent Jacques (laurent.jacques@uclouvain.be) received his Ph.D. degree in mathematics and physics from UCLouvain, Louvain-la-Neuve, Belgium, in 2004. He is a professor and senior research associate of the Fonds de la Recherche Scientique with the Image and Signal Processing Group, Institute of Information and Communication Technologies, Electronics, and Applied Mathematics, UCLouvain, Louvain-la-Neuve, B1348, Belgium. His research focuses on sparse signal representations, quantized compressive sensing theory, and computational imaging. He has coauthored more than 40 papers in international journals, 80 conference proceedings and presentations, and four book chapters.

Philip Schniter (schniter.1@osu.edu) received his Ph.D. degree in electrical engineering from Cornell University,

Ithaca, New York, USA, in 2000. He is a professor in the Department of Electrical and Computer Engineering, The Ohio State University, Columbus, Ohio, 43210, USA. He received the IEEE Signal Processing Society (SPS) Best Paper Award in 2016 and has served on several IEEE SPS technical committees, including those on signal processing for communications and networking, sensor arrays and multichannel communication, and computational imaging. He is a Fellow of IEEE.

References

[1] R. Gribonval, A. Chatalic, N. Keriven, V. Schellekens, L. Jacques, and P. Schniter, "Sketching datasets for large-scale learning (long version)," 2021, arXiv:2008.01839.

[2] F. Pérez-Cruz and O. Bousquet, "Kernel methods and their potential use in signal processing," *IEEE Signal Process. Mag.*, vol. 21, no. 3, p. 3, 2004.

[3] B. Schölkopf and A. Smola, "Learning with kernels," in *Adaptive Computation and Machine Learning*. Cambridge: MIT Press, 2002. [Online]. Available: https://mitpress.mit.edu/books/learning-kernels

[4] S. Foucart and H. Rauhut, *A Mathematical Introduction to Compressive Sensing*. BerlSpringer-Verlag, May 2012.

[5] R. Gribonval, G. Blanchard, N. Keriven, and Y. Traonmilin, "Compressive statistical learning with random feature moments," *Math. Stat. Learn.*, to be published.

[6] N. Keriven, A. Bourrier, R. Gribonval, and P. Pérez, "Sketching for large-scale learning of mixture models," *Inform. Inference*, vol. 7, no. 3, p. 3, 2017.

[7] A. Rahimi and B. Recht. "Random features for large scale kernel machines," in *Proc. Neural Inform. Process. Syst. Conf.*, 2007, 1177–1184.

[8] R. M. Gray and D. L. Neuhoff, "Quantization," *IEEE Trans. Inform. Theory*, vol. 44, no. 6, pp. 2325–2383, 1998. doi: 10.1109/18.720541.

[9] N. Keriven, N. Tremblay, Y. Traonmilin, and R. Gribonval. "Compressive K-means," in *Proc. IEEE Int. Conf. Acoust. Speech & Signal Process*, 2017, 6369–6373.

[10] E. Byrne, A. Chatalic, R. Gribonval, and P. Schniter, "Sketched clustering via hybrid approximate message passing," *IEEE Trans. Signal Process*, vol. 67, no. 17, p. 17, 2019. doi: 10.1109/TSP.2019.2924585.

[11] M. P. Sheehan, M. S. Kotzagiannidis, and M. E. Davies, "Compressive independent component analysis," in *Proc. 27th IEEE European Signal Process. Conf. (EUSIPCO)*, 2019, pp. 1–5. doi: 10.23919/EUSIPCO.2019.8903095.

[12] M. P. Sheehan, A. Gonon, and M. E. Davies, "Compressive learning for semiparametric models," 2019, arXiv:1910.10024.

[13] G. Cormode, M. Garofalakis, P. J. Haas, and C. Jermaine, "Synopses for massive data," *Found. Trends Databases*, vol. 4, nos. 1–3, p. 1, 2011. doi: 10.1561/1900000004.

