



HAL
open science

QuicK-means: Acceleration of K-means by learning a fast transform

Luc Giffon, Valentin Emiya, Liva Ralaivola, Hachem Kadri

► **To cite this version:**

Luc Giffon, Valentin Emiya, Liva Ralaivola, Hachem Kadri. QuicK-means: Acceleration of K-means by learning a fast transform. 2019. hal-02174845v1

HAL Id: hal-02174845

<https://hal.science/hal-02174845v1>

Preprint submitted on 26 Aug 2019 (v1), last revised 13 Jan 2021 (v4)

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

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

QuicK-means: Acceleration of K-means by learning a fast transform

Luc Giffon¹, Valentin Emiya¹, Liva Ralaivola^{2,1}, and Hachem Kadri¹

¹Aix Marseille Univ, CNRS, LIS, Marseille, France

²Criteo

August 26, 2019

Abstract

K-means— and the celebrated Lloyd algorithm — is more than the clustering method it was originally designed to be. It has indeed proven pivotal to help increase the speed of many machine learning and data analysis techniques such as indexing, nearest-neighbor search and prediction, data compression; its beneficial use has been shown to carry over to the acceleration of kernel machines (when using the Nyström method). Here, we propose a fast extension of **K-means**, dubbed **QuicK-means**, that rests on the idea of expressing the matrix of the K centroids as a product of sparse matrices, a feat made possible by recent results devoted to find approximations of matrices as a product of sparse factors. Using such a decomposition squashes the complexity of the matrix-vector product between the factorized $K \times D$ centroid matrix \mathbf{U} and any vector from $\mathcal{O}(KD)$ to $\mathcal{O}(A \log A + B)$, with $A = \min(K, D)$ and $B = \max(K, D)$, where D is the dimension of the training data. This drastic computational saving has a direct impact in the assignment process of a point to a cluster, meaning that it is not only tangible at prediction time, but also at training time, provided the factorization procedure is performed during Lloyd’s algorithm. We precisely show that resorting to a factorization step at each iteration does not impair the convergence of the optimization scheme and that, depending on the context, it may entail a reduction of the training time. Finally, we provide discussions and numerical simulations that show the versatility of our computationally-efficient **QuicK-means** algorithm.

1 Introduction

K-means is one of the most popular clustering algorithms [5, 6]. It can be used beyond clustering, for other tasks such as indexing, data compression, nearest-neighbor search and prediction, and local network community detection [13, 20]. **K-means** is also a pivotal process to help increase the speed and the accuracy of many machine learning techniques such as the Nyström approximation of kernel machines [19] and RBF networks [16]. The conventional **K-means** algorithm has a complexity of $\mathcal{O}(NKD)$ per iteration, where N is the number of data points, K the number of clusters and D is the dimension of the data points. However, the larger the number of clusters, the more iterations are needed to converge [1]. As data dimensionality and data sample size continue to grow, it is critical to produce viable and cost-effective alternatives to the computationally expensive conventional **K-means**. Previous attempts to alleviate the computational issues in **K-means** often relied on batch-, sparsity- and randomization-based methods [17, 3, 18, 11].

Fast transforms have recently received increased attention in machine learning community as they can be used to speed up random projections [8, 4] and to improve landmark-based approximations [19]. These works primarily focused on fast transforms such as Fourier and Hadamard transforms, which are fixed before the learning begins. An interesting question is whether one can go beyond that and learn the fast transform from data. In a recent paper [9], the authors introduced a sparse matrix approximation scheme aimed at reducing the complexity of applying linear operators in high dimension by approximately factorizing the corresponding

Symbol	Meaning
$\llbracket M \rrbracket$	set of integers from 1 to M
$\ \cdot\ $	L_2 -norm
$\ \cdot\ _F$	Frobenius norm
$\ \cdot\ _0$	L_0 -norm
$\ \cdot\ _2$	spectral norm
$\mathbf{D}_\mathbf{v}$	diagonal matrix with vector \mathbf{v} on the diagonal
N	number of data points
D	data dimension
K	number of clusters
Q	number of sparse factors
$\mathbf{x}_1, \dots, \mathbf{x}_N$	data points
$\mathbf{X} \in \mathbb{R}^{N \times D}$	data matrix
\mathbf{t}	cluster assignment vector
$\mathbf{u}_1, \dots, \mathbf{u}_K$	K-means centroids
$\mathbf{U} \in \mathbb{R}^{K \times D}$	K-means centroid matrix
$\mathbf{v}_1, \dots, \mathbf{v}_K$	QK-means centroids
$\mathbf{V} \in \mathbb{R}^{K \times D}$	QK-means centroid matrix
$\mathbf{S}_1, \dots, \mathbf{S}_Q$	sparse matrices
$\mathcal{E}_1, \dots, \mathcal{E}_Q$	sparsity constraint sets
$\delta_{\mathcal{E}}$	indicator functions for set \mathcal{E}
τ	current iteration

Table 1: Notation used in this paper.

matrix into few sparse factors. One interesting observation is that fast transforms, such as the Hadamard transform and the Discrete Cosine transform, can be exactly or approximately decomposed as a product of sparse matrices. In this paper, we take this idea further and investigate attractive and computationally less costly implementations of the **K-means** algorithm by learning a fast transform from data. Specifically, we make the following contributions:

- we introduce **QuickK-means**, a fast extension of **K-means** that rests on the idea of expressing the matrix of the K centroids as a product of sparse matrices, a feat made possible by recent results devoted to find approximations of matrices as a product of sparse factors,
- we show that each update step in one iteration of our algorithm reduces the overall objective, which is enough to guarantee the convergence of **QuickK-means**,
- we perform a complexity analysis of our algorithm, showing that the computational gain in **QuickK-means** has a direct impact in the assignment process of a point to a cluster, meaning that it is not only tangible at prediction time, but also at training time,
- we provide an empirical evaluation of **QuickK-means** performance which demonstrates its effectiveness on different datasets in the contexts of clustering and kernel Nyström approximation.

2 Preliminaries

We briefly review the basics of **K-means** and give background on learning fast transforms. To assist the reading, we list the notations used in the paper in Table 1.

2.1 K-means

The **K-means** algorithm is used to partition a set $\mathbf{X} = \{\mathbf{x}_1, \dots, \mathbf{x}_N\}$ of N vectors $\mathbf{x}_n \in \mathbb{R}^D$ into a predefined number K of clusters with the aim of minimizing the distance between each \mathbf{x}_n to the center $\mathbf{u}_k \in \mathbb{R}^D$ of

the cluster k it belongs to—the center \mathbf{u}_k of cluster k is the mean vector of the points assigned to cluster k . K-means attempts to solve

$$\arg \min_{\mathbf{U}, \mathbf{t}} \sum_{k \in \llbracket K \rrbracket} \sum_{n: t_n = k} \|\mathbf{x}_n - \mathbf{u}_k\|^2, \quad (1)$$

where $\mathbf{U} = \{\mathbf{u}_1, \dots, \mathbf{u}_K\}$ is the set of cluster centers and $\mathbf{t} \in \llbracket K \rrbracket^N$ is the assignment vector that puts \mathbf{x}_n in cluster k if $t_n = k$.

