

Quantized Compressive K-Means

Vincent Schellekens* and Laurent Jacques*

Abstract—The recent framework of *compressive statistical learning* proposes to design tractable learning algorithms that use only a heavily compressed representation—or *sketch*—of massive datasets. Compressive K-Means (CKM) is such a method: it aims at estimating the centroids of data clusters from pooled, nonlinear, random signatures of the learning examples. While this approach significantly reduces computational time on very large datasets, its digital implementation wastes acquisition resources because the learning examples are compressed only *after* the sensing stage. The present work generalizes the CKM sketching procedure to a large class of periodic nonlinearities including hardware-friendly implementations that *compressively acquire entire datasets*. This idea is exemplified in a Quantized Compressive K-Means procedure, a variant of CKM that leverages 1-bit universal quantization (*i.e.*, retaining the least significant bit of a standard uniform quantizer) as the periodic sketch nonlinearity. Trading for this resource-efficient signature (standard in most acquisition schemes) has almost no impact on the clustering performance, as illustrated by numerical experiments.

I. INTRODUCTION

Numerous scientific fields have recently experienced a paradigm shift towards data-driven approaches where mathematical models are inferred from a dataset of learning examples $X = \{\mathbf{x}_i \in \mathbb{R}^n\}_{i=1}^N$. *K-means clustering* (KMC) [1] is such a method widely used in, *e.g.*, data compression, pattern recognition, and bioinformatics [2], [3]. Given K , a prescribed number of clusters (groups of similar data), KMC seeks the centroids (or “cluster representatives”) $\mathcal{C} = \{\mathbf{c}_k \in \mathbb{R}^n\}_{k=1}^K$ minimizing the Sum of Squared Errors (SSE):

$$\mathcal{C}^* = \arg \min_{\mathcal{C}} \text{SSE}(\mathcal{C}) = \arg \min_{\mathcal{C}} \sum_{i=1}^N \min_{1 \leq k \leq K} \|\mathbf{x}_i - \mathbf{c}_k\|^2. \quad (1)$$

Solving (1) exactly is NP-hard [4], so in practice a tractable heuristic such as the popular *k-means* algorithm [5], [6] is widely used to find an approximate solution \mathcal{C}_{km} . However, *k-means* complexity scales poorly with the size of modern voluminous datasets where N is typically $\mathcal{O}(10^3 - 10^6)$, or grows continually for *data streams* processing. In fact, since *k-means* repeatedly requires—at each iteration—a thorough pass over X , this massive dataset must be stored and read several times, with prohibitive memory and time consumptions. Paradoxically, the large dataset size (*i.e.*, nN) dwarfs, and does not affect, the number of parameters learned by *k-means* (*i.e.*, nK). Ideally, larger datasets increase the model accuracy without requiring more training computational resources.

This goal motivates the recent *compressive learning* framework [7], where learning algorithms solely require access to a drastically compressed representation of the dataset called the

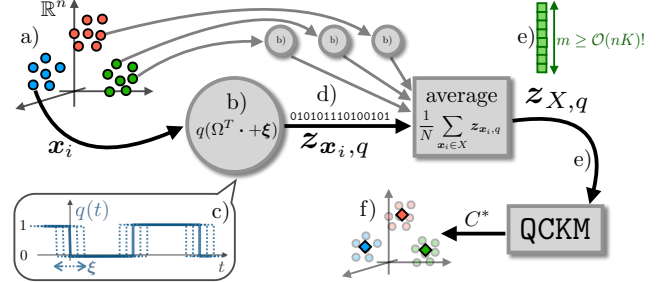


Fig. 1: **a)** A massive dataset X , composed of K clusters of examples $\mathbf{x}_i \in \mathbb{R}^n$, is not explicitly available but acquired via one (or a cloud of) low-power sensor(s) **b)** implementing random projections on frequencies Ω , a dithering ξ and **c)** the 1-bit universal quantization $q(\cdot)$ (-1 is encoded as 0). Only **d)** the sketch contributions of the examples (m bits) are acquired, and form after averaging **e)** the sketch, a highly compressed but meaningful representation of X . Our QCKM method then extracts **f)** the K cluster centroids from it.

sketch, a single vector $\mathbf{z}_X \in \mathbb{C}^m$, constructed by collecting m (empirical) generalized moments of the dataset X :

$$\mathbf{z}_X := \frac{1}{N} \sum_{i=1}^N \mathbf{z}_{\mathbf{x}_i} \text{ where } \mathbf{z}_{\mathbf{x}_i} := [\exp(-i\omega_j^T \mathbf{x}_i)]_{j=1}^m, \quad (2)$$

with “frequencies” $\omega_j \in \mathbb{R}^n$ sampled randomly according to a distribution Λ . The sketch is thus the *pooling* (average) of random projections of the data samples after passing through a nonlinear, periodic signature—the complex exponential. The Compressive K-Means (CKM) method [8] clusters X from \mathbf{z}_X by replacing (1) with a *sketch matching* optimization problem:

$$(\mathcal{C}_{\text{CKM}}, \alpha_{\text{CKM}}) = \arg \min_{(\mathcal{C}, \alpha)} \|\mathbf{z}_X - \sum_{k=1}^K \alpha_k \mathbf{z}_{\mathbf{c}_k}\|^2, \quad (3)$$