[14] D. Achlioptas, "Database-friendly random projections: Johnson–Lindenstrauss with binary coins," *J. Comput. Syst. Sci.*, vol. 66, no. 4, p. 4, 2003. doi: 10.1016/S0022-0000(03)00025-4.

[15] M. W. Mahoney, "Randomized algorithms for matrices and data," *Found. Trends Mach. Learning*, vol. 3, no. 2, pp. 123–224, 2010.

[16] C. Boutsidis, A. Zouzias, and P. Drineas, "Random projections for k-means clustering," in *Proc. Adv. Neural Inform. Process. Syst.*, 2010, pp. 1–9.

[17] D. P. Woodruff, "Sketching as a tool for numerical linear algebra," *Found. Trends Theoretical Comput. Sci.*, vol. 10, nos. 1–2, pp. 1–157, 2014. doi: 10.1561/0400000060.

[18] D. Achlioptas, F. McSherry, and B. Schölkopf, "Sampling techniques for kernel methods," in *Proc. Advances Neural Inform. Process. Syst.*, 2002, pp. 335–342.

[19] D. Feldman, M. Monemizadeh, and C. Sohler, "Coresets and sketches for high dimensional subspace approximation problems," in *Proc. Annu. ACM-SIAM SY~P. Discrete Algorithms*, 2010, vol. 1, pp. 630–649. doi: 10.1137/1.9781611973075.53.

[20] C. Williams and M. W. Seeger, "Using the Nystrom method to speed up kernel machines," in *Proc. Advances Neural Inform. Processing Syst.*, vol. 13, 2001, pp. 628–688.

[21] A. R. Hall, *Generalized Method of Moments*. London: Oxford Univ. Press, 2005.

[22] E. J. Candès and M. B. Wakin, "An introduction to compressive sampling," *IEEE Signal Process. Mag.*, vol. 25, no. 2, p. 2, 2008. doi: 10.1109/MSP.2007.914731.

[23] E. J. Candès and C. Fernandez-Granda, "Towards a mathematical theory of super-resolution," *Commun. Pure Appl. Math*, vol. 67, p. 6, 2014.

[24] M. Lustig, D. L. Donoho, J. M. Santos, and J. M. Pauly, "Compressed sensing MRI," *IEEE Signal Process. Mag.*, vol. 25, no. 2, p. 2, 2008. doi: 10.1109/MSP.2007.914728.

[25] J. A. Fessler, "Optimization methods for magnetic resonance image reconstruction: key models and optimization algorithms," *IEEE Signal Process. Mag.*, vol. 37, no. 1, p. 1, 2020. doi: 10.1109/MSP.2019.2943645.

[26] Q. Denoyelle, V. Duval, G. Peyrée, and E. Soubies, "The sliding Frank–Wolfe algorithm and its application to super-resolution microscopy," *Inverse Problems*, vol. 36, no. 1, p. 1, Jan. 2020. doi: 10.1088/1361-6420/ab2a29.

[27] 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," *IEEE Signal Process. Mag.*, vol. 35, no. 4, p. 4, 2018. doi: 10.1109/MSP.2018.2821706.

[28] R. Ge, J. D. Lee, and T. Ma, "Matrix completion has no spurious local minimum," in *Proc. Adv. Neural Inf. Process. Syst.*, New York: Curran Associates, 2016.

[29] V. Vapnik, *Nature Statistical Learning Theory*, vol. 6, Berlin, Germany: Springer, 2013.

[30] S. Shalev-Shwartz and S. Ben-David, *Understanding Machine Learning*. Cambridge, MA: Cambridge Univ. Press, 2009.

[31] W. B. Johnson and J. Lindenstrauss, "Extensions of Lipschitz mappings into a Hilbert space," *Contemporary Math.*, vol. 26, pp. 189–206, 1984.