Lloyd’s algorithm. The most popular procedure to (approximately) solve the K-means problem is the iterative Lloyd’s algorithm, which alternates i) an assignment step that decides the current cluster to which each point \mathbf{x}_n belongs and ii) a reestimation step which refines the clusters and their centers. In little more detail, the algorithm starts with an initialized set of K cluster centers $\mathbf{U}^{(0)}$ and proceeds as follows: at iteration τ , the assignments are updated as

$$\forall n \in \llbracket N \rrbracket, t_n^{(\tau)} \leftarrow \arg \min_{k \in \llbracket K \rrbracket} \left\| \mathbf{x}_n - \mathbf{u}_k^{(\tau-1)} \right\|_2^2 = \arg \min_{k \in \llbracket K \rrbracket} \left\| \mathbf{u}_k^{(\tau-1)} \right\|_2^2 - 2 \left\langle \mathbf{u}_k^{(\tau-1)}, \mathbf{x}_n \right\rangle, \quad (2)$$

the reestimation of the cluster centers is performed as

$$\forall k \in \llbracket K \rrbracket, \mathbf{u}_k^{(\tau)} \leftarrow \hat{\mathbf{x}}_k(\mathbf{t}^{(\tau)}) := \frac{1}{n_k^{(\tau)}} \sum_{n: t_n^{(\tau)} = k} \mathbf{x}_n \quad (3)$$

where $n_k^{(\tau)} := |\{n : t_n^{(\tau)} = k\}|$ is the number of points in cluster k at time τ and $\hat{\mathbf{x}}_k(\mathbf{t})$ is the mean vector of the elements of cluster k according to assignment \mathbf{t} .

Complexity of Lloyd’s algorithm. The assignment step (2) costs $\mathcal{O}(NDK)$ operations while the update of the centers (3) costs $\mathcal{O}(ND)$ operations. Hence, the bottleneck of the overall time complexity $\mathcal{O}(NDK)$ stems from the assignment step. Once the clusters have been defined, assigning N' new points to these clusters is performed via (2) at the cost of $\mathcal{O}(N'DK)$ operations.

The main contribution in this paper relies on the idea that (2) may be computed more efficiently by approximating \mathbf{U} as a fast operator.

2.2 Learning Fast Transforms as the Product of Sparse Matrices

Structured linear operators as products of sparse matrices. The popularity of some linear operators from \mathbb{R}^M to \mathbb{R}^M (with $M < \infty$) like Fourier or Hadamard transforms comes from both their mathematical properties and their ability to compute the mapping of some input $\mathbf{x} \in \mathbb{R}^M$ with efficiency, typically in $\mathcal{O}(M \log(M))$ rather than in $\mathcal{O}(M^2)$ operations. The main idea of the related fast algorithms is that the matrix $\mathbf{U} \in \mathbb{R}^{M \times M}$ characterizing such linear operators can be written as the product $\mathbf{U} = \prod_{q \in \llbracket Q \rrbracket} \mathbf{S}_q$ of Q sparse matrices \mathbf{S}_q , with $Q = \mathcal{O}(\log M)$ factors and $\|\mathbf{S}_q\|_0 = \mathcal{O}(M)$ non-zero coefficients per factor [9, 12]: for any vector $\mathbf{x} \in \mathbb{R}^M$, $\mathbf{U}\mathbf{x}$ can thus be computed as $\mathcal{O}(\log M)$ products $\mathbf{S}_0(\mathbf{S}_1(\dots(\mathbf{S}_{Q-1}\mathbf{x})))$ between a sparse matrix and a vector, the cost of each product being $\mathcal{O}(M)$. This gives a $\mathcal{O}(M \log M)$ time complexity for computing $\mathbf{U}\mathbf{x}$ in that case.

Learning a computationally-efficient decomposition approximating an arbitrary operator. When the linear operator \mathbf{U} is an arbitrary matrix, one may approximate it with such a sparse-product structure by learning the factors $\{\mathbf{S}_q\}_{q \in \llbracket Q \rrbracket}$ in order to benefit from a fast algorithm. A recent contribution [9] has proposed algorithmic strategies to learn such a factorization. Based on the proximal alternating linearized minimization (PALM) algorithm [2], the PALM for Multi-layer Sparse Approximation (palm4MSA) algorithm [9] aims at approximating a matrix $\mathbf{U} \in \mathbb{R}^{K \times D}$ as a product of sparse matrices by solving

$$\min_{\{\mathbf{S}_q\}_{q \in \llbracket Q \rrbracket}} \left\| \mathbf{U} - \prod_{q \in \llbracket Q \rrbracket} \mathbf{S}_q \right\|_F^2 + \sum_{q \in \llbracket Q \rrbracket} \delta_{\mathcal{E}_q}(\mathbf{S}_q) \quad (4)$$

Algorithm 1 QK-means algorithm and its time complexity.

Require: $\mathbf{X} \in \mathbb{R}^{N \times D}$, K , initialization $\{\mathbf{S}_q^{(0)} : \mathbf{S}_q^{(0)} \in \mathcal{E}_q\}_{q \in \llbracket Q \rrbracket}$ $A := \min(K, D)$

1: Set $\mathbf{V}^{(0)} : \mathbf{x} \mapsto \prod_{q \in \llbracket Q \rrbracket} \mathbf{S}_q^{(0)} \mathbf{x}$ $B := \max(K, D)$

2: **for** $\tau = 1, 2, \dots$ until convergence **do**

3: $\mathbf{t}^{(\tau)} := \arg \min_{\mathbf{t} \in \llbracket K \rrbracket^N} \sum_{n \in \llbracket N \rrbracket} \|\mathbf{x}_n - \mathbf{v}_{t_n}^{(\tau-1)}\|^2$ $\mathcal{O}(N(A \log A + B) + AB)$

4: $\forall k \in \llbracket K \rrbracket, \mathbf{u}_k := \frac{1}{n_k} \sum_{n: t_n^{(\tau)} = k} \mathbf{x}_n$ with $n_k := |\{n : t_n^{(\tau)} = k\}|$ $\mathcal{O}(ND)$

5: $\mathbf{A} := \mathbf{D}_{\sqrt{\mathbf{n}}} \times \mathbf{U}$ $\mathcal{O}(KD)$

6: $\mathcal{E}_0 := \{\mathbf{D}_{\sqrt{\mathbf{n}}}\}$

7: $\{\mathbf{S}_q^{(\tau)}\}_{q=0}^Q := \arg \min_{\{\mathbf{S}_q\}_{q=0}^Q} \|\mathbf{A} - \prod_{q=0}^Q \mathbf{S}_q\|_F^2 + \sum_{q=0}^Q \delta_{\mathcal{E}_q}(\mathbf{S}_q)$
 $\mathcal{O}(AB(\log^2 A + \log B))$ (or $\mathcal{O}(AB(\log^3 A + \log A \log B))$)

8: Set $\mathbf{V}^{(\tau)} : \mathbf{x} \mapsto \prod_{q \in \llbracket Q \rrbracket} \mathbf{S}_q^{(\tau)} \mathbf{x}$ $\mathcal{O}(1)$

9: **end for**

Ensure: assignment vector \mathbf{t} and sparse matrices $\{\mathbf{S}_q : \mathbf{S}_q \in \mathcal{E}_q\}_{q \in \llbracket Q \rrbracket}$ such that $\prod_{q \in \llbracket Q \rrbracket} \mathbf{S}_q \approx \mathbf{U}$ the K means of the N data points

where, for each $q \in \llbracket Q \rrbracket$, $\delta_{\mathcal{E}_q}(\mathbf{S}_q) = 0$ if $\mathbf{S}_q \in \mathcal{E}_q$ and $\delta_{\mathcal{E}_q}(\mathbf{S}_q) = +\infty$ otherwise, \mathcal{E}_q being a constraint set that typically impose a sparsity structure on its elements, as well as a scaling constraint. The `palM4MSA` algorithm and more related details are given in Appendix A.