with the cluster weights $\alpha_k \geq 0$ satisfying $\sum_k \alpha_k = 1$. It was shown empirically that $\text{SSE}(\mathcal{C}_{\text{CKM}}) \simeq \text{SSE}(\mathcal{C}_{\text{km}})$ provided $m = \mathcal{O}(nK)$, *i.e.*, the required sketch size m is only proportional to the number of parameters to learn, allowing for tractable memory consumption and training time *whatever the number of training examples* N . It was also theoretically proven that $m = \mathcal{O}(nK^2)$ is a sufficient condition for retrieving meaningful centroids from the sketch [7]. Interestingly, the sketch is (up to a re-scaling) linear¹ and \mathbf{z}_X can thus be computed in one pass over the data, possibly realized in parallel over several machines. This sketch is also easy to update when new examples are available (*e.g.*, in data streams).

The following limitation in CKMs sketching strategy motivate this work: *all signals* $\mathbf{x}_i \in X$ (or, equivalently, their projections onto the frequencies $\{\omega_j\}_{j=1}^m$, as in a compressive sensing scheme [9]) must be acquired and stored at full resolution (high bitrate) in order to expansively evaluate (in software) their contributions $\mathbf{z}_{\mathbf{x}_i}$ to the sketch. A resource-preserving (*e.g.*, computational or energy efficient) compressive learning sensor should directly and solely acquire $\mathbf{z}_{\mathbf{x}_i}$. While random projections can be cheap to compute (*e.g.*, by using fast structured random projections [10], [11] or, possibly,

* E-mail: {vincent.schellekens, laurent.jacques}@uclouvain.be. ISPGROUP, ELEN/ICTEAM, UCLouvain (UCL), B1348 Louvain-la-Neuve, Belgium. VS and LJ are funded by Belgian National Science Foundation (F.R.S.-FNRS).

¹A sketch Φ_S of a vector set S is said *linear* if $\Phi_{S \cup S'} = \Phi_S + \Phi_{S'}$.

by relying on optical random processes [12]) the evaluation of the complex exponential in (2) is complicated to implement in hardware—the costliest step in fast sketch computation [10].

Inspired by recent works concerning 1-bit random embeddings [13], we propose a new sketch procedure, illustrated in Fig. 1. This one is conceptually much simpler to integrate directly in hardware (*e.g.*, using voltage controlled oscillators [14]), bypassing the high-bitrate signal acquisition. We replace the costly $\exp(-i\cdot)$ signature function by *1-bit universal quantization* $q(\cdot) = \text{sign}(\cos(\cdot)) = 2(\lfloor \frac{\cdot}{2\pi} \rfloor \bmod 2) - 1$. This function (represented in Fig. 1c), corresponds to taking the least significant bit (LSB) of a uniform quantizer with quantizer stepsize π . We justify this modification by proving that the periodicity of the signature function is much more important than its particular shape. While our sketch is cheaper to compute, it retains the advantages of the original one, *i.e.*, it is linear (prone to distributed computing) and its required size m scales—as we show in our experiments—still as $\mathcal{O}(nK)$, with only a 15 to 25% increase compared to CKM.

Outline: Sec. II recapitulates how CKM performs clustering from the sketch z_X . We propose a generalized sketch $z_{X,f}$ in Sec. III where the signature $\exp(-i\cdot)$ is replaced by a generic periodic function $f(\cdot)$. Our main result states that, in the CKM clustering method, $z_{X,f}$ can be used instead of z_X , even though f is potentially non-differentiable, thanks to the addition of a random *dithering* on the argument of f . This claim is supported by the possibility to recover the cost function implicitly minimized in CKM from $z_{X,f}$ (Prop. 1). Based on this observation, in Sec. IV we define our Quantized Compressive K-Means (QCKM) method that solves KMC from $z_{X,q}$, a 1-bit sketch of the dataset associated with $f = q$. We validate experimentally that QCKM competes favorably with CKM in Sec. V, before concluding in Sec. VI.

Related work: Most fast clustering methods for massive datasets rely on sample-wise dimensionality reduction [15]–[18]. One notable exception is the coresets method [19] that proposes to subsample the dataset to both approximate the SSE and boost K-means. For the related kernel K-means problem, [20] uses Random Fourier Features [21], *i.e.*, the low-dimensional mapping $z_{(\cdot)}$ defined in (2). For $u, v \in \mathbb{R}^n$, the inner product $\langle z_u, z_v \rangle$ approximates a shift-invariant kernel $\kappa(u, v)$ associated with the frequency distribution Λ . CKM [8] actually *averages* individual RFF of data points. Interestingly, κ also defines a Reproducing Kernel Hilbert Space in which two probability density functions (pdfs) can be compared with a Maximum Mean Discrepancy (MMD) metric [22]–[26]. Equipped with the MMD metric, the Generalized Method of Moments [27] in (3) is equivalent to an infinite-dimensional Compressed Sensing [11] problem, where the “sparse” pdf underlying the data (*e.g.*, approximated by few Diracs) is reconstructed from a small number of compressive, random linear pdf measurements: the sketch [26]. The method to solve (3) is thus inspired by the OMP (R) CS recovery algorithm, *i.e.*, Orthogonal Matching Pursuit (with Replacement) [28], [29]. In this work, analogously to how RFF are generalized to any periodic signature of random projections in [13], [30], *pooled* RFF (the dataset sketch) are generalized to *pooled* periodic

signatures of random projections, with universal quantization (known to preserve local signal distances) as a particular case.

