

STOCHASTIC PARALLELIZABLE EIGENGAP DILATION FOR LARGE GRAPH CLUSTERING

Elise van der Pol*

University of Amsterdam, Amsterdam, Netherlands.
elisevanderpol@gmail.com

Ian Gemp

DeepMind, London, UK
imgemp@deepmind.com

Yoram Bachrach

DeepMind, London, UK

Richard Everett

DeepMind, London, UK

ABSTRACT

Large graphs commonly appear in social networks, knowledge graphs, recommender systems, life sciences, and decision making problems. Summarizing large graphs by their high level properties is helpful in solving problems in these settings. In spectral clustering, we aim to identify clusters of nodes where most edges fall within clusters and only few edges fall between clusters. This task is important for many downstream applications and exploratory analysis. A core step of spectral clustering is performing an eigendecomposition of the corresponding graph Laplacian matrix (or equivalently, a singular value decomposition, SVD, of the incidence matrix). The convergence of iterative singular value decomposition approaches depends on the eigengaps of the spectrum of the given matrix, i.e., the difference between consecutive eigenvalues. For a graph Laplacian corresponding to a well-clustered graph, the eigenvalues will be non-negative but very small (much less than 1) slowing convergence. This paper introduces a parallelizable approach to dilating the spectrum in order to accelerate SVD solvers and in turn, spectral clustering. This is accomplished via polynomial approximations to matrix operations that favorably transform the spectrum of a matrix without changing its eigenvectors. Experiments demonstrate that this approach significantly accelerates convergence, and we explain how this transformation can be parallelized and stochastically approximated to scale with available compute.

1 INTRODUCTION

Large graphs commonly appear in social networks, knowledge graphs, recommender systems, life sciences, and decision making problems. In spectral clustering, we aim to identify node clusters so that most edges fall within clusters and only few edges fall between clusters. This task is important for many downstream applications such as community detection in sociology or biology (Fortunato, 2010), image segmentation (Coleman & Andrews, 1979), generating or refining labels (Song et al., 2015), exploratory data analysis (Kumar & Bezdek, 2020), and more abstractly, approximating solutions to combinatorial graph problems like min-cut (Chung & Graham, 1997).

A key intermediate step in spectral clustering (Von Luxburg, 2007) performs an eigendecomposition of the corresponding graph Laplacian matrix. Specifically, the eigenvectors associated with the bottom- k eigenvalues of the graph Laplacian define k -dimensional embeddings for each node in the graph providing a “soft” clustering. This resulting embedding well-separates nodes that belong in separate clusters, making a final “hard” clustering step, e.g., with k -means, relatively trivial.

For large graphs, computing a full eigendecomposition can be computationally costly, as doing an eigendecomposition of the graph Laplacian has complexity $\mathcal{O}(|\mathcal{E}||\mathcal{V}|^2)$ where $|\mathcal{V}|$ is the number of nodes and $|\mathcal{E}|$ the number of edges (Allen-Zhu & Li, 2017).

In practice, we are only interested in a subset of k eigenvectors. The convergence rate of many iterative solvers depends on the normalized eigengap, i.e. the difference between consecutive eigenvalues

*Work completed while at DeepMind.

relative to the spectral radius (Balcan et al., 2016; Gemp et al., 2021a). For the bottom- k eigenvalues of a large graph Laplacian, the eigengap may be quite small compared to the spectral radius, which can inhibit convergence rates.

This paper introduces *stochastic parallelizable eigengap dilation* or SPED, an approach to accelerate eigendecompositions, particularly for the learning of embeddings for clustering large graphs. SPED improves convergence rates of bottom- k eigendecompositions for large graphs by manipulating the graph Laplacian with cheap eigenvector-preserving transformations during eigenvector computation.

Our primary contribution is the development of an approach that favorably transforms the spectrum of the graph Laplacian in a way that 1) can be computed cheaply, 2) admits unbiased estimates, 3) can be parallelized, and 4) critically, maintains the original eigenvectors along with their rank.