Although this problem is non-convex and the computation of a global optimum cannot be ascertained, the `palM4MSA` algorithm is able to find good local minima with convergence guarantees.

3 Quick-means

We here introduce our main contribution, `Quick-means` (abbreviated by `QK-means`), show its convergence property and analyze its computational complexity.

3.1 QK-means: Encoding Centroids as Products of Sparse Matrices

`Quick-means` is a variant of the `K-means` algorithm in which the matrix of centroids \mathbf{U} is approximated as a product $\mathbf{V} = \prod_{q \in \llbracket Q \rrbracket} \mathbf{S}_q$ of sparse matrices \mathbf{S}_q . Doing so will allow us to cope with the computational bulk imposed by the product $\mathbf{U}\mathbf{x}$ (cf. (2)) at the core of the cluster assignment process.

Building upon the `K-means` optimization problem (1) and fast-operator approximation problem (4) the `QK-means` optimization problem writes:

$$\arg \min_{\mathbf{S}_1, \dots, \mathbf{S}_Q, \mathbf{t}} g(\mathbf{S}_1, \dots, \mathbf{S}_Q, \mathbf{t}) := \sum_{k \in \llbracket K \rrbracket} \sum_{n: t_n = k} \|\mathbf{x}_n - \mathbf{v}_k\|^2 + \sum_{q \in \llbracket Q \rrbracket} \delta_{\mathcal{E}_q}(\mathbf{S}_q) \text{ s. t. } \mathbf{V} = \prod_{q \in \llbracket Q \rrbracket} \mathbf{S}_q \quad (5)$$

This is a regularized version of the `K-means` optimization problem (1) in which centroids \mathbf{v}_k are constrained to form a matrix \mathbf{V} with a fast-operator structure, the indicator functions $\delta_{\mathcal{E}_q}$ imposing the sparsity of matrices \mathbf{S}_q . More details on the modeling choices are given in the experimental part in section 4.1.

This problem can be solved using Algorithm 1, which proceeds in a similar way as Lloyd's algorithm by alternating an assignment step at line 3 and an update of the centroids at lines 4–8. The assignment step can be computed efficiently thanks to the fast-structure in matrix \mathbf{V} . The update of the centroids relies on learning a fast-structure operator \mathbf{V} that approximate of the true centroid matrix \mathbf{U} weighted by the number of examples n_k assigned to each cluster k .

3.2 Convergence of QK-means

Similarly to K-means, QK-means converges locally as stated in the following proposition.

Proposition (Convergence of QK-means). *The iterates $\{\mathbf{S}^{(\tau)}\}_{q \in \llbracket Q \rrbracket}$ and $\mathbf{t}^{(\tau)}$ in Algorithm 1 are such that the values*

$$g(\mathbf{S}_1^{(\tau)}, \dots, \mathbf{S}_Q^{(\tau)}, \mathbf{t}^{(\tau)}) = \sum_{k \in \llbracket K \rrbracket} \sum_{n: \mathbf{t}_n^{(\tau)} = k} \left\| \mathbf{x}_n - \mathbf{v}_k^{(\tau)} \right\|_2^2 + \sum_{q \in \llbracket Q \rrbracket} \delta_{\mathcal{E}_q}(\mathbf{S}_q^{(\tau)}) \quad \text{s.t. } \mathbf{V} = \prod_{q \in \llbracket Q \rrbracket} \mathbf{S}_q^{(\tau)} \quad (6)$$

of the objective function are non-increasing.

Proof. To prove this convergence, we show that each of the assignment and centroid update steps in one iteration τ of the algorithm actually reduces the overall objective.

Assignment step (Line 3) For a fixed $\mathbf{V}^{(\tau-1)}$, the optimization problem at Line 3 is separable for each example indexed by $n \in \llbracket N \rrbracket$ and the new indicator vector $\mathbf{t}^{(\tau)}$ is thus defined as:

$$t_n^{(\tau)} = \arg \min_{k \in \llbracket K \rrbracket} \left\| \mathbf{x}_n - \mathbf{v}_k^{(\tau-1)} \right\|_2^2. \quad (7)$$

This step minimizes the first term in (6) w.r.t. \mathbf{t} while the second term is constant so we have

$$g(\mathbf{S}_1^{(\tau-1)}, \dots, \mathbf{S}_Q^{(\tau-1)}, \mathbf{t}^{(\tau)}) \leq g(\mathbf{S}_1^{(\tau-1)}, \dots, \mathbf{S}_Q^{(\tau-1)}, \mathbf{t}^{(\tau-1)}).$$

Centroids update step (Lines 4–8). We now consider a fixed assignment vector \mathbf{t} . We first note that for any cluster k with true centroid \mathbf{u}_k and approximated centroid \mathbf{v}_k , we have

$$\begin{aligned} \sum_{n: t_n = k} \left\| \mathbf{x}_n - \mathbf{v}_k \right\|^2 &= \sum_{n: t_n = k} \left\| \mathbf{x}_n - \mathbf{u}_k + \mathbf{u}_k - \mathbf{v}_k \right\|^2 \\ &= \sum_{n: t_n = k} \left(\left\| \mathbf{x}_n - \mathbf{u}_k \right\|^2 + \left\| \mathbf{u}_k - \mathbf{v}_k \right\|^2 - 2 \langle \mathbf{x}_n - \mathbf{u}_k, \mathbf{u}_k - \mathbf{v}_k \rangle \right) \\ &= \sum_{n: t_n = k} \left\| \mathbf{x}_n - \mathbf{u}_k \right\|^2 + n_k \left\| \mathbf{u}_k - \mathbf{v}_k \right\|^2 - 2 \underbrace{\left\langle \sum_{n: t_n = k} (\mathbf{x}_n - \mathbf{u}_k), \mathbf{u}_k - \mathbf{v}_k \right\rangle}_{=0} \\ &= \sum_{n: t_n = k} \left\| \mathbf{x}_n - \mathbf{u}_k \right\|^2 + \|\sqrt{n_k}(\mathbf{u}_k - \mathbf{v}_k)\|^2 \end{aligned}$$

For a fixed \mathbf{t} , the new sparsely-factorized centroids are solutions of the following subproblem:

$$\begin{aligned} \arg \min_{\mathbf{S}_1, \dots, \mathbf{S}_Q} g(\mathbf{S}_1, \dots, \mathbf{S}_Q, \mathbf{t}) &= \arg \min_{\mathbf{S}_1, \dots, \mathbf{S}_Q} \sum_{k \in \llbracket K \rrbracket} \sum_{n: t_n = k} \left\| \mathbf{x}_n - \mathbf{v}_k \right\|_2^2 + \sum_{q \in \llbracket Q \rrbracket} \delta_q(\mathbf{S}_q) \quad \text{s.t. } \mathbf{V} = \prod_{q \in \llbracket Q \rrbracket} \mathbf{S}_q \\ &= \arg \min_{\mathbf{S}_1, \dots, \mathbf{S}_Q} \left\| \mathbf{D}_{\sqrt{\mathbf{n}}}(\mathbf{U} - \mathbf{V}) \right\|_F^2 + \sum_{k \in \llbracket K \rrbracket} c_k + \sum_{q \in \llbracket Q \rrbracket} \delta_q(\mathbf{S}_q) \quad \text{s.t. } \mathbf{V} = \prod_{q \in \llbracket Q \rrbracket} \mathbf{S}_q \\ &= \arg \min_{\mathbf{S}_1, \dots, \mathbf{S}_Q} \left\| \mathbf{A} - \mathbf{D}_{\sqrt{\mathbf{n}}} \prod_{q \in \llbracket Q \rrbracket} \mathbf{S}_q \right\|_F^2 + \sum_{q \in \llbracket Q \rrbracket} \delta_q(\mathbf{S}_q) \end{aligned} \quad (8)$$

where :

