

Mini-Batch Spectral Clustering

Yufei HAN

Symantec Research Labs
Biot Sophia-Antipolis, France
Email: yufei.han@symantec.com

Maurizio Filippone

Department of Data Science, EURECOM
Biot Sophia-Antipolis, France
Email: maurizio.filippone@eurecom.fr

Abstract—The cost of computing the spectrum of Laplacian matrices hinders the application of spectral clustering to large data sets. While approximations recover computational tractability, they can potentially affect clustering performance. This paper proposes a practical approach to learn spectral clustering, where the spectrum of the Laplacian is recovered following a constrained optimization problem that we solve using adaptive mini-batch-based stochastic gradient optimization on Stiefel manifolds. Crucially, the proposed approach is formulated so that the memory footprint of the algorithm is low, the cost of each iteration is linear in the number of samples, and convergence to critical points of the objective function is guaranteed. Extensive experimental validation on data sets with up to half a million samples demonstrate its scalability and its ability to outperform state-of-the-art approximate methods to learn spectral clustering for a given computational budget.

I. INTRODUCTION

Over the past two decades, spectral clustering has established itself as one of the most prominent clustering methods [1], [2]. The effectiveness of spectral clustering in identifying complex structures in data is a direct consequence of its close connection with kernel machines [3], [4]. Because of this connection, however, it is also apparent that spectral clustering inherits the scalability issues of kernel machines. In spectral clustering, the computational challenge is to determine the spectral properties of the so called Laplacian matrix [5]. Denoting by n the number of samples, storing the Laplacian matrix requires $\mathcal{O}(n^2)$ space while calculating its spectrum requires $\mathcal{O}(n^3)$ computations.

Several approaches have been proposed to reduce the complexity of spectral clustering, such as employing power methods to identify the principal eigenvectors of the Laplacian [6]. While this approach is exact in the limit of the iterations and does not require storing the Laplacian, the complexity is dominated by the iterative multiplication of the Laplacian matrix by vectors, leading to $\mathcal{O}(n^2)$ computations. In order to further reduce this complexity to $\mathcal{O}(n)$, a number of approximations are proposed in the literature. A popular technique based on the Nyström approximation relies on a small set of inducing points to approximate the spectrum of the Laplacian matrix [7]. Other recent approximations which attempt to compress the dataset appear in [8] and [9]. These approximations recover tractability and make it possible to apply spectral clustering to large data sets. However, approximations can generally affect the quality of the clustering solution, as we illustrate in the experiments.

This paper proposes a practical approach to learn spectral clustering as follows. Denoting by L the Laplacian matrix, the

idea hinges on the possibility to cast the algebraic problem of identifying its principal eigenvectors as the trace optimization problem

$$\begin{aligned} \arg \min_{W \in \mathcal{R}^{n \times k}} \left\{ \text{Tr} \left(-\frac{1}{2} W^\top L W \right) \right\}, \\ \text{subject to } W^\top W = I. \end{aligned} \quad (1)$$

We propose to solve the constrained optimization problem by means of stochastic gradient optimization. In view of the orthonormality constraint the elements of W lie on the so called Stiefel manifold, and appealing to theoretical guarantees of convergence of stochastic gradient optimization on manifolds, we can prove that our proposal converges to a critical point of the objective function in the limit of the iterations [10]. In order to simplify the tuning of the optimization procedure, we adapt Adagrad [11] for stochastic gradient optimization on the Stiefel manifold. The novelty of our proposal stems from the use of stochastic linear algebra techniques to compute stochastic gradients in $\mathcal{O}(n)$. This leads to computations of stochastic gradients that require processing of a limited number of columns of the full Laplacian matrix, motivating us to name our proposal Mini-Batch Spectral Clustering.

The results on a variety of clustering problems with n up to $580K$ give credence to the value of our proposal. We can tackle large-scale spectral clustering problems achieving the same level of accuracy of the approach that uses the exact spectrum of L at a fraction of the computing time. We also compare against approximate spectral clustering methods and show that approximations lead to faster solutions that are suboptimal compared to what we can achieve with the proposed method, especially for large data sets.

Summary of contributions (i) We propose adaptive stochastic gradient optimization to solve the constrained optimization problem in Eq. 1. (ii) We present a novel way of computing stochastic gradients linearly in the number of data that does not require storing the Laplacian matrix. (iii) We analyze the variance of the proposed estimator of the exact gradient to explain the impact of algorithm parameters. (iv) We make use of Nyström approximations to speedup the convergence of the proposed algorithm. (v) We demonstrate that our proposal allows us to tackle large-scale spectral clustering problems by reporting results on data sets of size up to $n = 580K$. Crucially, we can achieve clustering solutions of similar accuracy and orders of magnitude faster compared to the approach that computes the exact spectrum of L , and higher clustering

accuracy compared to approximate methods at a comparable cost.

II. SPECTRAL CLUSTERING

A. Background

Consider an unsupervised learning problem where $X = \{\mathbf{x}_1, \dots, \mathbf{x}_n\}$ is a set of n samples to be clustered. The formulation of spectral clustering introduces an undirected graph \mathcal{G} based on X , where the n nodes of \mathcal{G} represent the n input data in X , and the edges are weighted according to a similarity measure between the inputs. The graph \mathcal{G} is expressed by an $n \times n$ adjacency matrix A , where each entry a_{ij} determines the weight associated with the edge connecting inputs i and j . Typically, the elements of A are defined through off-the-shelf kernel functions, e.g., the Radial Basis Function (RBF) kernel [12].