In Section 2, we review the graph Laplacian and its connection to clustering via relaxation of min-cut from combinatorial graph theory. Next, Section 3 gives background on computing eigenvectors of large matrices with special attention to computing eigenvectors of graph Laplacians. We then describe our approach, SPED, in Section 4. Section 5 then evaluates the viability of SPED by examining it in the non-stochastic, sequential setting on several different domains: clustering large graphs, clustering large graphs with uncertain edges filled in by *link prediction*, and computing proto-value functions for use in reinforcement learning. Section 6 concludes. Appendix A presents additional experiments and Appendix B reviews related work¹.

2 SPECTRA OF GRAPH LAPLACIANS

The Laplacian of a graph $\mathcal{G} = (\mathcal{V}, \mathcal{E})$ with node set (vertices) \mathcal{V} and edge set \mathcal{E} is given by $L = D - A$ where D is a diagonal matrix containing the degree of each node and A is the adjacency matrix of the graph \mathcal{G} . The graph Laplacian can also be represented as $L = X^\top X = \sum_{e \in \mathcal{E}} x_e x_e^\top$ where the incidence matrix X encodes edges on its rows. Each row $x_e \in \mathbb{R}^{|\mathcal{V}|}$ corresponds to an edge (i, j) and contains two nonzero entries: a 1 at index $\min(i, j)$ and a -1 at index $\max(i, j)$. Note that $L\mathbf{1} = \sum_{e \in \mathcal{E}} x_e x_e^\top \mathbf{1} = \sum_{e \in \mathcal{E}} x_e (1 - 1) = \mathbf{0}$ implying the ones vector is an eigenvector of L with zero eigenvalue. We denote eigenvalues in increasing order by λ_i where $\lambda_i \leq \lambda_{i+1}$.

2.1 RELAXATION OF MIN CUT

Given the incidence matrix representation of the graph Laplacian, consider the following quadratic form (Von Luxburg, 2007):

$$\text{cut}(S, \bar{S}) = v^\top (X^\top X) v = \sum_{e \in \mathcal{E}} (x_e^\top v)^2 = \sum_{e=(i,j) \in \mathcal{E}} (v_i - v_j)^2 \quad (1)$$

where S is a set of nodes, \bar{S} is its complement, and v is a vector of length number of nodes that indicates whether a node i belongs to S ($v_{i \in S} = 1$) or not ($v_{i \notin S} = -1$). By inspection, equation 1 counts ($4\times$) the number of edges crossing between S and its complement \bar{S} .

Finding v such that the *cut* is minimized and $|v_i| = 1$ for all i is NP-Hard. Relaxing the feasible set to include all v such that $\|v\|_2^2 = |\mathcal{V}|$ results in an eigenvalue problem, in which case, $\|v\|_2^2 = |\mathcal{V}|$ can be replaced with $\|v\|_2^2 = 1$ w.l.o.g.:

$$\min_{v^\top \mathbf{1} = 0} v^\top L v \quad \text{s.t.} \quad v^\top v = 1 \quad (2)$$

where $v^\top \mathbf{1} = 0$ is included to avoid the trivial solution of assigning all nodes to the same cluster. Note the solution to equation 2 is the 2nd smallest eigenvector of L , also called the Fiedler vector. To assign hard clusters, we can then either threshold v or run k -means with $k = 2$ on v .

Cheeger’s inequality bounds a related objective called the *normalized cut* in terms of the eigenvalues of the graph Laplacian. Let

$$\phi_{\mathcal{G}}(S) = \frac{\text{cut}(S, \bar{S})}{\text{vol}(S)} \quad (3)$$

¹For the appendix, see <https://arxiv.org/pdf/2207.14589.pdf>.

where $\text{vol}(S)$ denotes the number of edges incident to nodes in S . And define the best cut of the graph to be the S (and \bar{S}) that solves

$$\rho_G = \min_{S \subset V} \max\{\phi_G(S), \phi_G(\bar{S})\}. \quad (4)$$

Then

$$\frac{\lambda_2}{2} \leq \rho_G \leq \sqrt{2\lambda_2} \quad (5)$$

where λ_2 is the 2nd smallest eigenvalue of L .

