Published in final edited form as:

Inf Process Med Imaging. 2015 July ; 9123: 687–698. doi:10.1007/978-3-319-19992-4_54.

Sampling from Determinantal Point Processes for Scalable Manifold Learning

Christian Wachinger^{1,2} and Polina Golland¹

Christian Wachinger: wachinger@csail.mit.edu

¹Computer Science and Artificial Intelligence Lab, MIT, Cambridge, USA

²Massachusetts General Hospital, Harvard Medical School, Boston, USA

Abstract

High computational costs of manifold learning prohibit its application for large datasets. A common strategy to overcome this problem is to perform dimensionality reduction on selected landmarks and to successively embed the entire dataset with the Nyström method. The two main challenges that arise are: (i) the landmarks selected in non-Euclidean geometries must result in a low reconstruction error, (ii) the graph constructed from sparsely sampled landmarks must approximate the manifold well. We propose to sample the landmarks from determinantal distributions on non-Euclidean spaces. Since current determinantal sampling algorithms have the same complexity as those for manifold learning, we present an efficient approximation with linear complexity. Further, we recover the local geometry after the sparsification by assigning each landmark a local covariance matrix, estimated from the original point set. The resulting neighborhood selection based on the Bhattacharyya distance improves the embedding of sparsely sampled manifolds. Our experiments show a significant performance improvement compared to state-of-the-art landmark selection techniques on synthetic and medical data.

1 Introduction

Spectral methods are central for a multitude of applications in medical image analysis, computer vision, and machine learning, such as dimensionality reduction, classification, and segmentation. A limiting factor for the spectral analysis on large datasets is the computational cost of the eigen decomposition. To overcome this limitation, the Nyström method [21] is commonly applied to approximate the spectral decomposition of the Gramian matrix. A subset of rows/columns is selected and based on the eigen decomposition of the resulting small sub-matrix, the spectrum of the original matrix can be approximated. While the Nyström extension is the standard method for the matrix reconstruction, the crucial challenge is the subset selection. In early work [21], uniform sampling without replacement was proposed. This was followed by numerous alternatives including K-means clustering [22], greedy approaches [12], and volume sampling [3,9]. A recent comparison is presented in [16].

Of particular interest for subset selection is volume sampling [9], equivalent to determinantal sampling [3], because reconstruction error bounds exist. It is, however, not used in practice because of the high computational complexity of sampling from the underlying distributions [16]. Independently, determinantal point processes (DPPs) have been proposed recently for tracking and pose estimation [15]. They were originally designed to model the repulsive interaction between particles. DPPs are well suited for modeling diversity in a point set. A sampling algorithm for DPPs was presented in [14,15], which has complexity $O(n^3)$ for n points. Since this algorithm has the same complexity as the spectral analysis, it cannot be directly used as a subset selection scheme.

In this paper, we focus on nonlinear dimensionality reduction for large datasets via manifold learning. Popular manifold learning techniques include kernel PCA, Isomap [19], and Laplacian eigenmaps [5]. All of these methods are based on a kernel matrix of size $O(n^2)$ that contains the information about the pairwise relationships between the input points. The spectral decomposition of the kernel matrix leads to the low-dimensional embedding of the points. For large n , one seeks to avoid the explicit construction and storage of the matrix. In contrast to general rank- k matrix approximation, this is possible by taking the nature of the non-linear dimensionality reduction into account and relating the entries of the kernel matrix directly to the original point set.

We propose to perform DPP sampling on the original point set to extract a diverse set of landmarks. Since the input points lie in a non-Euclidean space, ignoring the underlying geometry leads to poor results. To account for the non-Euclidean geometry of the input space, we replace the Euclidean distance with the geodesic distance along the manifold, which is approximated by the shortest path distance on the graph. Due to the high complexity of DPP sampling, we derive an efficient approximation that runs in $O(ndk)$ with input dimensionality d and subset cardinality k . The algorithm restricts the updates to be local, which enables sampling on complex geometries. This, together with its low computational complexity, makes the algorithm well suited for the subset selection in large scale manifold learning.