II. BACKGROUND: COMPRESSIVE K-MEANS (CKM)

Most unsupervised learning tasks amount to estimating (some parameters of) the unknown probability distribution \mathcal{P} from N learning examples $x_i \stackrel{\text{iid}}{\sim} \mathcal{P}$. Compressive learning aims at estimating \mathcal{P} from its sketch $\mathcal{A}(\mathcal{P}) \in \mathbb{C}^m$: a random sampling of its characteristic function $\phi_{\mathcal{P}}(\omega) := \mathbb{E}_{x \sim \mathcal{P}} e^{i\omega^T x}$ at m frequencies $\Omega = (\omega_1, \dots, \omega_m)$ drawn from a well-specified distribution $\omega_j \stackrel{\text{iid}}{\sim} \Lambda$. The *sketch operator* reads²

$$\begin{aligned} \mathcal{A}(\mathcal{P}) &:= \mathbb{E}_{x \sim \mathcal{P}} e^{-i\Omega^T x} = (\phi_{\mathcal{P}}^*(\omega_j))_{j=1}^m \\ &\simeq z_X := \mathcal{A}(\hat{\mathcal{P}}_X) = \frac{1}{N} \sum_{x_i \in X} e^{-i\Omega^T x_i}, \end{aligned} \quad (4)$$

where the sketch z_X of a dataset X actually refers to the sketch of its empirical pdf, $\hat{\mathcal{P}}_X := \frac{1}{N} \sum_{x_i \in X} \delta_{x_i}$, as announced in (2). Given its sketch $\mathcal{A}(\mathcal{P})$, \mathcal{P} can be approximated by a pdf \mathcal{Q} belonging to some simple (“sparse”) model set \mathcal{G} —where the approximation error is quantified by the MMD metric, that can in this particular context be written as $\gamma_{\Lambda}^2(\mathcal{P}, \mathcal{Q}) := \mathbb{E}_{\omega \sim \Lambda} |\phi_{\mathcal{P}}(\omega) - \phi_{\mathcal{Q}}(\omega)|^2$ [24], [31]. The ℓ_2 sketch distance serves as an estimate for $\gamma_{\Lambda}(\mathcal{P}, \mathcal{Q})$, hence in practice \mathcal{Q} is found by solving the *sketch matching problem*:

$$\mathcal{Q}^* \in \arg \min_{\mathcal{Q} \in \mathcal{G}} \|\mathcal{A}(\mathcal{P}) - \mathcal{A}(\mathcal{Q})\|^2 \simeq \arg \min_{\mathcal{Q} \in \mathcal{G}} \gamma_{\Lambda}^2(\mathcal{P}, \mathcal{Q}). \quad (5)$$

In CKM, $\mathcal{A}(\mathcal{P})$ is approximated by $\mathcal{A}(\hat{\mathcal{P}}_X)$, and \mathcal{Q} is a weighted mixture of K Diracs located at the the centroids $\mathbf{c}_k \in \mathcal{C} \subset \mathbb{R}^n$, *i.e.*, $\mathcal{G} := \{\sum_{k=1}^K \alpha_k \delta_{\mathbf{c}_k} : \mathbf{c}_k \in \mathcal{C}, \alpha_k \geq 0, \sum \alpha_k = 1\}$. From (5) the CKM objective function reads:

$$(\mathcal{C}_{\text{CKM}}, \alpha_{\text{CKM}}) \in \arg \min_{(\mathcal{C}, \alpha)} \|z_X - \mathcal{A}(\sum_{k=1}^K \alpha_k \delta_{\mathbf{c}_k})\|^2, \quad (6)$$

as announced in (3). This non-convex problem is hard to solve exactly, but the CKM algorithm (based on OMPR) [8], detailed in pseudocode below, seeks an approximate solution. More precisely, CKM greedily selects new centroids minimizing a residual $\mathbf{r} \in \mathbb{C}^m$ (Steps 1 and 2) inside a box with lower and upper bounds $l, u \in \mathbb{R}^n$, respectively, enclosing the data X , and eventually replacing bad centroids in Step 3. The centroid weights α_k are then computed and a global gradient descent initialized at the current values allows further decrease of the objective (Steps 4 and 5). CKM relies on solving several (not always convex) optimization sub-problems, in practice solved approximately (a *local* optimum is found) using a quasi-Newton optimization scheme.

CKM parameters: The frequency distribution $\Lambda(\omega)$ ought to define a meaningful metric γ_{Λ} , *i.e.*, the objective function of CKM. By Bochner’s theorem [32], Λ is associated with a positive definite, translation-invariant kernel (“similarity measure”) $\kappa(x, x') = K(x - x')$ through the Fourier transform $K(u) = \mathcal{F}(\Lambda)(u) := \int e^{-iu^T \omega} d\Lambda(\omega)$. Concretely, Λ limits the frequencies of \mathcal{P} we are able to observe in $\mathcal{A}(\mathcal{P})$, acting as a “low-pass filter” convolving \mathcal{P} with K . Λ thus implicitly controls CKM’s clustering scale, and requires some *a priori*

²Scalar functions (*e.g.*, \exp) are here applied component-wise on vectors.