- $\sqrt{\mathbf{n}} \in \mathbb{R}^K$ is the pair-wise square root of the vector indicating the number of observations $n_k := |\{n : t_n = k\}|$ in each cluster k ;

- $\mathbf{D}_{\sqrt{\mathbf{n}}} \in \mathbb{R}^{K \times K}$ refers to a diagonal matrix with vector $\sqrt{\mathbf{n}}$ on the diagonal;
- $\mathbf{U} \in \mathbb{R}^{K \times d}$ refers to the unconstrained centroid matrix obtained from the data matrix \mathbf{X} and the indicator vector \mathbf{t} : $\mathbf{u}_k := \frac{1}{n_k} \sum_{n:t_n=k} \mathbf{x}_n$ (see Line 4);
- $\mathbf{D}_{\sqrt{\mathbf{n}}}(\mathbf{U} - \mathbf{V})$ is the matrix with $\sqrt{n_k}(\mathbf{u}_k - \mathbf{v}_k)$ as k -th row;
- $c_k := \sum_{n:t_n=k} \|\mathbf{x}_n - \mathbf{u}_k\|$ is constant w.r.t. $\mathbf{S}_1, \dots, \mathbf{S}_Q$;
- $\mathbf{A} := \mathbf{D}_{\sqrt{\mathbf{n}}}\mathbf{U}$ is the unconstrained centroid matrix reweighted by the size of each cluster (see Line 5).

A local minimum of (8) is obtained by applying the `palM4MSA` algorithm or its hierarchical variant to approximate \mathbf{A} , as in Line 7. The first factor is forced to equal $\mathbf{D}_{\sqrt{\mathbf{n}}}$ by setting \mathcal{E}_0 to a singleton at Line 6. Using the previous estimate $\{\mathbf{S}_q^{(\tau-1)}\}_{q \in [Q]}$ to initialize this local minimization, we thus obtain that $g(\mathbf{S}_1^{(\tau)}, \dots, \mathbf{S}_Q^{(\tau)}, \mathbf{t}^{(\tau)}) \leq g(\mathbf{S}_1^{(\tau-1)}, \dots, \mathbf{S}_Q^{(\tau-1)}, \mathbf{t}^{(\tau)})$.

We finally have, for any τ ,

$$\begin{aligned} g(\mathbf{S}_1^{(\tau)}, \dots, \mathbf{S}_Q^{(\tau)}, \mathbf{t}^{(\tau)}) &\leq g(\mathbf{S}_1^{(\tau-1)}, \dots, \mathbf{S}_Q^{(\tau-1)}, \mathbf{t}^{(\tau)}) \leq g(\mathbf{S}_1^{(\tau-1)}, \dots, \mathbf{S}_Q^{(\tau-1)}, \mathbf{t}^{(\tau-1)}) \\ &\leq \dots \leq g(\mathbf{S}_1^{(0)}, \dots, \mathbf{S}_Q^{(0)}, \mathbf{t}^{(0)}) \end{aligned}$$

□

3.3 Complexity analysis

Since the space complexity of the proposed `QK-means` algorithm is comparable to that of `K-means`, we only detail its time complexity. We set $A = \min(K, D)$ and $B = \max(K, D)$, and assume that the number of factors satisfies $Q = \mathcal{O}(\log A)$.

The analysis is proposed under the following assumptions: the product between two dense matrices of shapes $N_1 \times N_2$ and $N_2 \times N_3$ can be done $\mathcal{O}(N_1 N_2 N_3)$ operations; the product between a sparse matrix with $\mathcal{O}(S)$ non-zero entries and a dense vector can be done in $\mathcal{O}(S)$ operations; the product between two sparse matrices of shapes $N_1 \times N_2$ and $N_2 \times N_3$, both having $\mathcal{O}(S)$ non-zero values can be done in $\mathcal{O}(S \min(N_1, N_3))$ and the number of non-zero entries in the resulting matrix is $\mathcal{O}(S^2)$.

Complexity of the `K-means` algorithm. We recall here that the `K-means` algorithm complexity is dominated by its cluster assignation step which requires $\mathcal{O}(NKD) = \mathcal{O}(NAB)$ operations (see Eq. (2)).

Complexity of algorithm `palM4MSA`. The procedure consists in an alternate optimization of each sparse factor. At each iteration, the whole set of Q factors is updated with at a cost in $\mathcal{O}(AB(\log^2 A + \log B))$, as detailed in Appendix A. The bottleneck is the computation of the gradient, which benefits from fast computations with sparse matrices. The hierarchical version of `palM4MSA` proposed in [9] consists in running `palM4MSA` $2Q$ times so that its time complexity is in $\mathcal{O}(AB(\log^3 A + \log A \log B))$.

Complexity of the `QK-means` algorithm. The overall complexity of `QK-means` is in $\mathcal{O}(N(A \log A + B) + AB \log^2 A)$ when used with `palM4MSA` and in $\mathcal{O}(N(A \log A + B) + AB \log^3 A)$ when used with the hierarchical version of `palM4MSA`. The time complexities of the main steps are given in Algorithm 1.

The assignation step (line 3 and Eq. (2)) benefits from the fast computation of \mathbf{VX} in $\mathcal{O}(N(A \log A + B))$ while the computation of the norms of the cluster centers is in $\mathcal{O}(AB)$. One can see that the computational bottleneck of `K-means` is here reduced, which shows the advantage of using `QK-means` when N , K and D are large.

The computation of the centers of each cluster, given in line 4, is the same as in `K-means` and takes $\mathcal{O}(ND)$ operations.

The update of the fast transform, in lines 5 to 8 is a computational overload compared to `K-means`. Its time complexity is dominated by the update of the sparse factors at line 7, in $\mathcal{O}(AB \log^2 A)$ if `palm4MSA` is called and in $\mathcal{O}(AB \log^3 A)$ if its hierarchical version is called. Note that this cost is dominated by the cost of the assignment step as soon as the number of examples N is greater than $\log^3 A$.

4 Experiments and applications

4.1 Experimental setting

Implementation details. The simulations have been conducted in Python, including for the `palm4MSA` algorithm. Running times are measured on computer grid with 3.8GHz-CPU (2.5GHz in Figure 1). Fast operators \mathbf{V} based on sparse matrices \mathbf{S}_q are implemented with `csr_matrix` objects from the `scipy.linalg` package. While more efficient implementations may be beneficial for larger deployment, our implementation is sufficient as a proof of concept for assessing the performance of the proposed approach. In particular, the running times of fast operators of the form $\prod_{q \in \llbracket Q \rrbracket} \mathbf{S}_q$ have been measured when applying to random vectors, for several sparsity levels: as shown in Figure 1, they are significantly faster than dense operators – implemented as a `numpy.ndarray` matrix –, especially when the data size is larger than 10^3 .