A consequence of the landmark selection is that the manifold is less densely sampled than before, making its approximation with neighborhood graphs more difficult. It was noted in [2], as a critical response to [19], that the approximation of manifolds with graphs is topologically unstable. In order to improve the graph construction, we retain the local geometry around each landmark by locally estimating the covariance matrix on the original point set. This allows us to compare multivariate Gaussian distributions with the Bhattacharyya distance for neighborhood selection, yielding improved embeddings.

2 Background

We assume n points in high dimensional space $x_1, \dots, x_n \in \mathbb{R}^d$ and let $X \in \mathbb{R}^{d \times n}$ be the matrix whose i -th column represents point x_i . Non-linear dimensionality reduction techniques are based on a positive semidefinite kernel K , with a typical choice of Gaussian or heat kernel $K_{i,j} = \exp(-\|x_i - x_j\|^2/2\sigma^2)$. The resulting kernel matrix is of size $O(n^2)$. The eigen decomposition of the kernel matrix is necessary for spectral analysis. Unfortunately,

its complexity is $O(n^3)$. Most techniques require only the top k eigenvectors. The problem can therefore also be viewed as finding the best rank- k approximation of the matrix K , with the optimal solution $K_k = \sum_{i=1}^k \lambda_i u_i u_i^\top$, where λ_i is the i -th largest eigenvalue and u_i is the corresponding eigenvector.

2.1 Nyström Method

Suppose $J \subseteq \{1, \dots, n\}$ is a subset of the original point set of size k and $\bar{J}$ is its complement. We can reorder the kernel matrix K such that

$$K = \begin{bmatrix} K_{J \times J} & K_{J \times \bar{J}} \\ K_{\bar{J} \times J}^\top & K_{\bar{J} \times \bar{J}} \end{bmatrix}, \quad \widetilde{K} = \begin{bmatrix} K_{J \times J} & K_{J \times \bar{J}} \\ K_{\bar{J} \times J}^\top & K_{\bar{J} \times \bar{J}}^\top K_{J \times J}^{-1} K_{J \times \bar{J}} \end{bmatrix} \quad (1)$$

where $\widetilde{K}$ is the matrix estimated via the Nyström method [21]. The Nyström extension leads to the approximation $K_{\bar{J} \times \bar{J}} \approx K_{\bar{J} \times J}^\top K_{J \times J}^{-1} K_{J \times \bar{J}}$. The matrix inverse is replaced by the Moore-Penrose generalized inverse in the case of rank deficiency. The Nyström method leads to the minimal kernel completion [3] conditioned on the selected landmarks and has been reported to perform well in numerous applications [8,13,18]. The challenge lies in finding landmarks that minimize the reconstruction error

$$\|K - \widetilde{K}\|_{tr} = tr(K_{\bar{J} \times \bar{J}}) - tr(K_{\bar{J} \times J}^\top K_{J \times J}^{-1} K_{J \times \bar{J}}). \quad (2)$$

The trace norm $\|\cdot\|_{tr}$ is applied because results only depend on the spectrum due to its unitary invariance.

2.2 Annealed Determinantal Sampling

A large variety of methods have been proposed for selecting the subset J . For general matrix approximation, this step is referred to as row/column selection of the matrix K , which is equivalent to selecting a subset of points X . This property is important because it avoids explicit computation of the $O(n^2)$ entries in the kernel matrix K . We focus on volume sampling for subset selection because of its theoretical advantages [9]. We employ the factorization $K_{J \times J} = Y_J^\top Y_J$, which exists because K_J is positive semidefinite. Columns in Y_J can be thought of as feature vectors describing the selected points. Based on this factorization, the volume $\text{Vol}(\{Y_i\}_{i \in J})$ of the simplex spanned by the origin and the feature vectors Y_J is calculated, which is equivalent to the volume of the parallelepiped spanned by Y_J . The subset J is then sampled proportionally to the squared volume. This is directly related to the calculation of the determinant with