Therefore, as the number of edges in the graph increases (denominator of $\phi_G(S)$) relative to the number of edges crossing the cut (numerator), i.e., in a “well-clustered” graph, λ_2 shrinks.

The results and connections discussed so far have been generalized to the setting where the graph is partitioned into $k > 2$ sets (Lee et al., 2014). In this case, the bottom- k eigenvectors of the graph Laplacian are particularly useful for clustering. Similarly to before,

$$\frac{\lambda_k}{2} \leq \rho_G(k) \leq ck^2 \sqrt{\lambda_k} \quad (6)$$

for some $c > 0$ with

$$\rho_G(k) = \min_{S_1, S_2, \dots, S_k \subset V} \max\{\phi_G(S_i) : i = 1, 2, \dots, k\} \quad (7)$$

where $\{S_i\}$ represent all possible k -way partitions of the graph. Similar reasoning still applies—if there are k clear clusters in the graph, we should expect k eigenvalues $\ll 1$.

In addition, if the graph edges are weighted, the above analysis is extended by summing edge weights rather than counting edges ($w_{ij} = 1$). The graph Laplacian can be written more generally as $X^\top W X$ where W is a diagonal matrix containing the edge weights.

3 EIGENDECOMPOSITION

The incidence matrix representation of the graph Laplacian makes it clear the matrix is positive semi-definite. It also makes clear its connection to the singular value decomposition given the fact that the eigenvectors of $X^\top X$ are the same as the right singular vectors of X . This means various algorithms for SVD can be applied to spectral clustering.

Most scalable SVD methods focus on finding the top- k eigenvectors of $X^\top X$. To find the bottom- k eigenvectors, a common approach uses

$$L^- = \lambda^* I - L \quad (8)$$

with $\lambda^* > \lambda_{|V|}$, which turns the bottom- k eigenvectors of L into the top- k eigenvectors of L^- , and allows the use of a top- k solver for finding eigenvectors for spectral clustering.

The convergence of most scalable SVD methods also depends inversely on the eigengaps of the matrix. The eigengaps are the differences between consecutive eigenvalues

$$g_i = \lambda_{i+1} - \lambda_i \quad (9)$$

For example, in Gemp et al. (2021a), the number of steps to convergence depends on the ratio $\frac{\lambda_{|V|}}{g_i}$ for each consecutive eigenvector i (larger ratio means more steps). Intuitively, algorithms struggle to rank eigenvectors when they have similar eigenvalues (small eigengaps), especially in the stochastic setting where eigenvalues / eigenvectors are estimated from random samples of the data. In the large data setting where the entire dataset cannot be stored in memory, approaches that sample data (in this case batches of edge vectors x_e) are often all that are viable.

4 METHOD

In this section, we describe our approach to accelerating the convergence of solvers in finding the bottom- k eigenvectors of well-clustered graphs. We are particularly interested in constructing an approach that can scale to large graphs. As suggested in the previous section, we aim to improve convergence rates by dilating (enlarging) the eigengaps of the graph Laplacian relative to its spectral radius. If successful, this should reduce the number of steps required for convergence.

4.1 EIGENVECTOR PRESERVING TRANSFORMATIONS

Any monotonically increasing polynomial function of a matrix preserves eigenvectors and their rank. If we know we only want to discover eigenvectors with eigenvalues below some threshold λ_c , the function f actually only needs to be monotonic below this threshold; it can then be non-monotonic above λ_c as long as $f(x) > f(\lambda_c)$ for $x > \lambda_c$. For example, in the spectral clustering setting when using the normalized graph Laplacian, we might only be interested in eigenvectors with eigenvalues below $\lambda_c = 1$, i.e., clusters with fewer edges crossing between clusters than edges within a cluster.

Once we have a monotonically increasing function f that is tailored to dilate the eigengaps of the problem at hand, we will reverse the spectrum with equation 8 to compute the bottom- k eigenvectors. This is the approach we take in experiments in Section 5.