Algorithm: CKM: Compressive K -Means clustering.

```

1  $\mathbf{r} \leftarrow \mathbf{z}_X, \mathcal{C} \leftarrow \emptyset$  (Initialize residual and centroids)
2 for  $t = 1, \dots, 2K$  do
3   Step 1 : gradient descent selects  $\mathbf{c}$  highly correlated with residual:
4    $\mathbf{c} = \underset{\mathbf{c}}{\text{maximize}} \Re \langle \frac{\mathcal{A}\delta_{\mathbf{c}}}{\|\mathcal{A}\delta_{\mathbf{c}}\|}, \mathbf{r} \rangle$  s.t.  $\mathbf{l} \leq \mathbf{c} \leq \mathbf{u}$ 
5   Step 2 : add it to the support:
6    $\mathcal{C} \leftarrow \mathcal{C} \cup \{\mathbf{c}\}$ 
7   Step 3 : Reduce support by Hard Thresholding:
8   if  $|\mathcal{C}| > K$  then
9      $\beta = \arg \min_{\beta} \|\mathbf{z}_X - \sum_{k=1}^{|\mathcal{C}|} \beta_k \frac{\mathcal{A}\delta_{\mathbf{c}_k}}{\|\mathcal{A}\delta_{\mathbf{c}_k}\|}\|$  s.t.  $\beta \in \mathbb{R}_+^{|\mathcal{C}|}$ 
10     $\mathcal{C} \leftarrow$  set of  $K$  centroids  $\mathbf{c}_k$  corresponding to  $K$  largest
        magnitude values of  $\beta$ .
11  end
12  Step 4 : Project to find optimal weights:
13   $\alpha = \arg \min_{\alpha} \|\mathbf{z}_X - \sum_{k=1}^{|\mathcal{C}|} \alpha_k \mathcal{A}\delta_{\mathbf{c}_k}\|$  s.t.  $\alpha \in \mathbb{R}_+^{|\mathcal{C}|}$ 
14  Step 5 : Global gradient descent
15   $(\mathcal{C}, \alpha) \leftarrow \underset{\mathcal{C}=\{\mathbf{c}_k\}, \alpha}{\text{minimize}} \|\mathbf{z}_X - \sum_k \alpha_k \mathcal{A}\delta_{\mathbf{c}_k}\|$  s.t.  $\mathbf{l} \leq \mathbf{c}_k \leq \mathbf{u}$ 
16   $\mathbf{r} = \mathbf{z}_X - \sum_{k=1}^{|\mathcal{C}|} \alpha_k \mathcal{A}\delta_{\mathbf{c}_k}$  (Update residual)
17 end
```

insight about \mathcal{P} . In practice CKM uses heuristics adjusting Λ from a subset of X [26]. For the required sketch dimension m , [7] provides theoretical guarantees when $m = \mathcal{O}(nK^2)$ but experiments strongly suggest $\mathcal{O}(nK)$ is sufficient [8].

III. SKETCHING WITH GENERAL SIGNATURE FUNCTIONS

We here generalize the sensing function $\exp(-i \cdot)$ of the sketch (4) to a general (e.g., discontinuous) periodic function $f(t)$ assumed (w.l.o.g.) 2π -periodic, centered and taking values in $[-1, 1]$. Therefore $f(t) = \sum_{k \neq 0} F_k e^{ikt}$, with Fourier series coefficients F_k such that $F_0 = 0$, and $F_{\pm 1} \neq 0$ (up to a rescaling of f). The *generalized sketch operator* \mathcal{A}_f is

$$\mathcal{A}_f(\mathcal{P}) := \mathbb{E}_{\mathbf{x} \sim \mathcal{P}} f(\Omega^T \mathbf{x} + \xi) \quad \text{for } \omega_j \stackrel{\text{iid}}{\sim} \Lambda(\omega), \quad (7)$$

with a uniform *dithering* $\xi_j \stackrel{\text{iid}}{\sim} \mathcal{U}([0, 2\pi])$.

Our main question is now: *given $\mathcal{A}_f(\mathcal{P})$, is it still possible to approximate \mathcal{P} by some low-complexity distribution \mathcal{Q} , as done in (5)?* Our answer is positive since we can still approximate the same objective function, i.e., the MMD metric $\gamma_\Lambda(\mathcal{P}, \mathcal{Q})$, from this new sketch. Intuitively, the dithering ξ allows us to “separate” $\mathcal{A}_f(\mathcal{P})$ into two terms: one associated with the low frequencies of f , that contributes to the target objective $\gamma_\Lambda(\mathcal{P}, \mathcal{Q})$, and one “high-frequency” term that is constant for the relevant optimization problem. This is formally proven in the following proposition.