Figure 1: Running times, averaged over 30 runs, when applying dense or fast $D \times D$ operators to a set of 100 random vectors. The number of factors in fast operators equals $\log_2(D)$ and the sparsity level denotes the number of non-zero coefficients per row and per column in each factor.

Datasets. We present results on real-world and toy datasets summarized in Table 2. On the one hand, the real world datasets `MNIST` [10] and `Fashion-Mnist` [15] are used to show — quantitatively and qualitatively — the good quality of the obtained centroids when using our method `QK-means`. On the other hand, we use the `blobs` synthetic dataset from `sklearn.dataset` to show the speed up offered by our method `QK-means` when the number of clusters and the dimensionality of the data are sufficiently large.

Algorithm settings. The `QK-means` algorithm is used with $Q := \log_2(A)$ sparse factors, where $A = \min(K, D)$. All factors \mathbf{S}_q are with shape $A \times A$ except, depending on the shape of \mathbf{A} , the leftmost one ($K \times A$) or the rightmost one ($A \times D$). The sparsity constraint of each factor \mathbf{S}_q is set in \mathcal{E}_q and is governed

Dataset	Data dim. D	# classes	Training set size N	Test set size N'
MNIST	784	10	60 000	10 000
Fashion-MNIST	784	10	60 000	10 000
Blobs (clusters std: 12)	2000	1000	29000	1000


Table 2: Datasets statistics

by a global parameter denoted as *sparsity level*, which indicates the desired number of non-zero coefficients in each row and in each column of \mathbf{S}_q . Since the projection onto this set of structured-sparsity constraints may be computationally expensive, this projection is relaxed in the implementation of `palm4MSA` and only guarantees that the number of non-zero coefficients in each row and each column is at least the sparsity level, as in [9]. The actual number of non-zero coefficients in the sparse factors is measured at the end of the optimization process and reported in the results. The sparse factors are updated using the `palm4MSA` rather than its hierarchical version, since we observed that this was a better choice in terms of computational cost, with satisfying approximation results (See Figure 2a and 2b). Additional details about `palm4MSA` are given in Appendix A. The stopping criterion of `K-means` and `QK-means` consists of a tolerance set to 10^{-6} on the relative variation of the objective function and a maximum number of iterations set to 10 for the `Blobs` dataset and to 20 for others. The same principle governs the stopping criterion of `palm4MSA` with a tolerance set to 10^{-6} and a maximum number of iterations set to 300. Each experiment have been replicated using different seed values for random initialisation. Competing techniques share the same seed values, hence share the same initialisation of centroids.


4.2 Clustering

Approximation quality. One important question is the ability of the fast-structure model to fit arbitrary data. Indeed, no theoretical result about the expressivity of such models is currently available. In order to assess this approximation quality, the MNIST and Fashion-MNIST data have been clustered into $K = 30$ clusters by `K-means`, `QK-means` and a variant of `QK-means` using the hierarchical version of `palm4MSA`, with several sparsity levels. Results are reported in Figure 2. In Figures 2a and 2b, one can observe that the objective function of `QK-means` is decreasing in a similar way as `K-means` over iterations. In particular, the use of the fast-structure model does not seem to increase the number of iteration necessary before convergence. At the end of the iterations, the value of objective function for `QK-means` is slightly above that of `K-means`. As expected, the sparser the model, the more degradation in the objective function. However, even very sparse models do not degrade the results significantly. These Figures also demonstrate the convergence property of the `QK-means` algorithm when using the standard, proved convergent, *Palm4MSA* algorithm: in this case, the objective function is always non-increasing whereas the `QK-means` version with *Hierarchical Palm4MSA*, not guaranteed to converge, suffers a small bump in its objective function (see Figure 2b iteration 6). The approximation quality can be assessed visually, in a more subjective and interpretable way, in Figures 2c to 2h where the obtained centroids are displayed as images. Although some degradation may be observed in some images, one can note that each image obtained with `QK-means` clearly represents a single visual item without noticeable interference with other items.


Clustering assignation times. Higher dimensions are required to assess the computational benefits of the proposed approach, as shown here. The assignation times of the clustering procedure were measured on the `Blobs` dataset. The centroid matrices are with shape $K \times D$ with $D = 2000$ and $K \in \{128, 256, 512\}$. Results reported in Figure 3 show that in this setting and with the current implementation, the computational advantage of `QK-means` is observed in high dimension, for $K = 256$ and $K = 512$ clusters. It is worth noticing that when K increases, the running times are not affected that much for `QK-means` while it significantly grows for `K-means`. These trends are directly related to the number of model parameters that are reported in the figure.




(a) MNIST, $K = 30$: objective function.



(b) Fashion-MNIST, $K = 30$: objective function.




(c) K-means centroids.



(d) K-means centroids.




(e) QK-means centroids.



(f) QK-means centroids.



(g) Hierarchical-palm4MSA QK-means centroids.



(h) Hierarchical-palm4MSA QK-means centroids.

Figure 2: Clustering results on MNIST (left) and Fashion-MNIST (right) for $K = 30$ clusters.

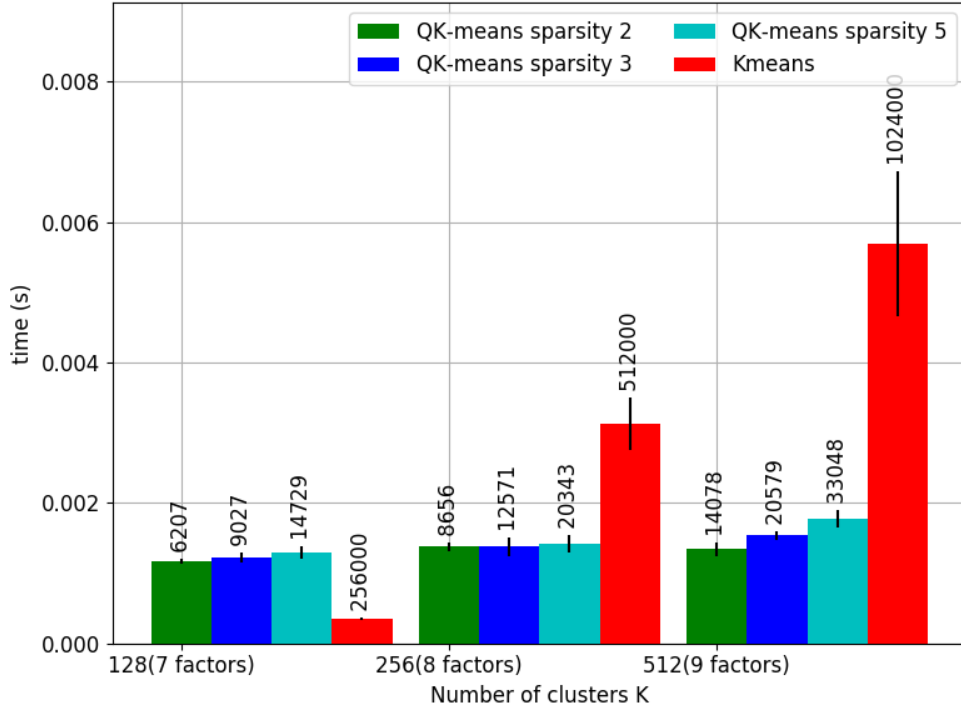


Figure 3: Clustering Blobs data: running times of the assignment step, averaged over 5 runs. The vertical black lines are the standard deviation w.r.t. the runs and the average number of parameters actually learned in the models are reported above those lines.