Spectral clustering attempts to cluster the elements of X by analyzing the graph \mathcal{G} . In particular, the objective of spectral clustering is to partition the graph so as to minimize some graph cut criterion, e.g., the normalized cut [1]. The graph cut problem is generally NP-hard, but its relaxation leads to the definition of the clustering problem as the solution of an algebraic problem [13]. In particular, following the spectral clustering algorithm proposed by [2], the graph Laplacian is defined as a normalized version of the adjacency matrix

$$L = D^{-\frac{1}{2}} A D^{-\frac{1}{2}}, \quad (2)$$

where D is the diagonal matrix of the degrees of the n nodes. Spectral clustering represents each data point using the corresponding component of the top k eigenvectors of the Laplacian L , and computes the solution to the clustering problem by applying k -means in this representation. The difficulty in solving the graph cut problem then becomes calculating the spectrum of the Laplacian matrix L ; this requires $\mathcal{O}(n^3)$ computations and $\mathcal{O}(n^2)$ space, making it prohibitive – if not unfeasible – for large data sets. The aim of this paper is to address this scaling issue of spectral clustering without sacrificing the accuracy of the solution, as explained next.

B. Constrained optimization problem formulation

The first step to reduce the complexity in finding the top k eigenvectors of L , is to cast this algebraic operation as solving the constrained optimization problem in Eq. 1. This is an optimization problem involving $n \times k$ parameters representing the n components of the top k eigenvectors of L . The objective function rewards maximization of a score that, at convergence, is the sum of the k largest eigenvalues. The constraint $W^\top W = I$ gives rise to the so called Stiefel manifold in $\mathcal{R}^{n \times d}$; this imposes orthonormality on the columns of W which, at convergence, represent the eigenvectors associated with the k largest eigenvalues.

The constrained optimization problem in Eq. 1 can be solved by formulating standard optimization algorithms to deal with the Riemannian metric on the manifold [14], which typically rely on search operations along geodesics. Alternative schemes, such as the Cayley transform, have been proposed to tackle




Fig. 1: Left: Exact gradient G and a few realizations of stochastic gradients \tilde{G} ; the projection of these gradients onto the tangent space $T_W M$ of a manifold M at a given W give H and \tilde{H} , respectively. Right: Retraction scheme that approximates the exponential map at W .

optimization problems on Riemannian manifolds [15]. All these optimization schemes require calculating the gradient of the objective function

$$G = \nabla_W \text{Tr} \left(-\frac{1}{2} W^\top L W \right) = -LW. \quad (3)$$

This costs $\mathcal{O}(n^2)$ and does not require storing L anymore, which is an improvement with respect to computing the full spectrum of L . However, while casting spectral clustering as a constrained optimization problem improves scalability, when n is large it may still take a prohibitive amount of time to be practical. In the next section we present our proposal to solve the constrained optimization problem using iterative updates that cost $\mathcal{O}(n)$ operations, with the guarantee of convergence to critical points of the objective function in Eq. 1.

III. MINI-BATCH SPECTRAL CLUSTERING (MBSC)

The intuition behind our proposal is that it is possible to solve the constrained optimization problem in Eq. 1 relying exclusively on stochastic gradients [10]. By inspecting the expression of the exact gradient in Eq. 3, we show how stochastic linear algebra techniques can be employed to unbiasedly estimate exact gradients ($\mathcal{O}(n^2)$ cost) with stochastic gradients at $\mathcal{O}(n)$ cost.

A. Stochastic optimization on Stiefel manifolds

Stochastic gradient optimization on manifolds is based on the notion of Riemannian gradients, which are elements of the tangent space at a given W that determine the direction of steepest increase of the objective function on the manifold. For the Stiefel manifold, the Riemannian gradient is

$$H = (I - WW^\top)G. \quad (4)$$

In the case where an unbiased version \tilde{G} of the gradient G is available, namely $E[\tilde{G}] = G$, an intuitively sensible strategy is to perform stochastic gradient optimization on the manifold. Formally, this would amount in moving along geodesics for a given length based on the noisy version of the gradient. Defining γ_t as the equivalent of the usual step-sizes in stochastic gradient optimization, and $\exp_W(\cdot)$ as the exponential map at W , the

Algorithm 1: Stochastic Riemannian gradient on Stiefel Manifold using Mini-Batches

Ht i l d e (L, p, N_r, W)
 Initialize $\tilde{G} \in R^{n \times k}$ with elements equal to zero
for $i = 1$ **to** N_r **do**
 | Draw the components of \mathbf{r}_i
 | $\tilde{G} += \frac{1}{N_r} L \mathbf{r}_i \mathbf{r}_i^\top W$
end
 Return $\tilde{H} = (I - WW^\top) \tilde{G}$

update equation would then be $W' = \exp_W(-\gamma_t \tilde{H})$, where $\tilde{H} = (I - WW^\top) \tilde{G}$ is the unbiased Riemannian gradient of the objective function. This approach can be shown to converge to critical points of the objective function in Eq. 1 [10]. However, computing the exponential map to simulate the trajectory of the solver on the manifold requires solving potentially expensive differential equations.