Proposition 1. *Given pdfs \mathcal{P}, \mathcal{Q} , denoting f 's 1st harmonic as $f_1(t) := \sum_{k \in \{\pm 1\}} F_k e^{ikt}$, there is a constant $c_{\mathcal{P}} > 0$ s.t.*

$$|(2m|F_1|^2)^{-1} \|\mathcal{A}_f(\mathcal{P}) - \mathcal{A}_{f_1}(\mathcal{Q})\|^2 - \gamma_\Lambda^2(\mathcal{P}, \mathcal{Q}) - c_{\mathcal{P}}| \leq \epsilon, \quad (8)$$

with probability exceeding $1 - 2e^{-C_f m \epsilon^2}$ on the draw of Ω and ξ , for $C_f = 8|F_1|^4 (1 + 2|F_1|)^{-4}$.

Proof. Note that $\|\mathcal{A}_f(\mathcal{P}) - \mathcal{A}_{f_1}(\mathcal{Q})\|^2 = \sum_{j=1}^m Z_j$ with $Z_j := |\mathbb{E}_{\mathbf{x} \sim \mathcal{P}} f(\omega_j^T \mathbf{x} + \xi_j) - \mathbb{E}_{\mathbf{y} \sim \mathcal{Q}} f_1(\omega_j^T \mathbf{y} + \xi_j)|^2 \leq (1 + 2|F_1|)^2$. Define $\tilde{F}_k = F_k(\delta_{k', 1} + \delta_{k', -1})$ with $\delta_{p, q}$ the Kronecker delta. Since $\mathbb{E}_{\xi} e^{ik'\xi} = \delta_{k', 0}$, $\mu_Z := \mathbb{E}_{\omega_j, \xi_j} Z_j$ reads

$$\begin{aligned} \mu_Z &= \mathbb{E}_{\omega} \mathbb{E}_{\xi} \left| \sum_{k \neq 0} e^{ik\xi} (F_k \mathbb{E}_{\mathbf{x} \sim \mathcal{P}} e^{ik\omega^T \mathbf{x}} - \tilde{F}_k \mathbb{E}_{\mathbf{x}' \sim \mathcal{Q}} e^{ik\omega^T \mathbf{x}'}) \right|^2 \\ &= \mathbb{E}_{\omega \sim \Lambda} \sum_{k \neq 0} |F_k \phi_{\mathcal{P}}(k\omega) - \tilde{F}_k \phi_{\mathcal{Q}}(k\omega)|^2 \\ &= 2|F_1|^2 \gamma_\Lambda^2(\mathcal{P}, \mathcal{Q}) + \mathbb{E}_{\omega} \sum_{|k| \geq 2} |F_k|^2 |\phi_{\mathcal{P}}(k\omega)|^2. \end{aligned}$$

Hoeffding's inequality applied on the bounded Z_j s gives (8) with $c_{\mathcal{P}} = \mathbb{E}_{\omega} \sum_{|k| \geq 2} \frac{|F_k|^2}{2|F_1|^2} |\phi_{\mathcal{P}}(k\omega)|^2$ constant w.r.t. \mathcal{Q} . Note that $c_{\mathcal{P}} < \infty$ since f is bounded and $F_1 \neq 0$. \square

Prop. 1 allows us—at the price of adding a dithering—to sketch a dataset into $\mathbf{z}_{X, f} = \mathcal{A}_f(\hat{\mathcal{P}}_X)$ with a very large class of functions f , e.g., that model a realistic sensing scheme. Indeed, it shows that, for a fixed pair of pdfs, replacing CKMs objective in (5) by $\|\mathcal{A}_f(\mathcal{P}) - \mathcal{A}_{f_1}(\mathcal{Q})\|^2$ still approximates (up to a harmless constant $c_{\mathcal{P}}$) the MMD metric $\gamma_\Lambda^2(\mathcal{P}, \mathcal{Q})$: an intuitively good cost function as justified by previous work. Moreover, thanks to the fact that f_1 is a cosine, the relevant gradients of this new cost function enjoy the same nice analytic expressions as CKM. Fixing the probability of (8), we also see that (again, for a fixed pair of distributions) the approximation error ϵ decays like $\mathcal{O}(1/\sqrt{m})$ as m increases. However, it is yet unclear how large m must be to characterize the quality of the solution of (10). This would require Prop. 1 to hold for all \mathcal{Q} in a “low-dimensional” set $\mathcal{G} \ni \mathcal{P}$ as done in [7], a generalization that we postpone to a future work.

IV. QUANTIZED COMPRESSIVE K-MEANS (QCKM)

We now instantiate the results of Sec. III to the 1-bit universal quantization $f(t) = q(t) = \text{sign}(\cos(t)) \in \{-1, +1\}$ to construct a hardware-friendly *quantized sketch*. The 1-bit universal quantizer q is a square wave and can be seen as the Least Significant Bit of a uniform quantizer with quantization stepsize π [13], [30]. The resulting sketch operator \mathcal{A}_q on a pdf \mathcal{P} (resp. the sketch $\mathbf{z}_{X, q}$ of a dataset X) is

$$\mathcal{A}_q(\mathcal{P}) := \mathbb{E}_{\mathbf{x} \sim \mathcal{P}} q(\Omega^T \mathbf{x} + \xi) \simeq \mathbf{z}_{X, q} := \mathcal{A}_q(\hat{\mathcal{P}}_X), \quad (9)$$

and the clustering problem in (6) is now replaced by

$$(\mathcal{C}_{\text{QCKM}}, \alpha_{\text{QCKM}}) \in \arg \min_{(\mathcal{C}, \alpha)} \|\mathbf{z}_{X, q} - \mathcal{A}_{q_1}(\sum_k \alpha_k \delta_{\mathbf{c}_k})\|^2. \quad (10)$$

where q_1 denotes the first harmonic of q (a cosine). Interestingly, the contribution $\mathbf{z}_{\mathbf{x}_i, q} = \mathcal{A}_q(\delta_{\mathbf{x}_i}) \in \{-1, 1\}^m$ of each signal \mathbf{x}_i can be encoded by only m bits, as illustrated Fig. 1.