[32] Z.-W. Pan, D.-H. Xiang, Q.-W. Xiao, and D.-X. Zhou, "Parzen windows for multi-class classification," *J. Complexity*, vol. 24, nos. 5–6, pp. 5–6, 2008. doi: 10.1016/j.jco.2008.07.001.

[33] 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, pp. 1975–1983.

[34] 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, pp. 4714–4718.

[35] V. Schellekens and L. Jacques, "Quantized compressive k-means," *IEEE Signal Process. Lett.*, vol. 25, no. 8, pp. 1211–1215, Aug. 2018. doi: 10.1109/LSP.2018.2847908.

[36] V. Schellekens and L. Jacques, "Breaking the waves: Asymmetric random periodic features for low-bitrate kernel machines," *Inf. Inference*, to be published. doi: https://doi.org/10.1093/imaiai/iaab008

[37] B. C. M. Fung, K. Wang, R. Chen, and P. S. Yu, "Privacy-preserving data publishing: A survey of recent developments," *ACM Comput. Survey*, vol. 42, no. 4, p. 4, June 2010. doi: 10.1145/1749603.1749605.

[38] C. Dwork and A. Roth, "The algorithmic foundations of differential privacy," *Theoretical Comput. Sci.*, vol. 9, no. 3–4, pp. 3–4, 2014.

[39] A. D. Sarwate and K. Chaudhuri, "Signal processing and machine learning with differential privacy: Algorithms and challenges for continuous data," *IEEE Signal Process. Mag.*, vol. 30, no. 5, p. 5, 2013. doi: 10.1109/MSP.2013.2259911.

[40] M. Testa, D. Valsesia, T. Bianchi, and E. Magli, *Compressed Sensing for Privacy-Preserving Data Processing*. Berlin: Springer-Verlag, 2019.

[41] S. Rane and P. T. Boufounos, "Privacy-preserving nearest neighbor methods: Comparing signals without revealing them," *IEEE Signal Process. Mag.*, vol. 30, no. 2, p. 2, 2013. doi: 10.1109/MSP.2012.2230221.

[42] H. V. Poor, *An Introduction to Signal Detection and Estimation*. Springer Science & Business Media, 2013.

[43] 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, pp. 7933–7937.

[44] A. Chatalic, V. Schellekens, F. Houssiau, Y.-A. de Montjoye, L. Jacques, and R. Gribonval, "Compressive learning with privacy guarantees," *Information and Inference: A Journal of the IMA*, May 2021. doi: 10.1093/imaiai/iaab005. [Online]. Available: https://doi.org/10.1093/imaiai/iaab005

[45] A. Saade et al., "Random projections through multiple optical scattering," in *Proc. IEEE Int. Conf. Acoust. Speech & Signal Process*, 2016, 6215–6219.

[46] N. Boyd, G. Schiebinger, and B. Recht, "The alternating descent conditional gradient method for sparse inverse problems," *SIAM J. Optim.*, vol. 27, no. 2, p. 2, 2017. doi: 10.1137/15M1035793.

[47] Y. Traonmilin and J.-F. Aujol, "The basins of attraction of the global minimizers of the non-convex sparse spikes estimation problem," 2018, arXiv:1811.12000.

[48] L. Chizat and F. Bach, "On the global convergence of gradient descent for over-parameterized models using optimal transport," in *Proc. Adv. Neural Inform. Process. Syst.*, 2018.

[49] V. Schellekens and L. Jacques, "Compressive classification (machine learning without learning)," 2018, arXiv:1812.01410.

[50] M. Udell, C. Horn, R. Zadeh, and S. Boyd, "Generalized low rank models," *Found. Trends Mach. Learn.*, vol. 9, no. 1, p. 1, 2016. doi: 10.1561/2200000055.

[51] R. Gribonval, G. Blanchard, N. Keriven and Y. Traonmilin, "Statistical learning guarantees for compressive clustering and compressive mixture modeling," *Math. Stat. Learn.*, to be published.

SP