In order to dilate the eigengaps of small eigenvalues relative to those of large eigenvalues, we require nonlinear operators. Naively, one could leverage the eigendecomposition of L and compute, for example e^L as

$$e^L = e^{V\Lambda V^\top} = Ve^\Lambda V^\top = V\text{diag}(e^{\lambda_i})V^\top. \quad (10)$$

But obviously, this requires the eigendecomposition which we do not have and are trying to compute.

4.2 SERIES EXPANSIONS OF MATRIX TRANSFORMATIONS

Instead, we can explore power series approximations to nonlinear matrix operators. For example, the matrix exponential is given by

$$e^L = \sum_{i=0}^{\infty} \frac{L^i}{i!}. \quad (11)$$

This operation dilates eigengaps for large eigenvalues, but using a similar series we can compute $-e^{-L}$ which shrinks large eigenvalues relative to small ones. By shrinking large eigenvalues relative to smaller ones, the scale of the spectrum is reduced. Therefore, the ratio $\frac{\lambda|V_i|}{g_i}$ is smaller, and convergence is sped up. Note that the maximum eigenvalue under this transformation is < 0 (i.e., we can set $\lambda^* = 0$ in equation 8) and the spectral radius is 1.

4.3 SCALING: STOCHASTIC & PARALLEL

Recall that the incidence matrix X encodes edges on its rows and that the graph Laplacian can be written as $L = X^\top X = \sum_e x_e x_e^\top$. Therefore, any power ℓ of the graph Laplacian can be written as

$$\begin{aligned} L^\ell &= \left(\sum_{e_1 \in E} x_{e_1} x_{e_1}^\top \right) \dots \left(\sum_{e_\ell \in E} x_{e_\ell} x_{e_\ell}^\top \right) \\ &= \sum_{e_1 \in E} \dots \sum_{e_\ell \in E} (x_{e_1} x_{e_1}^\top) \dots (x_{e_\ell} x_{e_\ell}^\top) \\ &= \sum_{e_1 \in E} \dots \sum_{e_\ell \in E} \underbrace{\left[\prod_{j=1}^{\ell-1} x_{e_j}^\top x_{e_{j+1}} \right]}_{\alpha_c} x_{e_1} x_{e_\ell}^\top \\ &= \sum_{c_1 \in E^\ell} \alpha_c x_{e_1} x_{e_\ell}^\top \end{aligned} \quad (12)$$

where E^ℓ represents the Cartesian product of the edge set. Note that the scalar, denoted by α_c , is only non-zero if each of its constituent inner products $x_{e_j}^\top x_{e_{j+1}}$ is non-zero. And the inner product of two edge vectors is only non-zero if those two edges are incident. Table 1 lists the possible values for inner products of pairs of edges assuming x_e is encoded such that $x_{e_i} = 1$ and $x_{e_j} = -1$ if $i < j$ where $e = (i, j)$.

disconnected edges	$i \rightarrow j, k \rightarrow l$	0
serial	$i \rightarrow j \rightarrow l$	-1
converging	$i \rightarrow j \leftarrow l$	1
diverging	$i \leftarrow j \rightarrow l$	1
repeated	$i \Rightarrow j$	2

Table 1: Recall that edge directions are defined by node numbers, not by the direction of a walk on the graph.

In other words, α_c is only non-zero if the sequence of edges $e_1 \dots e_\ell$ forms a chain (hence the subscript c of α_c). Therefore, computing a power of the graph Laplacian is equivalent to computing a weighted sum of outer products over length- ℓ walks in the *edge incidence graph*².

This suggests a highly parallelizable method. Assuming access to d graph “walkers“, let each walker, in parallel, sample a node and conduct a random walk of length- ℓ on the graph. Average each walker’s calculation of the summand in equation 12 to estimate $\frac{1}{|\mathcal{E}|} L^\ell$. Note that if each walk is unbiased, then the average is unbiased.