$\det(K_{J \times J}) = \det(Y_J^\top Y_J) = \det(Y_J)^2 = \text{Vol}^2(\{Y_i\}_{i \in J})$. These ideas were further generalized in [3] based on annealed determinantal distributions

$$p^s(J) \propto \det(K_{J \times J})^s = \det(Y_J^\top Y_J)^s = \det(Y_J)^{2s}. \quad (3)$$

This distribution is well defined because the principal submatrices of a positive semidefinite matrix are themselves positive semidefinite. Varying the exponent $s \geq 0$ results in a family of distributions, modeling the annealing behavior as used in stochastic computations. For $s = 0$ this is equivalent to uniform sampling [21]. In the following derivations, we focus on $s = 1$. It was shown in [9] that for $J \sim p(J)$, $|J| = k$

$$\mathbb{E} \left[\|K - \widetilde{K}\|_{tr} \right] \leq (k+1) \|K - K_k\|_F^2, \quad (4)$$

where $\widetilde{K}$ is the Nyström reconstruction of the kernel based on the subset J , K_k the best rank- k approximation achieved by selecting the largest eigenvectors, and $\|\cdot\|_F$ the Frobenius norm. It was further shown that the factor $k + 1$ is the best possible for a k -subset. Related bounds were presented in [4].

3 Method

We first analyze the sampling from determinantal distributions on non-Euclidean geometries. We then introduce an efficient algorithm for approximate DPP sampling on manifolds. Finally, we present our approach for robust graph construction on sparsely sampled manifolds.

3.1 DPP Sampling on Manifolds

As described in Sect. 2.2, sampling from determinantal distributions is used for row/column selection. Independently, determinantal point processes (DPPs) were introduced for modeling probabilistic mutual exclusion. They present an attractive scheme for ensuring diversity in the selected subset. Here we work with the construction of DPPs based on L-ensembles [7]. Given a positive semidefinite matrix $L \in \mathbb{R}^{n \times n}$, the likelihood for selecting the subset $J \subseteq \{1, \dots, n\}$ is

$$P_L(J) = \frac{\det(L_{J \times J})}{\det(L+I)}, \quad (5)$$

where I is the identity matrix and $L_{J \times J}$ is the sub-matrix of L containing the rows and columns indexed by J . By associating the L-ensemble matrix L with the kernel matrix K , we can apply DPPs to sample subsets from the point set X .

To date, applications using determinantal point processes have assumed Euclidean geometry [15]. For non-linear dimensionality reduction, we assume that the data points lie in a non-Euclidean space, such as the Swiss roll in Fig. 1(a). To evaluate the performance of DPPs on manifolds, we sample from the Swiss roll. Since we know the construction rule in this case, we can invert it and display the sampled 3D points in the underlying 2D space. The result in Fig. 1(b) shows that the inner part of the roll is almost entirely neglected, as a consequence of not taking the manifold structure into account. A common solution is to use geodesic distances [19], which can be approximated by the graph shortest path algorithm. We replace the Euclidean distance $\|\cdot\|$ in the construction of the kernel matrix K with the geodesic distance $K_{i,j} = \exp(-\|x_i - x_j\|_{geo}^2 / 2\sigma^2)$ to obtain the result in Fig. 1(c). We observe a clear

improvement in the diversity of the sampling, now also including points in the interior part of the Swiss roll.

3.2 Efficient Approximation of DPP Sampling on Manifolds

While it is possible to adapt determinantal sampling to non-Euclidean geometries and to characterize the error for the subset selection, we are missing an efficient sampling algorithm for handling large point sets. In [9], an approximative sampling based on the Markov chain Monte Carlo method is proposed to circumvent the combinatorial problem

with $\binom{n}{k}$ possible subsets. Further approximations include sampling proportionally to the diagonal elements K_{ii} or their squared version K_{ii}^2 , leading to additive error bounds [4,11]. In [10], an algorithm is proposed that yields a $k!$ approximation to volume sampling, worsening the approximation from $(k+1)$ to $(k+1)!$.