An alternative that avoids computing the exponential map altogether is to replace this calculation with an approximation $W' = R_W(-\gamma_t \tilde{H})$ that is much easier to calculate. If the so called retraction R_W satisfies the property that $d(R_W(\delta v), \exp_W(\delta v)) = \mathcal{O}(\delta^2)$, where v is an element of the tangent space, then it is still possible to prove convergence [10]. A simple and computationally convenient retraction function that satisfies this property is

$$W' = \text{QR}_Q \left(W - \gamma_t (I - WW^\top) \tilde{G} \right), \quad (5)$$

where QR_Q extracts the orthonormal factor Q of a QR decomposition. This simple retraction moves the optimization in the direction of the stochastic Riemannian gradient $\tilde{H} = (I - WW^\top) \tilde{G}$ and applies an orthonormalization step to ensure that the update is projected back onto the manifold. Under this choice of retraction, we can appeal to the theoretical results in [10] that ensure convergence to critical points of the objective function in the limit of the iterations, similarly to standard stochastic gradient optimization. An illustration of the retraction scheme is provided in Fig.1.

B. Calculation of stochastic gradients in $\mathcal{O}(n)$

The introduction of stochasticity in the calculation of \tilde{G} follows on from ideas that have been proposed to calculate unbiased stochastic approximations to algebraic quantities, such as traces and log-determinants [16], [17]. In particular, we define a vector \mathbf{r} such that $\text{E}[\mathbf{r} \mathbf{r}^\top] = I$ and we rewrite the expensive matrix product as

$$G = LW = LIW = LE[\mathbf{r} \mathbf{r}^\top]W = \text{E}[L \mathbf{r} \mathbf{r}^\top W], \quad (6)$$

which suggests that we can replace the exact calculation of LW with the estimator

$$\tilde{G} = \frac{1}{N_r} \sum_{i=1}^{N_r} L \mathbf{r}_i \mathbf{r}_i^\top W. \quad (7)$$

The key to making computations linear in n is to define the components of the random vectors \mathbf{r}_i as drawn from

Algorithm 2: Mini-Batch Spectral Clustering

Input: Normalized Laplacian Matrix $L \in R^{n \times n}$, number of clusters k , regularization factor ε

Parameters: Step-size λ , max # iterations T

Output: Cluster labels of each data points

Initialize $W^{(0)} \in R^{n \times k}$ as a random orthonormal matrix

Initialize $M^{(0)} \in R^{n \times k}$ with elements equal to zero

for $t = 1$ **to** T **do**

 | $\tilde{H}^{(t)} = \text{Ht i l d e}(L, p, N_r, W^{(t-1)})$

 | $M_{ij}^{(t)} = M_{ij}^{(t-1)} + |\tilde{H}_{ij}^{(t)}|^2$

 | $\hat{H}_{ij}^{(t)} = \frac{\tilde{H}_{ij}^{(t)}}{\varepsilon + \sqrt{M_{ij}^{(t)}}}$

 | $W^{(t)} = W^{(t-1)} - \lambda \hat{H}^{(t)}$

 | $W^{(t)} = \text{QR}_Q(W^{(t)})$

end

Apply k -means on $W^{(T)}$ to get the cluster labels

the set $\{-p^{-\frac{1}{2}}, 0, +p^{-\frac{1}{2}}\}$ with probabilities $(p/2, 1-p, p/2)$ respectively. It is straightforward to verify that $\text{E}[\mathbf{r}_i \mathbf{r}_i^\top] = I$, and p can be chosen to enforce any proportion of zeros in the \mathbf{r}_i vectors. With this mechanism to inject stochasticity in the calculation of the gradients, we are effectively ignoring some columns of the matrix L whenever there is a zero in the corresponding positions of the \mathbf{r}_i vectors. This makes it possible to update the parameters W during the solution of the constrained optimization problem in Eq. 1, by only selecting a few columns of the full Laplacian. If the average number of non-zero elements is chosen to be independent of n , the calculation of the stochastic gradient is $\mathcal{O}(n)$, making the proposed iterative solver linear in the number of samples. The memory footprint of the algorithm is a distinctive feature of our proposal; if the degree matrix D is precomputed, calculating the necessary columns of L requires evaluating and normalizing just $\mathcal{O}(n)$ elements of the adjacency matrix A .

Given that only a subset of the columns in L are used at each iteration to calculate stochastic gradients, we term our proposal Mini-Batch Spectral Clustering (MBSC). Instead of defining a probability p to select columns and to repeat this N_r times, we can fix the number and indices of columns that are selected at each iteration (size of the mini-batch) to be $m = lN_r$ and interpret $p = l/n$. While this is intuitively sensible, fixing the indices of the mini-batches would violate the property that $\text{E}[\mathbf{r}_i \mathbf{r}_i^\top] = I$. One easy way around this issue is to constantly change the way data are split into mini-batches, e.g., by shuffling the data, and this recovers the property $\text{E}[\mathbf{r}_i \mathbf{r}_i^\top] = I$. Even though it is not the focus of the current paper, we envisage the possibility to develop a distributed version of the proposed MBSC algorithm based on our formulation. The proposed MBSC algorithm is sketched in Algorithms 1 and 2, where, for the sake of clarity, L is assumed to be stored. For memory constrained systems where storing the whole Laplacian matrix is unfeasible, our proposal can easily be adapted to avoid storing it, and we report on the performance of this

variant in the experiments.