To solve (10), we adapt the CKM algorithm to account for the changes in objective function, which we call the Quantized Compressive K-Means (QCKM) algorithm. More precisely, \mathbf{z}_X is replaced by $\mathbf{z}_{X, q}$ at initialization and in Steps 3, 4 and 5, and $\mathcal{A}\delta_{\mathbf{c}}$ is replaced by $\mathcal{A}_{q_1}\delta_{\mathbf{c}}$ in Steps 1, 3, 4 and 5. Small modifications also take into account the addition of ξ .

V. EXPERIMENTS TO VALIDATE QCKM

We show now empirically that QCKM requires only $m = \mathcal{O}(nK)$ measurements to find good centroids, with a hidden multiplicative constant only slightly higher (15 to 25%) than for CKM—remembering that QCKM receives m -bit sketch contributions whereas CKM uses full-precision contributions. We validate QCKM on both synthetic and real datasets and compare the performance with k -means (built-in MATLAB function) as well as CKM (from the SketchMLbox toolbox [33]).

Synthetic data: We compute *phase transition diagrams* (Fig. 2) to highlight the relationship between the required

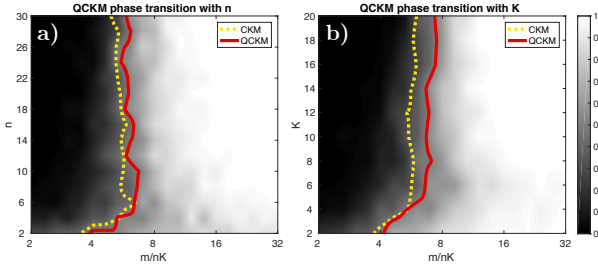


Fig. 2: Empirical success rate—from 0% (black) to 100% (white)—evolution of QCKM with m/nK , and (a) n , or (b) K . The red solid line shows the transition to a success rate above 50% for QCKM; for comparison, the yellow dotted line shows the same transition for CKM.

amount of measurements m , and the sample dimension n or the number of clusters K . For this, we arbitrary say that (Q)CKM is *successful* if $SSE_{(Q)CKM} \leq 1.2SSE_{k\text{-means}}$, where $SSE_{k\text{-means}}$ is the best out of 5 $k\text{-means}$ runs. These diagrams show how the empirical success rate (averaged over 100 trials) of QCKM evolves with m , as n or K varies. For fair comparison with the complex exponential sketch (composed of a cosine and sine in its real and imaginary part, respectively), the j^{th} measurement of the quantized sketch is, in our experiment, composed of two measurements with the same frequency ω_j but two dithering values ξ_j and $\xi_j + \frac{\pi}{2}$. First, we draw $N = 10000$ samples uniformly from $K = 2$ isotropic Gaussians in *varying dimension* n , with means $\pm(1, \dots, 1)^T \in \mathbb{R}^n$ and covariance matrix $\frac{n}{20}\text{Id}$. The phase transition diagram is reported Fig. 2a, along with lines showing the transition to more than 50% success rate of QCKM (red solid) and, for comparison, of CKM (yellow dotted). This transition happens (except for a deviation at small dimensions) at a constant value of m/nK : as CKM, QCKM requires m to be proportional to n . In this experiment, QCKM requires about 1.13 more measurements than CKMs (complex and full precision) measurements. Fig. 2b is the phase transition for *varying numbers of centroids* K while fixing $n = 5$. Samples are drawn from K Gaussians with means chosen randomly in $\{\pm 1\}^n$, other parameters being identical to the previous experiment. Successful estimation occurs when m scales linearly with K , with a factor of about 1.23 between QCKM and CKM sample complexities. These experiments suggest that CKMs empirical rule $m = \mathcal{O}(nK)$ holds for QCKM, with a slightly higher multiplicative constant.

Real datasets: The performance of QCKM are also assessed on real (and non-Gaussian) data: the spectral clustering (SC) [34] of the MNIST dataset (70000 28×28 pixel images of handwritten digits [35]). This experiment aims at detecting, in an unsupervised setting, the 10 clusters corresponding to the digits 0 – 9 from their representation in a 10-dimensional feature space³. We run the compressive clustering algorithms with $m = 1000$ frequencies. To avoid bad local minima, several replicates of $k\text{-means}$ are usually run and the solution that achieves the best SSE is then selected. We thus also perform several replicates of (Q)CKM, but since computing the SSE requires access to whole dataset (which is not supposed available to the compressive algorithms), we select the solution of CKM (resp. QCKM) minimizing (6) (resp. (10)) [8].