Algorithm 1. DPP sampling equivalent to [15]

Require: Eigen decomposition of $K: \{(v_i, \lambda_i)\}_{i=1}^n$

```

1: Initialize  $V = \emptyset$ 
2: for  $i = 1$  to  $n$  do
3:   Add eigenvector  $v_i$  with probability  $\frac{\lambda_i}{\lambda_i + 1}$  to  $V$ 
4: end for
5:  $B = V^\top$ 
6: for  $1$  to  $|V|$  do
7:   Select  $i \in 1 \dots n$  with probability  $P(i) \propto \|B_i\|^2$ 
8:    $J \leftarrow J \cup i$ 
9:    $B_j \leftarrow \text{Proj}_{\perp B_i} B_j$  for all  $j \in \{1, \dots, n\}$ 
10: end for
11: return  $J$ 

```

An exact sampling algorithm for DPPs was presented in [14,15], which requires the eigen decomposition of $K = \sum_{i=1}^n \lambda_i v_i v_i^\top$. Algorithm 1 states an equivalent formulation of this sampling approach. First, eigenvectors are selected proportionally to the magnitude of their eigenvalues and stored as columns in V . Assuming m vectors are selected, $V \in \mathbb{R}^{n \times m}$. By setting $B = V^\top$, we use $B_i \in \mathbb{R}^m$ to denote the rows of V . In each iteration, we select one of the n points where point i is selected proportionally to the squared norm $\|B_i\|^2$. The selected point is added to the subset J . After the selection of i , all vectors B_j are projected to the orthogonal space of B_i . Since $\text{Proj}_{\perp B_i} B_i = 0$, the same point is almost surely not selected twice. The update formulation differs from [15], where an orthonormal basis of the eigenvectors in V perpendicular to the i -th basis vector $e_i \in \mathbb{R}^n$ is constructed. Both formulations are equivalent but provide a different point of view on the algorithm. This modification is essential to motivate the proposed efficient sampling procedure. The following proposition characterizes the behavior of the update rule in the algorithm.

Proposition 1—Let $B_i, B_j \in \mathbb{R}^m \setminus \{0\}$ be two non-zero vectors in $\mathbb{R}^m$, and $\theta = \angle(B_i, B_j)$ be the angle between them. Then

$$\|Proj_{\perp B_i} B_j\|^2 = \|B_j\|^2 \sin^2 \theta, \quad (6)$$

where $Proj_{\perp B_i} B_j$ is the projection of B_j on the subspace perpendicular to B_i . For $B_i = 0$ and $B_j = 0$ the projection is $\|Proj_{\perp B_i} B_j\|^2 = 0$.

Sampling from a determinantal distribution is not only advantageous because of the presented error bounds but it also makes intuitive sense that a selection of a diverse set of points yields a more accurate matrix reconstruction. The computational complexity of Algorithm 1 is, however, similar or even higher than that of manifold learning because the spectral decomposition of a dense graph is required, whereas Laplacian eigenmaps operate on sparse matrices. An approach for efficient sampling proposed in [15] works with the dual representation of $K = Y^T Y$ to obtain $Q = Y Y^T$, with Q hopefully smaller than the matrix K . Considering that we work with a Gaussian kernel matrix, this factorization corresponds to the inner product in feature space $\phi(x_i)^T \phi(x_j)$ of the original points x_i, x_j . The Gaussian kernel corresponds to an infinite dimensional feature space. Since we work with symmetric, positive definite matrices, we can calculate a Cholesky decomposition. In this case, the dual representation has the same size as the original matrix and therefore yields no improvement.