C. Variance of stochastic gradients

Here we are interested in quantifying the impact of the choice of p and N_r in the variance of the proposed estimator of the exact gradients. Without loss of generality, we focus on the variance of a given column of the stochastic gradient, say $\tilde{G}_{\cdot s} = L\mathbf{r}\mathbf{r}^\top W_{\cdot s}$ that is the one associated with the s th eigenvector. In order to avoid cluttering the presentation of the paper, we leave the detailed derivation of the variance and its bound to the next section and we only report the final expression here. We express the bound in terms of the mini-batch size $m = lN_r$, for which $p = l/n$, and when N_r vectors \mathbf{r}_i are used to calculate stochastic gradients as in Eq. 7.

$$\begin{aligned} \text{Tr}[\text{cov}(L\mathbf{r}\mathbf{r}^\top \mathbf{w})] &\leq \frac{1}{N_r} G_{\cdot s}^\top G_{\cdot s} + \frac{1}{N_r} \text{Tr}(LL^\top) \\ &\quad + \frac{n}{lN_r} \text{Tr}[L\text{diag}(\mathbf{w})\text{diag}(\mathbf{w})L^\top]. \end{aligned}$$

This reveals that we have two ways of reducing the variance of stochastic gradients; one is to increase N_r and another is to increase l . Imagine that we fix the mini-batch size $m = lN_r$; is it better to increase l and reduce N_r , or the other way around? For the last term it does not matter. For the first two instead, given that they depend only on N_r , it is clear that we should favor increasing N_r and reducing l . This entails that we should consider averaging stochastic gradients over several subsets of a mini-batch instead of a few large ones. This result is interesting because in other popular mini-batch approaches increasing the mini-batch size or the number of repetitions is equivalent. In our proposed MBSC, because of the nonlinearity of the estimator with respect to the vectors \mathbf{r}_i , the bound on the variance shows an unintuitive asymmetry between N_r and l . A further consideration we can make is that this suggestion is most useful during the first phase of the optimization where the gradients are likely to be large.

D. Detailed derivation of the variance of stochastic gradients

We report here a detailed derivation of the variance of a given column of the stochastic gradient \tilde{G} , namely the one associated with a given eigenvector, say $\mathbf{w} := W_{\cdot s}$. Recall that the exact gradient with respect to \mathbf{w} would be $G_{\cdot s} = L\mathbf{w}$ and assume that we use a single vector \mathbf{r} to unbiasedly estimate this as $\tilde{G}_{\cdot s} = L\mathbf{r}\mathbf{r}^\top \mathbf{w}$. The covariance of $\tilde{G}_{\cdot s}$ is

$$\text{E}[(L\mathbf{r}\mathbf{r}^\top \mathbf{w})(L\mathbf{r}\mathbf{r}^\top \mathbf{w})^\top] - (\text{E}[L\mathbf{r}\mathbf{r}^\top \mathbf{w}]) (\text{E}[L\mathbf{r}\mathbf{r}^\top \mathbf{w}])^\top.$$

We are interested in analyzing the sum of the variances of the components of the stochastic gradient:

$$\begin{aligned} \text{Tr}[\text{cov}(L\mathbf{r}\mathbf{r}^\top \mathbf{w})] &= \text{E}[(L\mathbf{r}\mathbf{r}^\top \mathbf{w})^\top (L\mathbf{r}\mathbf{r}^\top \mathbf{w})] \\ &\quad - (\text{E}[L\mathbf{r}\mathbf{r}^\top \mathbf{w}])^\top (\text{E}[L\mathbf{r}\mathbf{r}^\top \mathbf{w}]). \end{aligned} \quad (8)$$

By expanding the first term and realizing that the second term in the right hand side is just the norm of the s th column of \tilde{G} , we obtain that the sum variances of the components of the stochastic gradient is:

$$\text{Tr}[\text{cov}(L\mathbf{r}\mathbf{r}^\top \mathbf{w})] = \text{Tr}[LE[\mathbf{r}\mathbf{r}^\top \mathbf{w}\mathbf{w}^\top \mathbf{r}\mathbf{r}^\top]L^\top] - \mathbf{w}^\top L^\top L\mathbf{w}.$$

The expectation on the right hand side can be computed analytically; its derivation is rather lengthy and, for the sake of clarity, we just report the final expression here:

$$\text{E}(\mathbf{r}\mathbf{r}^\top \mathbf{w}\mathbf{w}^\top \mathbf{r}\mathbf{r}^\top) = 2\mathbf{w}\mathbf{w}^\top + \left(\frac{1}{p} - 3\right) \text{diag}(\mathbf{w})\text{diag}(\mathbf{w}) + I.$$