³We thank the authors of [8] for having shared this SC dataset.




Fig. 3: Mean with standard deviation over 100 experiments of the performance (SSE/N and ARI on left and right, respectively) of the different compared clustering algorithms ($k\text{-means}$ in blue, CKM in yellow, QCKM in green), both for 1 and 5 repetitions of the learning algorithms.

We use two performance metrics to assess the clustering algorithms: the SSE in (1), an obvious KMC quality measure, and the Adjusted Rand Index (ARI) [36] that compares the clusters produced by the different algorithms with the ground truth digits. A higher ARI means the clusters are closer to the ground truth, with $ARI = 1$ if the partitions are identical and $ARI = 0$ (on average) if the clusters are assigned at random.

Fig. 3 reports the mean and standard deviation (excluding a few clear outliers for CKM and QCKM, occurring about 5% of the time on average) obtained for both performance metrics. Globally, QCKM performs similarly to CKM, retaining its advantages over $k\text{-means}$. First, the compressive learning algorithms are more stable: their performance exhibit small variance, in contrast with $k\text{-means}$ that hence benefits the most from several replicates. In addition, while for several replicates $k\text{-means}$ outperforms (Q)CKM in terms of SSE, the solutions of (Q)CKM are closer (as quantified by the ARI score) to the ground truth labels: this suggests that the objectives (6) or (10) are better suited than the SSE for (at least) this task. Note that QCKM performances have moderately higher variance than those of CKM: this is probably due to the increased measurement rate of QCKM required to reach similar performance (as suggested by the first experiments) while here both algorithms ran with $m = 1000$.

VI. CONCLUSION

In the context of compressive learning, we have shown that replacing the complex exponential by any periodic function f in the sketch procedure can be compensated by *i*) adding a dithering term to its input, and *ii*) retaining only the first harmonic of f at reconstruction. However, Prop. 1 is valid for *fixed* distributions: future work should provide guarantees for *all* distributions, *e.g.*, belonging to some low-complexity set \mathcal{G} . Still, we believe this result can simplify the design of *low-power sensors acquiring only the minimal information required for some learning task* (*e.g.*, the sketch contribution). To illustrate this idea, we proposed QCKM, a compressive clustering method based on CKM [8] but using hardware-friendly, 1-bit sketches of the learning data, and validated this approach through experiments. However, while yielding promising results in practice, the greedy algorithms that try to solve the non-convex sketch matching optimization problem (*e.g.*, CKM and QCKM) still lack theoretical convergence guarantees. Future work could also consider the binarization of new learning tasks (*e.g.*, Compressive PCA [7]), or explore other sketching mechanisms.