To overcome the high computational costs of exact DPP sampling, we present an efficient approximation in Algorithm 2. The computational complexity is $O(ndk)$. Vector $D \in \mathbb{R}^n$ models the probabilities for the selection of points as does $\|B_i\|^2$ in the original DPP sampling algorithm. The algorithm proceeds by sampling k points. At each iteration we select one point x_i with probability $p(i) \propto D_i$. Next we compute distances $\{d_{ij}\}_{j=1 \dots n}$ of the selected point x_i to all points in X . Based on these distances we identify a local neighborhood N_i of m nearest neighbors around the selected point x_i . The update of the probabilities D is restricted to the neighborhood N_i , which proves advantageous for sampling on manifolds. In contrast, Algorithm 1 updates probabilities for all points. If we seek a similar behavior to Algorithm 1, the local neighborhood should include all points. The update function f takes distances as input, where we consider $f(d) = \sin^2(d/\tau)$ and $f(d) = (1 - \exp(-d^2/2\sigma^2))$, as motivated below. In subsequent iterations of the algorithm, points close to x_i are selected with lower probability.

Algorithm 2. Efficient approximation of DPP sampling

Require: Point set X , subset cardinality k , nearest neighbor count m , update function f

- 1: Initialize $D = \mathbf{1}_n$ and $J = \emptyset$
- 2: **for** 1 to k **do**
- 3: Select $i \in 1 \dots n$ with probability $p(i) \propto D_i$
- 4: $J \leftarrow J \cup \{i\}$
- 5: Compute $d_{ij} = \|x_i - x_j\|, \forall j \in 1 \dots n$
- 6: Set m nearest neighbors of x_i as neighborhood N_i based on $\{d_{ij}\}_{j=1 \dots n}$
- 7: $D_j \leftarrow D_j \cdot f(d_{ij}), \forall j \in N_i$
- 8: Optional: Compute covariance C_i in local neighborhood N_i around x_i

```

9:   end for
10:  return  $J$  and optionally  $\{Ci\}_{i \in J}$ 

```

We initialize the vector $D = \mathbf{1}_n$, since it was noted in [15] that the squared norm of the vectors $\|B_i\|^2$ gives rise initially to a fairly uniform distribution because no points have yet been selected. For the update step, Proposition 1 implies that the update of the probabilities of selecting specific points acts as $\sin^2(\theta)$. The angle $\theta = \angle(B_i, B_j)$ correlates strongly with the distance $\|B_i - B_j\|$, since $\|B_j\|$ and $\|B_i\|$ are initially the same. In Algorithm 1, the largest eigenvalues are selected with high probability. We can therefore draw the analogy to multidimensional scaling (MDS) [20] with a Gaussian kernel, where MDS selects the top eigenvectors. Consequently, vectors B_i correspond to low-dimensional embeddings produced by multidimensional scaling of original points x_i . MDS preserves pairwise distances between the original space and the embedding space, enabling the approximation $\|B_i - B_j\| \approx \|x_i - x_j\|$. We approximate the update based on the distance of points in the original space, $\sin^2(\theta) \approx \sin^2(\|x_i - x_j\|/\tau)$. The scaling factor τ ensures that values are in the range $[-\pi/2; \pi/2]$. This update is similar to the Welsch function $(1 - \exp(-\|x_i - x_j\|^2/2\sigma^2))$, which is directly related to the weights in the kernel matrix and is commonly used in machine learning. For subsequent iterations of the algorithm, the assumption of a similar norm of all vectors B_i is violated, because the projection on the orthogonal space changes their lengths. However, this change is locally restricted around the currently selected point. Since this region is less likely to be sampled in the subsequent iterations, the assumption still holds for parts of the space that contain most probability.

Remark—The proposed algorithm bears similarities to K-means++ [1], which replaces the initialization through uniform sampling of K-means by a new seeding algorithm. K-means++ seeding is a heuristic that samples points based on their distance to the closest selected landmark. Initially, when only one landmark is selected, our algorithm has a nearly identical update rule for a maximal neighborhood N_i . In later iterations, the algorithms differ because K-means++ bases the selection only on the distance to the nearest landmark, while all landmarks influence the probability space in our algorithm. Consequently, our approach potentially yields subsets with higher diversity.