Plugging the expression of the expectation into the expression of the variance of the components of the stochastic gradient, we obtain the following expression:

$$\begin{aligned} \text{Tr}[\text{cov}(L\mathbf{r}\mathbf{r}^\top \mathbf{w})] &= G_{\cdot s}^\top G_{\cdot s} + \text{Tr}(LL^\top) \\ &\quad + \left(\frac{1}{p} - 3\right) \text{Tr}[L\text{diag}(\mathbf{w})\text{diag}(\mathbf{w})L^\top], \end{aligned}$$

which we can bound as follows

$$\begin{aligned} \text{Tr}[\text{cov}(L\mathbf{r}\mathbf{r}^\top \mathbf{w})] &\leq G_{\cdot s}^\top G_{\cdot s} + \text{Tr}(LL^\top) \\ &\quad + \frac{1}{p} \text{Tr}[L\text{diag}(\mathbf{w})\text{diag}(\mathbf{w})L^\top]. \end{aligned}$$

E. Other practical considerations

In the following, we employ a “warm start” initialization for the proposed MBSC method by carrying out a Nyström approximation and setting the initial $W^{(0)}$ to the top k approximate eigenvectors [7]. We will illustrate the impact of the warm start initialization in the experiments section.

Another practical consideration is that, for memory constrained systems where L is not stored, there is a difficulty in computing the degree matrix D , which is required to normalize the Laplacian (Eq. 2). We obtain an initial estimate \hat{D} through the Nyström approximation, and along the iterations, \hat{D} is updated whenever columns of L are computed. After one pass through the whole dataset, $\hat{D} = D$.

F. Other valid mini-batch formulations

We note here that there are other ways to formulate a mini-batch strategy for the optimization problem in Eq. 1. In particular, consider the following estimator

$$\tilde{G} = (L \circ H)W, \quad (9)$$

where \circ denoted the Hadamard (element-wise) product, and H is a matrix such that $\text{E}[H] = \mathbf{1}\mathbf{1}^\top$. One way to define the elements of H is to pick a set of columns with probability p and set their elements to $1/p$ and the rest to zero. In this way, H acts as a “mask” matrix that selects only the columns of the Laplacian matrix pertaining to a subset of the data. We found empirically that the variance of the resulting stochastic gradient is of the same order as the one that we propose, and that none of the two approaches gives a consistently lower variance compared to the other. A similar estimator was proposed in [18], where the elements of H to set to $1/p$ were not selected by columns, but just randomly from the whole matrix; again the variance of the stochastic gradients is comparable to the one we propose.

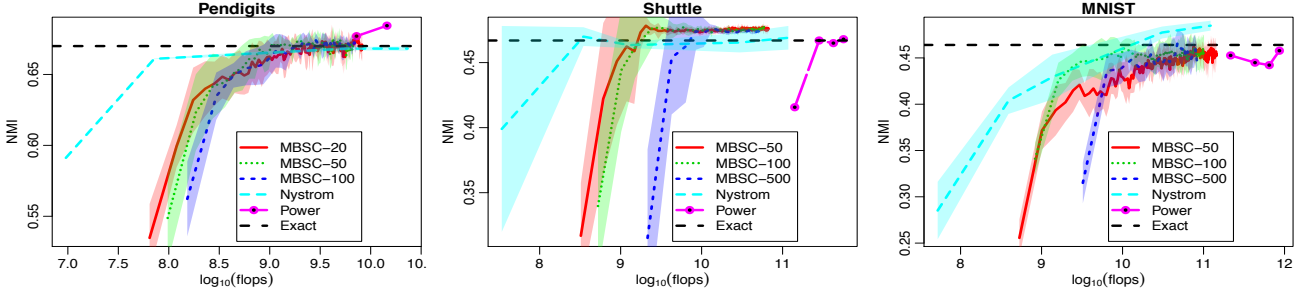


Fig. 2: NMI versus counts of floating-point operations on the `Pendigits`, `Shuttle` and `MNIST` data sets. NMI for the exact spectral clustering is shown as a constant dashed line. MBSC with the largest mini-batches runs for 3.8s, 107s and 112s to achieve an average NMI score of 0.67, 0.48 and 0.47 on the three data sets, whereas the power method takes 20s, 935s and 987s to get a stable clustering output. The Nyström approximation needs 1000 samples to obtain the optimal clustering accuracy, which takes 15.2s, 198s and 211s on the three data sets.

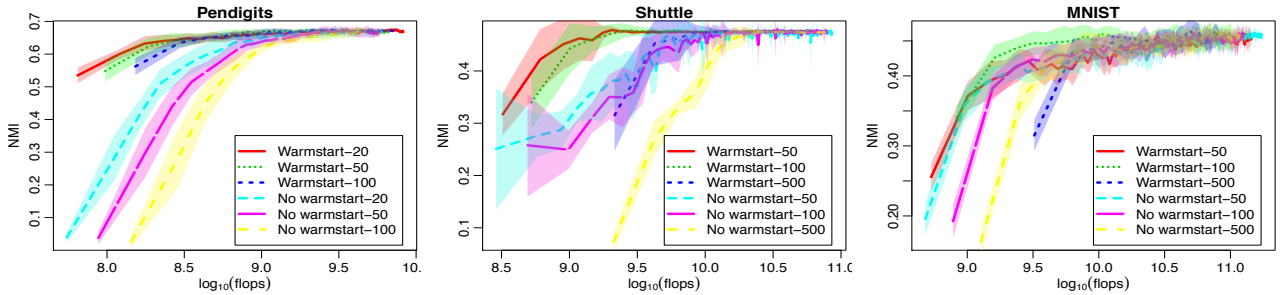


Fig. 3: NMI versus counts of floating-point operations on the `Pendigits`, `Shuttle` and `MNIST` data sets. We compare performance of MBSC with Nyström approximation based warm start (noted as "warmstart" in the figure) to that with random initialization (noted as "no warm start" in the figure)

TABLE I: Summary of the characteristics of the data sets considered in the experiments.