4.3 Nearest-neighbor search in a large dataset

The Nearest-neighbor search is a fundamental task that suffers from computational limitations when the dataset is large. Fast strategies have been proposed, e.g., using kd trees or ball trees. One may also use a clustering strategy to perform an approximate nearest-neighbor search: the query is first compared to K centroids computed beforehand by clustering the whole dataset, and the nearest neighbor search is then performed among a lower number of data points, within the related cluster. We compare this strategy using **K-means** and **QK-means** against the **scikit-learn** implementation [15] of the nearest-neighbor search (brute force search, kd tree, ball tree). Inference time results on the **Blobs** dataset are reported in Figure 4 and accuracy results are displayed in Table 3. The running times reported in Figure 4 show a dramatic advantage of using a clustering-based approximate search and this advantage is even stronger with the clustering obtained by our **QK-means** method. This speed-up comes at a cost though, we can see a drop in classification performance in Table 3.

4.4 Nyström approximation

In this sub-section, we show how we can take advantage of the fast-operator obtained as output of our **QK-means** algorithm in order to speed-up the computation in the Nyström approximation. We start by giving background knowledge on the Nyström approximation then we present some recent work aiming at accelerating it using well know fast-transform method. We finally stem on this work to present a novel approach based on our **QK-means** algorithm.

	Accuracy Blobs
1NN Brute force search	N/A
1NN KD Tree	N/A
1NN Ball Tree	N/A
<hr/>	
1NN K-means 128 Clusters	0.96
1NN K-means 256 Clusters	0.97
1NN K-means 512 Clusters	0.99
<hr/>	
1NN QK-means 128 Clusters	0.74
1NN QK-means 256 Clusters	0.66
1NN QK-means 512 Clusters	0.66
<hr/>	
Nyström K-means + SVM 128 Clusters	0.98
Nyström K-means + SVM 256 Clusters	1.0
Nyström K-means + SVM 512 Clusters	1.0
<hr/>	
Nyström QK-means + SVM 128 Clusters	0.95
Nyström QK-means + SVM 256 Clusters	1.0
Nyström QK-means + SVM 512 Clusters	1.0

Table 3: Results on the classification task on **Blobs** dataset. Results are averaged over 5 runs. “N/A” denotes experiments that did not finish. Only results with sparsity value 5 are displayed for **QK-means** experiments. For the **QK-means** results, only those obtained with sparsity level = 5 are displayed.

	Accuracy Fashion-MNIST	Accuracy MNIST
1NN Brute force search	0.85	0.97
1NN KD Tree	0.85	0.97
1NN Ball Tree	0.85	0.97
<hr/>		
1NN K-means 10 Clusters	0.84	0.96
1NN K-means 16 Clusters	0.84	0.96
1NN K-means 30 Clusters	0.84	0.96
<hr/>		
1NN QK-means 10 Clusters	0.84	0.96
1NN QK-means 16 Clusters	0.84	0.96
1NN QK-means 30 Clusters	0.84	0.96
<hr/>		
Nyström K-means + SVM 10 Clusters	0.71	0.74
Nyström K-means + SVM 16 Clusters	0.75	0.83
Nyström K-means + SVM 30 Clusters	0.78	0.88
<hr/>		
Nyström QK-means + SVM 10 Clusters	0.71	0.74
Nyström QK-means + SVM 16 Clusters	0.74	0.82
Nyström QK-means + SVM 30 Clusters	0.77	0.88

Table 4: Results on the classification task on the **MNIST** and **Fashion-MNIST** datasets. Results are averaged over 5 runs. “N/A” denotes experiments that did not finish. For the **QK-means** results, only those obtained with sparsity level = 5 are displayed.




Figure 4: Running time of nearest neighbor search on blobs data. Results are averaged over 5 runs (vertical lines: standard deviation) and the average number of parameters actually learned is reported above each bar. The results for the Brute Force Search, KD Tree and Ball Tree are not displayed because they were longer than 10 times the K-means search version.

4.4.1 Background on the Nyström approximation

Standard kernel machines are often impossible to use in large-scale applications because of their high computational cost associated with the kernel matrix \mathbf{K} which has $O(n^2)$ storage and $O(n^2d)$ computational complexity: $\forall i, j \in \llbracket N \rrbracket, \mathbf{K}_{i,j} = k(\mathbf{x}_i, \mathbf{x}_j)$. A well-known strategy to overcome this problem is to use the Nyström method which computes a low-rank approximation of the kernel matrix on the basis of some pre-selected landmark points.

Given $K \ll n$ landmark points $\{\mathbf{U}_i\}_{i=1}^K$, the Nyström method gives the following approximation of the full kernel matrix:

$$\mathbf{K} \approx \tilde{\mathbf{K}} = \mathbf{C}\mathbf{W}^\dagger\mathbf{C}^T, \quad (9)$$

with $\mathbf{W} \in \mathbb{R}^{K \times K}$ containing all the kernel values between landmarks: $\forall i, j \in \llbracket K \rrbracket \mathbf{W}_{i,j} = k(\mathbf{U}_i, \mathbf{U}_j)$; \mathbf{W}^\dagger being the pseudo-inverse of \mathbf{W} and $\mathbf{C} \in \mathbb{R}^{n \times K}$ containing the kernel values between landmark points and all data points: $\forall i \in \llbracket n \rrbracket, \forall j \in \llbracket K \rrbracket \mathbf{C}_{i,j} = k(\mathbf{X}_i, \mathbf{U}_j)$.

4.4.2 Efficient Nyström approximation

A substantial amount of research has been conducted toward landmark point selection methods for improved approximation accuracy [7] [14], but much less has been done to improve computation speed. In [19], the authors propose an algorithm to learn the matrix of landmark points with some structure constraint, so that its utilisation is fast, taking advantage of fast-transforms. This results in an efficient Nyström approximation that is faster to use both in the training and testing phases of some ulterior machine learning application.

Remarking that the main computation cost of the Nyström approximation comes from the computation of the kernel function between the train/test samples and the landmark points, [19] aim at accelerating this step. In particular, they focus on a family of kernel functions that has the following form:

$$k(\mathbf{x}_i, \mathbf{x}_j) = f(\mathbf{x}_i)f(\mathbf{x}_j)g(\mathbf{x}_i^T \mathbf{x}_j), \tag{10}$$

where $f : \mathbb{R}^d \rightarrow \mathbb{R}$ and $g : \mathbb{R} \rightarrow \mathbb{R}$. They show that this family of functions contains some widely used kernels such as the Gaussian and the polynomial kernel. Given a set of K landmark points $\mathbf{U} \in \mathbb{R}^{K \times d}$ and a sample \mathbf{x} , the computational time for computing the kernel between \mathbf{x} and each row of \mathbf{U} (necessary for the Nyström approximation) is bottlenecked by the computation of the product $\mathbf{U}\mathbf{x}$. They hence propose to write the \mathbf{U} matrix as the concatenation of structured $s = K/d$ product of matrices:

$$\mathbf{U} = [\mathbf{V}_1 \mathbf{H}^T, \dots, \mathbf{V}_s \mathbf{H}^T]^T, \tag{11}$$

where the \mathbf{H} is a $d \times d$ matrix associated with a fast transform such as the *Haar* or *Hadamard* matrix, and the \mathbf{V}_i s are some $d \times d$ diagonal matrices to be either chosen with a standard landmark selection method or learned using an algorithm they provide.