3.3 Robust Landmark-Based Graph Construction

After selecting the landmarks, the next step in the spectral analysis consists of building a graph that approximates the manifold. Common techniques for the graph construction include selection of nearest neighbors or ε -balls around each node. Both approaches require the setting of a parameter, either the number of neighbors or the size of the ball, which is crucial for the performance. Setting the parameter too low leads to a large number of disconnected components, while for many applications one is interested in having all points connected to obtain a consistent embedding of all points. Choosing too high values of the parameters leads to shortcuts, yielding a poor approximation of the manifold. The appropriate selection of the parameters is more challenging on sparsely sampled manifolds, obtained after the subset selection.

To address this issue, we propose a new technique for graph construction that takes the initial distribution of the points into account. For each landmark x_i , we estimate the covariance matrix C_i around this point from its nearest neighbors N_i , as indicated as optional step in Algorithm 2. This step implies a multivariate Gaussian distribution $G(x_i, C_i)$ centered at the landmark x_i . A commonly used distance to compare distributions is the Bhattacharyya distance, which in case of Gaussian distributions corresponds to

$$B(G_i, G_j) = \frac{1}{8}(x_i - x_j)^T C^{-1}(x_i - x_j) + \frac{1}{2} \ln \left(\frac{|C|}{\sqrt{|C_i||C_j|}} \right), \quad (7)$$

with $C = \frac{C_i + C_j}{2}$. This distance is less likely to produce shortcuts across the manifold because points that fit the local geometry appear much closer than points that appear as outliers from the local geometry. Consequently, we replace the Euclidean distance for neighborhood selection in manifold learning with the Bhattacharyya distance. Space requirements of this step are $O(d^2 k)$. An alternative for limited space and large d is to only use the diagonal entries of the covariance matrix, requiring $O(dk)$ space.

4 Experiments

In our first experiment, we show that the proposed efficient DPP sampling algorithm is well suited for subset selection in non-Euclidean spaces. The algorithm restricts the update of the sampling probability D to a local neighborhood N_i around the current point x_i . This is in line with the motivation of many manifold learning algorithms that assume that the space behaves locally like a Euclidean space. In our experiment, we set the local neighborhood N_i to be the 20 nearest neighbors around the selected point x_i . The sampling result is shown in Fig. 1(d). We obtain a point set with high diversity, covering the entire manifold. This illustrates that the proposed algorithm provides diversity on complex geometries and is therefore appropriate for subset selection in the context of non-linear dimensionality reduction.

In our second experiment, we quantify the reconstruction error for matrix completion as formulated in Eq. (2). We compare the efficient DPP sampling result with uniform sampling [21] and K-means clustering with uniform seeding [22], which represents current state-of-the-art [16]. Moreover, we compare to selecting the subset with the K-means++ seeding and the K-means++ algorithm, which we have not been used for landmark selection before. We construct a Gaussian kernel matrix from 1,000 points on a Swiss roll (Fig. 1). We select subsets of varying size between 25 and 100 and set the parameters $\sigma = 1$ and $m = 30$ for the Swiss roll. Note that a further improvement can be achieved by adapting these parameters to the size of the subset. For smaller subsets, larger σ and m lead to improvements. We report the reconstruction error for the different methods and datasets in Table 1, averaged over 50 different runs. The diverse set of landmarks selected with efficient DPP sampling leads to the lowest average reconstruction errors in almost all settings.

In our third experiment, we perform manifold learning with Laplacian eigenmaps on a point set consisting of 10 million points on a Swiss roll. The dataset is too large to apply manifold

learning directly. We select 2,500 landmark points with the efficient DPP sampling algorithm and estimate the local covariance matrices, which we provide to the manifold learning algorithm. We vary the number of nearest neighbors in the graph from 25 to 500. We compare the graph construction with Euclidean and Bhattacharyya neighborhood selection with embedding results shown in Fig. 2. The results show that the Bhattacharyya based neighborhood selection is more robust with respect to the number of neighbors.