Data set	# samples	# feat.	# classes	σ
<code>Pendigits</code>	10992	16	10	223.61
<code>Shuttle</code>	58000	9	7	0.45
<code>MNIST</code>	60000	780	10	4.08
<code>Covtype-I</code>	100000	54	5	1.15
<code>Covtype-II</code>	581012	54	7	1.15

IV. EXPERIMENTS

Throughout the experiments, we make use of the Radial Basis Function (RBF) adjacency function:

$$a_{ij} = \exp\left(-\frac{\|\mathbf{x}_i - \mathbf{x}_j\|}{\sigma}\right) \quad i \neq j,$$

where we set $a_{ij} = 0$ when $i = j$. The parameter σ determines the rate of decay of the adjacency function, which can be set, e.g., using statistics on the distances between pairs of points [19].

We assess clustering performance using the normalized mutual information (NMI) score between the cluster labels and the ground truth class labels. To reliably measure computational cost of all methods involved in the comparison, we count the amount of floating-point addition and multiplication operations

they require, given the affinity matrix, and we also report running time statistics. We implemented all the algorithms in Python using the `numpy` and `scikit-learn` packages. All our experiments are conducted under Ubuntu Linux 14.04 with 10-core CPU and 20GB memory.

Table I summarizes the statistics of the data sets, taken from LibSVM [20], that we consider in the experiments. To construct the `Covtype-I` set, we randomly sample 14129 samples from the first two classes of the original `Covtype` data set and merge the data samples of classes 4, 5 and 6 into one single class to avoid severe imbalances between classes. `Covtype-II`, instead, uses all of the original `Covtype` data.

We organize the experimental study in two parts. In both parts, we aim to show the performance of the proposed MBSC algorithm against other state-of-the-art algorithms for solving spectral clustering. In particular, we compare our proposal with the solution of the power iteration-based algorithm in [6] and the Nyström approximation-based spectral clustering in [7]. We also include results obtained by the spectral clustering algorithm in [18], where, to ensure fairness with our proposal, we sample the same number of non-zero entries from L and employ Adagrad optimization. We refer to this method as PSGD. A comparison between PSGD and MBSC is illustrated in Figure 4

In the first part, experiments are carried out on moderately




Fig. 4: NMI versus counts of floating-point operations on the `Pendigits`, `Shuttle` and `MNIST` data sets. In this figure, we compare the clustering performance of the proposed MBSC method with PSGD. For better visualization, we only compare PSGD with MBSC-20, MBSC-50 and MBSC-50 on the three datasets respectively, noted as PSGD-20, PSGD-50 and PSGD-50 in the corresponding figures. PSGD uses the same computational cost per iteration as MBSC.

large data sets comprising $10K$, $58K$ and $60K$ samples. For these, we can also report the performance of the spectral clustering approach in [1] where the spectrum of the normalized Laplacian is computed exactly (denoted as “Exact” in the plots). In the second part, we repeat the same comparison on two larger sized data sets, comprising $100K$ and $580K$ samples, where we cannot compute the exact spectrum of L .

A. Comparative evaluation

We conduct a comparative evaluation of the proposed MBSC with state-of-the-art spectral clustering algorithms on the `Pendigits`, `Shuttle` and `MNIST` data sets. In the proposed MBSC approach, we experiment with different choices of the mini-batch size to assess its impact on performance. The warm start initializations on the three databases are built using the Nyström approximation, by sampling only 5 data points from the whole dataset. To comprehensively analyze the performance of the proposed method, the iterative stochastic gradient descent on the manifold runs until we make one full pass through the whole data set. In practice, the stochastic gradient descent process can be stopped either if the difference between the spectral embedding W derived at two successive iteration steps is less than a threshold, or if a fixed number of iteration is achieved. For the Nyström approximation, we report a few choices on the number of samples selected to construct the approximate eigenvectors, namely 10, 50, 100, 500 and 1000.

Figure 2 illustrates the NMI scores achieved by the various methods versus the amount of floating-point operations. In the figure, for the proposed MBSC and the Nyström approximation we report the average plus and minus one standard deviation of the NMI score over 10 repetitions. In Figure 3 we show the impact of the warm start.

Recalling that m is the mini-batch size, k is the number of top eigenvectors and the number of clusters, each iteration of MBSC requires $2nkm + 6nk^2 - k^2$ floating point operations. The power method starts by generating an n -by- k Gaussian random matrix S , which costs $\mathcal{O}(nk)$ operations. It then computes a matrix product between the n -by- n affinity matrix and S , and iteratively applies the same multiplication for a total cost of $4n^2k - 2nk$ floating point operations. The final step

of the power method performs Singular Value Decomposition (SVD) on an n -by- k matrix. Since $n \gg k$, we adopt the estimate of SVD complexity in [21], which costs $2nk^2 + 2k^3$ operations. Finally, in order to calculate the number of floating point operations for the Nyström approximation we follow the pseudo code in [7], for which the total count of floating point operations is $6nk + 8k^3 - 3k^2 + 4nk^2 - 3k + 2nkm + nm + 2nm^2 + m^2 + m^3 - n$.