Depending on the \mathbf{H} matrix chosen, it is possible to improve the time complexity for the computation of $\mathbf{U}\mathbf{x}$ from $O(Kd)$ to $O(K \log d)$ (*Fast Hadamard transform*) or $O(K)$ (*Fast Haar Transform*).

4.4.3 QK-means in Nyström

We propose to use our QK-means algorithm in order to learn directly the \mathbf{U} matrix in the Nyström approximation so that the matrix-vector multiplication $\mathbf{U}\mathbf{x}$ is cheap to compute, but the structure of \mathbf{U} is not constrained by some pre-defined transform matrix. We propose to take the objective \mathbf{U} matrix as the K-means matrix of \mathbf{X} since it has been shown to achieve good reconstruction accuracy in the Nyström method [7].

As shown in the next sub-section, our algorithm allow to obtain an efficient Nyström approximation, while not reducing too much the quality of the K-means landmark points which are encoded as a factorization of sparse matrix.

4.4.4 Results

The Figure 5 summarizes the results achieved in the Nyström approximation setting.

The Figures on the right display the average time for computing one line of the approximated matrix in Equation 9. In Figure 5b, we clearly see the speed-up offered using the QK-means method on the Blobs dataset. On the Mnist and Fashion-MNIST dataset (Figure 5d and 5f), this speed-up is sensible but not as clear because the standard deviation is much larger.

The Figures on the left show the approximation error of the Nyström approximation based on different sampling schemes w.r.t. the real kernel matrix. This error is computed by the Froebenius norm of the difference between the matrices and then normalized:

$$error = \frac{\|\mathbf{K} - \tilde{\mathbf{K}}\|_F}{\|\mathbf{K}\|_F} \tag{12}$$

. The QK-means approach gives better reconstruction error than the Nyström method based on uniform sampling although they are slightly worse than the one obtained with the K-means centroids. We see that that

the difference in approximation error between **K-means** and **QK-means** is almost negligible when compared to the approximation error obtained with the uniform sampling scheme.


From a more practical point of view, we show in Table 3 and Table 4 that the Nyström approximation based on **QK-means** can then be used in a linear SVM and achieve as good performance as the one based on the **K-means** approach.

5 Conclusion


In this paper, we have proposed a variant of the **K-means** algorithm, named **QK-means**, designed to achieve a similar goal – clustering data points around K learned centroids – with a much lower computational complexity as the dimension of the data, the number of examples and the number of clusters get high. Our approach is based on the approximation of the centroid matrix by an operator structured as a product of a small number of sparse matrices, resulting in a low time and space complexity when applied to data vectors. We have shown the convergence properties of the proposed algorithm and provided its complexity analysis.

An implementation prototype has been run in several core machine learning use cases including clustering, nearest-neighbor search and Nyström approximation. The experimental results illustrate the computational gain in high dimension at inference time as well as the good approximation qualities of the proposed model.

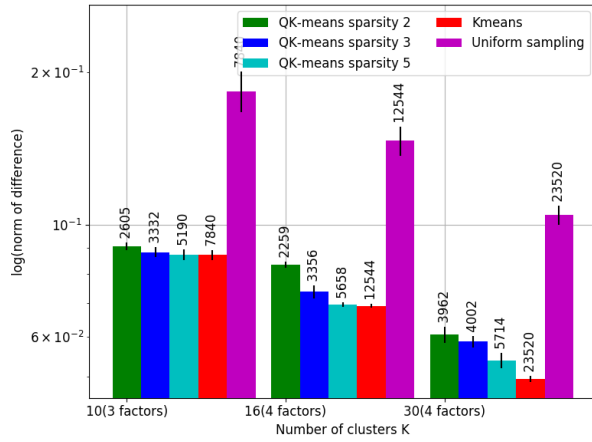
Beyond these modeling, algorithmic and experimental contributions to low-complexity high-dimensional machine learning, we have identified several important questions that are still to be addressed. First, although learning the fast-structure operator has been nicely integrated in the training algorithm with an advantageous theoretical time and space complexity, exhibiting gains in actual running times has not been achieved yet for the **QK-means** learning procedure, compared to **K-means**. This may be obtained in even higher dimensions than in the proposed experimental settings, which may require a new version of **QK-means** using batches of data in order to process amounts of data that do not fit in memory. Second, the expressiveness of the fast-structure model is still to be theoretically studied, while our experiments seems to show that arbitrary matrices may be well fitted by such models. Third, we believe that learning fast-structure linear operators during the training procedure may be generalized to many core machine learning methods in order to speed them up and make them scale to larger dimensions.




(a) Blobs: Nyström reconstruction error.




(b) Blobs: Nyström inference time.




(c) MNIST: Nyström reconstruction error.



(d) MNIST: Nyström inference time.



(e) Fashion-MNIST: Nyström reconstruction error.



(f) Fashion-MNIST: Nyström inference time.

Figure 5: Nyström approximation results: accuracy (left) and running times (right). The uniform sampling based Nyström approximation running times are not displayed because they are the same as for the Nyström approximation based on K-means centroids. Every experiment results are averaged over 5 runs. The vertical black lines are the standard deviation w.r.t. the runs.

References

- [1] David Arthur and Sergei Vassilvitskii. How slow is the k-means method? In *Symposium on Computational Geometry*, pages 1–10, 2006.
- [2] Jérôme Bolte, Shoham Sabach, and Marc Teboulle. Proximal alternating linearized minimization or nonconvex and nonsmooth problems. *Mathematical Programming*, 146(1-2):459–494, 2014.
- [3] Christos Boutsidis, Anastasios Zouzias, Michael W Mahoney, and Petros Drineas. Randomized dimensionality reduction for k -means clustering. *IEEE Transactions on Information Theory*, 61(2):1045–1062, 2014.
- [4] Alex Gittens and Michael W Mahoney. Revisiting the nystrom method for improved large-scale machine learning. *The Journal of Machine Learning Research*, 17(1):3977–4041, 2016.
- [5] John A Hartigan and Manchek A Wong. Algorithm as 136: A k -means clustering algorithm. *Journal of the Royal Statistical Society. Series C (Applied Statistics)*, 28(1):100–108, 1979.
- [6] Anil K Jain. Data clustering: 50 years beyond k -means. *Pattern recognition letters*, 31(8):651–666, 2010.
- [7] Sanjiv Kumar, Mehryar Mohri, and Ameet Talwalkar. Sampling methods for the nystrom method. *Journal of Machine Learning Research*, 13(Apr):981–1006, 2012.
- [8] Quoc Le, Tamás Sarlós, and Alex Smola. Fastfood—approximating kernel expansions in loglinear time. In *International Conference on Machine Learning*, 2013.
- [9] Luc Le Magoarou and Rémi Gribonval. Flexible multilayer sparse approximations of matrices and applications. *IEEE Journal of Selected Topics in Signal Processing*, 10(4):688–700, 2016.
- [10] Yann LeCun and Corinna Cortes. MNIST handwritten digit database, 2010.
- [11] Weiwei Liu, Xiaobo Shen, and Ivor Tsang. Sparse embedded k -means clustering. In *Advances in Neural Information Processing Systems*, pages 3319–3327, 2017.
- [12] Jacques Morgenstern. The Linear Complexity of Computation. *Journal of the ACM*, 22(2):184–194, April 1975.
- [13] Marius Muja and David G Lowe. Scalable nearest neighbor algorithms for high dimensional data. *IEEE transactions on pattern analysis and machine intelligence*, 36(11):2227–2240, 2014.
- [14] Cameron Musco and Christopher Musco. Recursive sampling for the nystrom method. In *Advances in Neural Information Processing Systems*, pages 3833–3845, 2017.
- [15] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and E. Duchesnay. Scikit-learn: Machine learning in Python. *Journal of Machine Learning Research*, 12:2825–2830, 2011.
- [16] Qichao Que and Mikhail Belkin. Back to the future: Radial basis function networks revisited. In *AISTATS*, pages 1375–1383, 2016.
- [17] David Sculley. Web-scale k -means clustering. In *Proceedings of the 19th international conference on World wide web*, pages 1177–1178. ACM, 2010.
- [18] Xiaobo Shen, Weiwei Liu, Ivor Tsang, Fumin Shen, and Quan-Sen Sun. Compressed k -means for large-scale clustering. In *Thirty-First AAAI Conference on Artificial Intelligence*, 2017.
- [19] Si Si, Cho-Jui Hsieh, and Inderjit Dhillon. Computationally efficient nystrom approximation using fast transforms. In *International Conference on Machine Learning*, pages 2655–2663, 2016.