REFERENCES

- [1] H. Steinhaus, "Sur la division des corps matériels en parties," *British Journal of Mathematical and Statistical Psychology*, Classe III, vol. IV, no. 12, 801-804, 1956.
- [2] A. K. Jain, "Data clustering: 50 years beyond K-means," *Pattern Recognition Letters*, vol. 31, no. 8, pp. 651-666, 2010.
- [3] D. Steinley, "K-means clustering: A half-century synthesis," *British Journal of Mathematical and Statistical Psychology*, vol. 59, no. 1, pp. 1-34, 2006.
- [4] D. Aloise, A. Deshpande, P. Hansen, and P. Papat, "NP-hardness of Euclidean sum-of-squares clustering," *Machine Learning*, vol. 75, no. 2, pp. 245-248, 2009.
- [5] S. P. Lloyd, "Least squares quantization in PCM," *IEEE Transactions on Information Theory*, vol. 28, no. 2, pp. 129-137, 1982.
- [6] D. Arthur and S. Vassilvitskii, "k-means++: The Advantages of Careful Seeding," *ACM-SIAM symposium on Discrete algorithms*, pp. 1027-103, 2007.
- [7] R. Gribonval, G. Blanchard, N. Keriven, and Y. Traonmilin, "Compressive Statistical Learning with Random Feature Moments," *arXiv preprint arXiv:1706.07180*, Jun. 2017.
- [8] N. Keriven, N. Tremblay, Y. Traonmilin, and R. Gribonval, "Compressive K-means," *ICASSP 2017 - IEEE International Conference on Acoustics, Speech and Signal Processing*, 2017.
- [9] E. J. Candès, J. K. Romberg, and T. Tao, "Stable signal recovery from incomplete and inaccurate measurements," *Communications on Pure and Applied Mathematics*, vol. 59, no. 8, pp. 1207-1223, 2006.
- [10] A. Chatalic, R. Gribonval, and N. Keriven, "Large-Scale High-Dimensional Clustering with Fast Sketching," *ICASSP 2018 - IEEE International Conference on Acoustics, Speech and Signal Processing*, Apr 2018.
- [11] E. J. Candès and M. B. Wakin, "An Introduction To Compressive Sampling," *IEEE Signal Processing Magazine*, vol. 25, no. 2, pp. 21-30, March 2008.
- [12] A. Saade, F. Caltagirone, I. Carron, L. Daudet, A. Drémeau, S. Gigan, and F. Krzakala, "Random projections through multiple optical scattering: Approximating Kernels at the speed of light," *ICASSP 2016 - IEEE International Conference on Acoustics, Speech and Signal Processing*, 2016, pp. 6215-6219.
- [13] P. T. Boufounos and S. Rane, "Efficient Coding of Signal Distances Using Universal Quantized Embeddings," in *2013 Data Compression Conference*, March 2013, pp. 251-260.
- [14] Y.-G. Yoon, J. Kim, T.-K. Jang, and S. Cho, "A Time-Based Bandpass ADC Using Time-Interleaved Voltage-Controlled Oscillators," *IEEE Transactions on Circuits and Systems I: Regular Papers*, vol. 55, no. 11, pp. 3571-3581, 2008.
- [15] C. Boutsidis, P. Drineas, and M. W. Mahoney, "Unsupervised feature selection for the k-means clustering problem," in *Advances in Neural Information Processing Systems 22*, Y. Bengio, D. Schuurmans, J. D. Lafferty, C. K. I. Williams, and A. Culotta, Eds. Curran Associates, Inc., 2009, pp. 153-161.
- [16] C. Boutsidis, A. Zouzias, and P. Drineas, "Random Projections for k-means Clustering," *Advances in Neural Information Processing Systems*, 2010, pp. 298-306.
- [17] D. Needell, R. Saab, and T. Woolf, "Simple classification using binary data," *CoRR*, vol. abs/1707.01945, 2017.
- [18] E. Dupraz, "K-means Algorithm over Compressed Binary Data," *Data Compression Conference (DCC)*, Mar 2018.
- [19] G. Frahling and C. Sohler, "A fast k-means implementation using coresets," *International Journal of Computational Geometry & Applications*, vol. 18, no. 06, pp. 605-625, 2008.
- [20] R. Chitta, R. Jin, and A. K. Jain, "Efficient Kernel Clustering Using Random Fourier Features," in *2012 IEEE 12th International Conference on Data Mining*, Dec 2012, pp. 161-170.
- [21] A. Rahimi and B. Recht, "Random Features for Large-Scale Kernel Machines," in *Advances in Neural Information Processing Systems 20*, J. C. Platt, D. Koller, Y. Singer, and S. T. Roweis, Eds. Curran Associates, Inc., 2008, pp. 1177-1184.
- [22] A. Gretton, K.M. Borgwardt, M. Rasch, B. Scholkopf, and A.J. Smola "A kernel method for the two-sample-problem", *NIPS 2007 - Advances in neural information processing systems*, 2007, pp. 513-520.
- [23] A. Smola, A. Gretton, L. Song, B. Scholkopf, "A Hilbert space embedding for distributions", *International Conference on Algorithmic Learning Theory*, Springer, Berlin, Heidelberg, 2007.
- [24] B. K. Sriperumbudur, "Mixture density estimation via Hilbert space embedding of measures," in *2011 IEEE International Symposium on Information Theory Proceedings*, July 2011, pp. 1027-1030.
- [25] N. Aronszajn, "Theory of reproducing kernels," *Transactions of the American Mathematical Society*, no. 68, pp. 337-404, 1950.
- [26] N. Keriven, A. Bourrier, R. Gribonval, and P. Pérez, "Sketching for Large-Scale Learning of Mixture Models," *Information and Inference: A Journal of the IMA*, 2017.
- [27] A. R. Hall, *Generalized method of moments*. Oxford University Press, 2005.
- [28] Y. C. Pati, R. Rezaifar, Y. C. P. R. Rezaifar, and P. S. Krishnaprasad, "Orthogonal Matching Pursuit: Recursive Function Approximation with Applications to Wavelet Decomposition," in *Proceedings of the 27th Annual Asilomar Conference on Signals, Systems, and Computers*, 1993, pp. 40-44.
- [29] P. Jain, A. Tewari, and I. S. Dhillon, "Orthogonal Matching Pursuit with Replacement," *Advances in Neural Information Processing Systems*, 2011, pp. 1215-1223.
- [30] P. T. Boufounos, S. Rane, and H. Mansour, "Representation and Coding of Signal Geometry," *Information and Inference: A Journal of the IMA*, vol. 6, no. 4, pp. 349-388, 2017.
- [31] B. K. Sriperumbudur, A. Gretton, K. Fukumizu, B. Schölkopf, and G. R. Lanckriet, "Hilbert Space Embeddings and Metrics on Probability Measures," *J. Mach. Learn. Res.*, vol. 11, pp. 1517-1561, Aug. 2010.
- [32] W. Rudin, *Fourier Analysis on Groups*. Interscience Publishers, 1962.
- [33] N. Keriven, "SketchMLbox, a MATLAB toolbox for large-scale mixture learning," <http://sketchml.gforge.inria.fr/>, Accessed: 2018-01-25.
- [34] U. Von Luxburg, "A tutorial on spectral clustering," *Statistics and computing*, vol. 17, no. 4, pp. 395-416, 2007.
- [35] Y. LeCun, C. Cortes, and C. J. Burges, "The MNIST database of handwritten digits," <http://yann.lecun.com/exdb/mnist/>, Accessed: 2018-01-25.
- [36] N. X. Vinh, J. Epps, and J. Bailey, "Information theoretic measures for clusterings comparison: Variants, properties, normalization and correction for chance," *Journal of Machine Learning Research*, vol. 11, no. Oct, pp. 2837-2854, 2010.