In Figure 2, we see how the variance of the clustering performance of MBSC diminishes throughout iterations. Larger mini-batches lead to faster variance reduction of the stochastic gradient, thus producing faster convergence to the solution. Compared with the power method, the proposed MBSC needs distinctively less computations, while achieving higher or similar clustering accuracy. Another interesting observation is that the proposed MBSC achieves stable clustering accuracy before it makes a full pass through the whole data set. The Nyström approximation is computationally fast on all three data sets. Nevertheless, its time and space complexity drastically increase with the number of inducing points. The PSGD method, as shown in Figure 4, converges more slowly than MBSC, producing comparable NMI scores on `Pendigits` and `Shuttle` data sets, and a lower NMI score on the `MNIST` data. Though stochastic gradients in PSGD have similar variance as in our proposal, MBSC samples columns rather than randomly selecting entries, and we speculate that this allows to preserve more global information of dense affinity structures. On the `Pendigits` and `Shuttle` data sets, with less running time, the proposed MBSC method requires smaller mini batches to conduct clustering and achieves better clustering performance than the Nyström approximation. On the `MNIST` data set, instead, the Nyström approximation produces strikingly good clustering results when the number of selected samples is larger than 500. However, the proposed MBSC method converges to the clustering accuracy of the exact spectral clustering even when the size of the mini-batch size is small, e.g., less than 100. As shown in Figure 3, the warm start initialization is effective in leading to faster convergence and a reduced variance, with minor extra cost.

TABLE II: Running time comparison on Covtype-I and Covtype-II data. T is the number of iterations.

Covtype-I					Covtype-II				
Algorithm	T	Mean NMI	Std NMI	time (s)	Algorithm	T	Mean NMI	Std NMI	time (s)
MBSC-E-400	230	0.40	0.01	1088	MBSC-E-1000	325	0.11	0.01	8400
MBSC-E-800	230	0.40	0.01	2190	MBSC-E-2000	325	0.15	0.01	13510
Nyström-400	-	0.38	0.01	78	Nyström-1000	-	0.09	0.01	2520
Nyström-800	-	0.38	0.01	965	Nyström-2000	-	0.11	0.01	11420
Power method	3	0.40	-	9300	Power Method	Too expensive			

B. Use case: spectral clustering on larger data sets

To demonstrate that the proposed approach can tackle large-scale spectral clustering problems, we implemented the proposed MBSC algorithm computing the necessary columns of L to construct stochastic gradients on-the-fly, and applied it to two data sets comprising 100K and 580K samples. Table II compares this version of MBSC, that we denote as MBSC-E, with the Nyström approximation and the power method. MBSC-E produces stable clustering results after 200 iterations. The running time and NMI are reported as average and standard deviation over 5 repetitions. In the comparison, the Nyström approximation selects a subset of the same size of the mini-batch in MBSC-E. On the Covtype-II data set, the power method fails to obtain clustering results within an acceptable time, and we omit it from Table II.

While running time is heavily dependent on implementation and system architecture, we argue that this is probably in favor of the Nyström approximation, for which we are using well optimized scientific computing packages. In any case, the purpose of this experiment is to demonstrate that the proposed MBSC algorithm, avoiding approximations to the Laplacian, can achieve higher performance than approximate methods on large-scale spectral clustering problems. Crucially, we demonstrate that this is possible at a comparable computational cost with approximate methods.

On both Covtype-I and Covtype-II data sets, MBSC achieves consistently better clustering accuracy than the Nyström approximation. Because the computational cost of the Nyström approximation rapidly increases with the size of the approximating set, it requires longer than MBSC-E to achieve a comparable clustering accuracy. Furthermore, compared with the power method, MBSC-E shows superior computational efficiency given large-scale data sets. Remarkably, MBSC-E requires less than 1GB and 3.6GB memory to run on the two data sets, respectively.

V. CONCLUSIONS

With the aim of improving scalability of normalized cut-based spectral clustering, we formulated it as an optimization problem with an orthonormality constraint, and solved it using stochastic gradient over Stiefel manifolds. We proposed a novel adaptive stochastic gradient optimization framework on Stiefel manifolds to compute the spectrum of Laplacian matrices, with computation of stochastic gradients linear in the number of samples. We provided theoretical justifications and empirical analyses to demonstrate how our proposal tackles large-scale spectral clustering in a practical way.

Our proposal is characterized by attractive robustness to parameter selection and scalability. It leads to the same clustering accuracy of spectral clustering approaches that use the exact spectrum of the Laplacian at a fraction of the cost. The results also support the motivation behind our proposal that approximate methods can potentially affect clustering performance. In cases where approximate methods perform well, as we reported in one of the experiments, we can see our proposal as a practical way to obtain clustering solutions as good as the gold-standard of the “exact” approach at a reasonable cost. Furthermore, the proposed method can compute the whole Laplacian matrix incrementally, making it an attractive option to tackle large-scale spectral clustering problems with a limited memory footprint.