- [20] Twan Van Laarhoven and Elena Marchiori. Local network community detection with continuous optimization of conductance and weighted kernel k-means. *The Journal of Machine Learning Research*, 17(1):5148–5175, 2016.

A palm4MSA algorithm

The palm4MSA algorithm [9] is given in Algorithm 2 together with the time complexity of each line, using $A = \min(K, D)$ and $B = \max(K, D)$. Even more general constraints can be used, the constraint sets \mathcal{E}_q are typically defined as the intersection of the set of unit Frobenius-norm matrices and of a set of sparse matrices. The unit Frobenius norm is used together with the λ factor to avoid a scaling indeterminacy. Note that to simplify the model presentation, factor λ is used internally in palm4MSA and is integrated in factor \mathbf{S}_1 at the end of the algorithm (Line 14) so that \mathbf{S}_1 does not satisfy the unit Frobenius norm in \mathcal{E}_1 at the end of the algorithm. The sparse constraints we used, as in [9], consist of trying to have a given number of non-zero coefficients in each row and in each column. This number of non-zero coefficients is called sparsity level in this paper. In practice, the projection function at Line 9 keeps the largest non-zero coefficients in each row and in each column, which only guarantees the actual number of non-zero coefficients is at least equal to the sparsity level.

Algorithm 2 palm4MSA algorithm

Require: The matrix to factorize $\mathbf{U} \in \mathbb{R}^{K \times D}$, the desired number of factors Q , the constraint sets \mathcal{E}_q , $q \in \llbracket Q \rrbracket$ and a stopping criterion (e.g., here, a number of iterations I).

```

1:  $\lambda \leftarrow \|\mathbf{S}_1\|_F$   $\mathcal{O}(B)$ 
2:  $\mathbf{S}_1 \leftarrow \frac{1}{\lambda} \mathbf{S}_1$   $\mathcal{O}(B)$ 
3: for  $i \in \llbracket I \rrbracket$  while the stopping criterion is not met do
4:   for  $q = Q$  down to 1 do
5:      $\mathbf{L}_q \leftarrow \prod_{l=1}^{q-1} \mathbf{S}_l^{(i)}$ 
6:      $\mathbf{R}_q \leftarrow \prod_{l=q+1}^Q \mathbf{S}_l^{(i+1)}$ 
7:     Choose  $c > \lambda^2 \|\mathbf{R}_q\|_2^2 \|\mathbf{L}_q\|_2^2$   $\mathcal{O}(A \log A + B)$ 
8:      $\mathbf{D} \leftarrow \mathbf{S}_q^i - \frac{1}{c} \lambda \mathbf{L}_q^T (\lambda \mathbf{L}_q \mathbf{S}_q^i \mathbf{R}_q - \mathbf{U}) \mathbf{R}_q^T$   $\mathcal{O}(AB \log A)$ 
9:      $\mathbf{S}_q^{(i+1)} \leftarrow P_{\mathcal{E}_q}(\mathbf{D})$   $\mathcal{O}(A^2 \log A)$  or  $\mathcal{O}(AB \log B)$ 
10:   end for
11:    $\hat{\mathbf{U}} := \prod_{j=1}^Q \mathbf{S}_j^{(i+1)}$   $\mathcal{O}(A^2 \log A + AB)$ 
12:    $\lambda \leftarrow \frac{\text{Trace}(\mathbf{U}^T \hat{\mathbf{U}})}{\text{Trace}(\hat{\mathbf{U}}^T \hat{\mathbf{U}})}$   $\mathcal{O}(AB)$ 
13: end for
14:  $\mathbf{S}_1 \leftarrow \lambda \mathbf{S}_1$   $\mathcal{O}(B)$ 

```

Ensure: $\{\mathbf{S}_q : \mathbf{S}_q \in \mathcal{E}_q\}_{q \in \llbracket Q \rrbracket}$ such that $\prod_{q \in \llbracket Q \rrbracket} \mathbf{S}_q \approx \mathbf{U}$

The complexity analysis is proposed under the following assumptions, which are satisfied in the mentioned applications and experiments: the number of factors is $Q = \mathcal{O}(\log A)$; all but one sparse factors are of shape $A \times A$ and have $\mathcal{O}(A)$ non-zero entries while one of them is of shape $A \times B$ or $B \times A$ with $\mathcal{O}(B)$ non-zero entries. In such conditions, the complexity of each line is:

Lines 1-2 Computing these normalization steps is linear in the number of non-zeros coefficients in \mathbf{S}_1 .

Lines 5-6 Fast operators \mathbf{L} and \mathbf{R} are defined for subsequent use without computing explicitly the product.

Line 7 The spectral norm of \mathbf{L} and \mathbf{R} is obtained via a power method by iteratively applying each operator, benefiting from the fast transform.

Line 8 The cost of the gradient step is dominated by the product of sparse matrices.

- Line 9 The projection onto a sparse-constraint set takes $\mathcal{O}(A^2 \log A)$ for all the $A \times A$ matrices and $\mathcal{O}(AB \log B)$ for the rectangular matrix at the leftmost or the rightmost position.
- Line 11 The reconstructed matrix $\hat{\mathbf{U}}$ is computed using $\mathcal{O}(\log A)$ products between $A \times A$ sparse matrices, in $\mathcal{O}(A^2)$ operations each, and one product with a sparse matrix in $\mathcal{O}(AB)$.
- Line 12 The numerator and denominator can be computed using a Hadamard product between the matrices followed by a sum over all the entries.
- Line 14 Computing renormalization step is linear in the number of non-zeros coefficients in \mathbf{S}_1 .

Hence, the overall time complexity of `palms4MSA` is in $\mathcal{O}(AB(\log^2 A + \log B))$, due to Lines 8 and 9.