4.1 Image Data

After having evaluated each of the steps of the proposed approach separately, we now present results for scalable manifold learning on image data. We work with two datasets, one consisting of handwritten digits and a second one consisting of patches extracted from 3D medical images. Each dataset is too large to apply manifold learning directly. Consequently, we select landmarks with the discussed method, perform manifold learning on the landmarks with the Bhattacharyya distance, and use the Nyström method to embed the entire point set. We only consider the diagonal entries of the covariance matrices due to space limitations. The exact formula for the out-of-sample extension with the Nyström method for Laplacian eigenmaps is derived in [6]. To evaluate the quality of the embedding, we use the labels associated to the image data to perform nearest neighbor classification in the low dimensional space. We expect advantages for the DPP landmark selection scheme because a diverse set of landmarks spreads the entire point set in the embedding space and helps the classification. We avoid sophisticated pre-processing of the data, since we are only interested in the relative performances across the different landmark selection methods.

We evaluate the method on the MNIST dataset [17], consisting of 60,000 binary images of handwritten digits for training and 10,000 for testing with a resolution of 28×28 pixels, and on CT scans with a resolution of $512 \times 512 \times 145$ voxels. The CT scans were acquired for radiation therapy of patients with head and neck tumors. The figure on the right shows one cross sectional slice with segmentations of three structures of risk: left parotid (red), right parotid (blue), and brainstem (green). The segmentation of these structures during treatment planning is of high clinical importance to ensure that they obtain a low radiation dose. We extract image patches from the left and right parotid glands, the brainstem and the surrounding background. We aim to classify patches into these four groups, where the outcome can readily serve in segmentation algorithms. We work with patches of size $7 \times 7 \times 3$ to reflect the physical resolution of the data which is $0.98 \times 0.98 \times 2.5 \text{ mm}^3$. This results in roughly 150,000 patches extracted from three scans. 80,000 patches are used for training and the remaining ones for testing. We set $m = 5000$ and $\sigma = 5$ in this experiment. We embed the images into 100 dimensional space with Laplacian eigenmaps. Figure 3 reports classification results over 20 repetitions for several landmark selection schemes across different numbers of landmarks k , as well as the Bhattacharyya based graph construction. The results suggest that the K-means++ seeding outperforms the uniform initialization and K-means++ cannot further improve the initialization. Moreover, we observe a significant improvement in classification performance for approximate DPP sampling compared to K-Means++ seeding. Finally, the Bhattacharyya based graph construction further improves the results. In addition to the significant improvement, our runtime measurements showed that

our unoptimized Matlab code for efficient DPP sampling runs approximately 15% faster than an optimized MEX version of K-means.

5 Conclusion

We have presented contributions for two crucial issues of scalable manifold learning in large datasets: (i) efficient sampling of diverse subsets from manifolds and (ii) robust graph construction on sparsely sampled manifolds. We analyzed the sampling from determinantal distributions in non-Euclidean spaces and proposed an efficient approximation of DPP sampling. Furthermore, we proposed the local covariance estimation around landmarks to capture the local geometry of the space. This enabled a more robust graph construction with the Bhattacharyya distance and yielded low dimensional embeddings of higher quality.

Acknowledgments

This work was supported by the Humboldt foundation, the National Alliance for Medical Image Computing (U54-EB005149), the NeuroImaging Analysis Center (P41-EB015902), and the Wistron Corporation.