We are investigating the possibility to combine our framework with other approximate methods, for example, to be able to afford more inducing points when conducting spectral clustering using the Nyström approximation. Another extension is to leverage approximations to reduce the variance of the stochastic gradients without introducing any bias, in order to accelerate stochastic gradient optimization. A Spark/TensorFlow implementation of MBSC is under development.

ACKNOWLEDGMENT

MF gratefully acknowledges support from the AXA Research Fund. We also thanks Dr. Yun Shen from Symantec Research Labs for helps on implementation and useful discussions.

REFERENCES

- [1] J. Shi and J. Malik, “Normalized Cuts and Image Segmentation,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 22, no. 8, pp. 888–905, 2000.
- [2] A. Y. Ng, M. I. Jordan, and Y. Weiss, “On Spectral Clustering: Analysis and an algorithm,” in *Advances in Neural Information Processing Systems 14*, Cambridge, MA, 2002.
- [3] I. S. Dhillon, Y. Guan, and B. Kulis, “Kernel k-means: spectral clustering and normalized cuts,” in *Proceedings of the tenth ACM SIGKDD international conference on Knowledge discovery and data mining*, New York, NY, USA, 2004, pp. 551–556.
- [4] M. Filippone, F. Camastra, F. Masulli, and S. Rovetta, “A Survey of Kernel and Spectral Methods for Clustering,” *Pattern Recognition*, vol. 41, no. 1, pp. 176–190, 2008.
- [5] U. von Luxburg, “A tutorial on spectral clustering,” *Statistics and Computing*, vol. 17, no. 4, pp. 395–416, 2007.
- [6] C. Boutsidis, P. Kambadur, and A. Gittens, “Spectral Clustering via the Power Method,” in *Proceedings of the 32nd International Conference on Machine Learning*, 2015, pp. 40–48.
- [7] C. Fowlkes, S. Belongie, F. Chung, and J. Malik, “Spectral Grouping Using the Nyström Method,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 26, no. 2, pp. 214–225, 2004.
- [8] D. Yan, L. Huang, and M. I. Jordan, “Fast Approximate Spectral Clustering,” in *Proceedings of the 15th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, New York, NY, USA, 2009, pp. 907–916.

- [9] Y. Li, J. Huang, and W. Liu, "Scalable Sequential Spectral Clustering," in *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence, February 12-17, 2016, Phoenix, Arizona, USA.*, 2016, pp. 1809–1815.
- [10] S. Bonnabel, "Stochastic gradient descent on Riemannian manifolds," *IEEE Transactions on Automatic Control*, vol. 58, no. 9, pp. 2217–2229, 2013.
- [11] J. Duchi, E. Hazan, and Y. Singer, "Adaptive Subgradient Methods for Online Learning and Stochastic Optimization," *Journal of Machine Learning Research*, vol. 12, pp. 2121–2159, 2011.
- [12] B. Schölkopf and A. J. Smola, *Learning with Kernels: Support Vector Machines, Regularization, Optimization, and Beyond*. Cambridge, MA, USA: MIT Press, 2001.
- [13] F. R. K. Chung, *Spectral Graph Theory (CBMS Regional Conference Series in Mathematics, No. 92)*. American Mathematical Society, 1997.
- [14] S. T. Smith, "Optimization Techniques on Riemannian Manifolds," 2014, arXiv:1407.5965.
- [15] Z. Wen and W. Yin, "A feasible method for optimization with orthogonality constraints," *Mathematical Programming*, vol. 142, no. 1, pp. 397–434, 2013.
- [16] J. Chen, M. Anitescu, and Y. Saad, "Computing $f(A)b$ via Least Squares Polynomial Approximations," *SIAM Journal on Scientific Computing*, vol. 33, no. 1, pp. 195–222, 2011.
- [17] M. Filippone and R. Engler, "Enabling scalable stochastic gradient-based inference for Gaussian processes by employing the Unbiased Linear System Solver (ULISSE)," in *Proceedings of the 32nd International Conference on Machine Learning*, 2015, pp. 1015–1024.
- [18] M. D. Collins, J. Liu, J. Xu, L. Mukherjee, and V. Singh, "Spectral clustering with a convex regularizer on millions of images," in *Computer Vision - ECCV 2014 - 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part III*, ser. Lecture Notes in Computer Science, D. J. Fleet, T. Pajdla, B. Schiele, and T. Tuytelaars, Eds., vol. 8691. Springer, 2014, pp. 282–298. [Online]. Available: http://dx.doi.org/10.1007/978-3-319-10578-9_19
- [19] L. Zelnik-Manor and P. Perona, "Self-tuning spectral clustering," in *Advances in Neural Information Processing Systems 17*, 2004, pp. 1601–1608.
- [20] C.-C. Chang and C.-J. Lin, "LIBSVM: A library for support vector machines," *ACM Transactions on Intelligent Systems and Technology*, vol. 2, pp. 27:1–27:27, 2011.
- [21] G. H. Golub and C. F. Van Loan, *Matrix computations*, 3rd ed. The Johns Hopkins University Press, 1996.