References

1. Arthur D, Vassilvitskii S. k-means++: The advantages of careful seeding. SODA. 2007:1027–1035.
2. Balasubramanian M, Schwartz EL. The Isomap algorithm and topological stability. Science. 2002; 295(5552):7. [PubMed: 11778013]
3. Belabbas M, Wolfe P. On landmark selection and sampling in high-dimensional data analysis. PNAS. 2009
4. Belabbas M, Wolfe P. Spectral methods in machine learning and new strategies for very large datasets. PNAS. 2009; 106(2):369–374. [PubMed: 19129490]
5. Belkin M, Niyogi P. Laplacian eigenmaps for dimensionality reduction and data representation. Neural Comput. 2003; 15(6):1373–1396.
6. Bengio Y, Paiement J, Vincent P, Delalleau O, Le Roux N, Ouimet M. Out-of-sample extensions for lle, isomap, mds, eigenmaps, and spectral clustering. NIPS. 2004
7. Borodin, A. Determinantal point processes. In: Akemann, G.; Baik, J.; Francesco, PD., editors. The Oxford Handbook of Random Matrix Theory. Oxford University Press; Oxford: 2011. p. 231–249.
8. Chen W, Song Y, Bai H, Lin C, Chang E. Parallel spectral clustering in distributed systems. TPAMI. 2011; 33(3):568–586.
9. Deshpande A, Rademacher L, Vempala S, Wang G. Matrix approximation and projective clustering via volume sampling. SODA. 2006
10. Deshpande, A.; Vempala, SS. Adaptive sampling and fast low-rank matrix approximation. In: Díaz, J.; Jansen, K.; Rolim, JDP.; Zwick, U., editors. APPROX 2006 and RANDOM 2006. Vol. 4110. Springer; Heidelberg: 2006. p. 292–303. LNCS
11. Drineas P, Mahoney MW. On the nystrom method for approximating a gram matrix for improved kernel-based learning. JMLR. 2005; 6:2153–2175.

12. Farahat A, Ghodsi A, Kamel M. A novel greedy algorithm for nyström approximation. IASTATS. 2011
13. Fowlkes C, Belongie S, Chung F, Malik J. Spectral grouping using the nyström method. TPAMI. 2004; 26(2):214–225.
14. Hough J, Krishnapur M, Peres Y, Virág B. Determinantal processes and independence. Probab Surv. 2006; 3:206–229.
15. Kulesza A, Taskar B. Structured determinantal point processes. NIPS. 2010
16. Kumar S, Mohri M, Talwalkar A. Sampling techniques for the nystrom method. J Mach Learn Res. 2012; 13:981–1006.
17. LeCun Y, Bottou L, Bengio Y, Haffner P. Gradient-based learning applied to document recognition. Proc IEEE. 1998; 86(11):2278–2324.
18. Talwalkar A, Kumar S, Rowley H. Large-scale manifold learning. CVPR. 2008
19. Tenenbaum J, Silva V, Langford J. A global geometric framework for nonlinear dimensionality reduction. Science. 2000; 290(5500):2319–2323. [PubMed: 11125149]
20. Williams C. On a connection between kernel PCA and metric multidimensional scaling. NIPS. 2001:675–681.
21. Williams C, Seeger M. Using the nyström method to speed up kernel machines. NIPS. 2001:682–688.
22. Zhang, K.; Tsang, I.; Kwok, J. ICML. ACM; 2008. Improved nyström low-rank approximation and error analysis; p. 1232-1239.

Fig. 1.

DPP sampling from 1,000 points lying on a manifold. We show results for standard DPP sampling, geodesic DPP sampling, and efficient DPP sampling. Note that the sampling is performed in 3D, but we can plot the underlying 2D manifold by reversing the construction of the Swiss roll. Geodesic and efficient sampling yields a diverse subset from the manifold.

Fig. 2.

Selection of 2,500 landmarks from a set of 10 million points. Embedding of landmarks into 2D with Laplacian eigenmaps. Results for Euclidean (first row) and Bhattacharyya neighborhood selection (second row) are shown.

Fig. 3. Classification accuracy for MNIST and head and neck data. Comparison of different subset selection schemes for varying numbers of selected landmarks. Bars indicate mean classification performance and error bars correspond to standard deviation. *, **, and *** indicate significance levels at 0.05, 0.01, and 0.001.

Table 1

Average reconstruction errors over 50 runs for several sampling schemes with subset sizes varying from 25 to 100.

Sampling	25	50	60	70	80	90	100
Uniform	70.384	8.006	4.838	2.785	1.676	0.731	0.442
K-means Uniform	28.124	3.848	2.319	1.393	0.756	0.403	0.235
K-means++ Seeding	50.114	5.832	3.033	1.655	1.013	0.683	0.347
K-means++	24.954	3.575	1.915	1.018	0.711	0.383	0.222
Efficient DPP	33.036	3.371	1.466	0.844	0.488	0.312	0.202

Best results are highlighted in bold face.