In order to sample a walk uniformly at random, one could use rejection sampling. Pick an edge in the graph uniformly at random. Then perform a random walk on the edge incidence graph keeping track of the degree of every node visited in the edge incidence graph. Let $\deg(e_i)$ be the degree of node i in the edge incidence graph. Then the probability of a given length- ℓ walk is

$$p_\ell = \frac{1}{|\mathcal{E}|} \prod_{i=1}^{\ell} \frac{1}{\deg(e_i)}. \quad (13)$$

Also, let \deg^* be an upper bound on the maximum degree of any node in the original graph. Then $\deg_{inc}^* = 2\deg^* - 1$ is an upper bound on the degree of any node in the edge incidence graph. Let

$$p_{min} = \frac{(\deg_{inc}^*)^{-\ell}}{|\mathcal{E}|}. \quad (14)$$

Finally, if walks are rejected with probability $\frac{p_{min}}{p_\ell}$, then all walks in the edge incidence graph will occur with probability p_{min} ensuring uniform sampling.

Linearity of expectation says you can use same length- ℓ walk for unbiased estimates of all shorter walks because $\mathbb{E}[\sum_i \gamma_i L^i] = \sum_i \gamma_i \mathbb{E}[L^i]$. This means L^i and L^j can be correlated. Therefore, to generate a single unbiased estimate for all L^i with $1 \leq i \leq \ell$, simply continue performing random walks of length ℓ until at least one *subwalk* of each length less than or equal to ℓ has passed the rejection step. Note that with sufficient compute a large batch of random walks can be generated in parallel to reduce runtime. Furthermore, the generation of a single random walk can be parallelized as well (Lacki et al., 2020; Kapralov et al., 2021).

In general, a product of ℓ matrices can be computed in parallel via a binary tree in $\log(\ell)$ time but here we take advantage of the structure of L to compute it in a different way. We mention this as our approach can generalize beyond graph Laplacians given general knowledge of the spectrum

5 EXPERIMENTS

We compute the bottom- k eigenvectors of matrices with spectra that have several eigenvalues less than 1. This is meant to simulate conditions common in well separated graphs and clustering problems (Peng et al., 2015). The transformations we consider here are not robust to all graph types. Different classes of graphs will benefit from different transformations, but with prior knowledge of the problem setting and spectrum of the graph Laplacian, it should be possible to generalize our approach to other settings.

²By *edge incidence graph*, we mean to define a new undirected graph whose nodes represent the edges in the original graph and whose edges indicate whether two edges (now nodes) are incident. Note every node in the new graph also has a self-loop edge.

Matrix Logarithm (add $\epsilon \ll 1$)	$\begin{aligned} & \log(L + \epsilon) \\ & \sum_{i=1}^{\ell} (-1)^{i+1} \frac{(L+\epsilon I - I)^i}{i} \\ & -e^{-L} \\ & -\sum_{i=0}^{\ell} \frac{(-L)^i}{i!} \\ & -(I - L/\ell)^\ell \end{aligned}$
Taylor Series of $\log(L + \epsilon I)$	
Negative Decaying Exponential	
Taylor Series of $-e^{-L}$	
Limit Approximation of $-e^{-L}$	

Table 2: Transformation Functions (ℓ is odd)

5.1 MANIPULATING SPECTRA OF GRAPH LAPLACIANS

In this paper, we propose manipulating the spectrum of graph Laplacians in order to reduce the ratio $\frac{\lambda_{|V|}}{g_i}$ especially for the bottom- k eigenvalues. The goal is to transform the Laplacian in such a manner that larger eigengaps are shifted to the bottom of the eigenvalue spectrum. If we then only compute the bottom- k eigenvectors we can potentially reduce computation significantly, especially if we use cheap approximations of the transformations as discussed in Section 4.2.

In the following experiments, we compute both exact and series approximations to nonlinear transformations of the graph Laplacian and observe its effect on convergence of two representative, iterative SVD solvers: Oja’s algorithm (Shamir, 2015) and μ -EG (Gemp et al., 2021b).

5.2 EVALUATING CONVERGENCE

We consider two types of evaluations: normalized subspace error and longest eigenvector streak.

Subspace Error We use the definition of subspace error in (Tang, 2019; Gemp et al., 2021a):

$$\delta^t = 1 - \frac{1}{k} \text{tr}(U^* P_t) \tag{15}$$

where $U^* = V^* V^{*\top}$ is the ground truth orthogonal projector of the principal subspace and $P^t = V V^\top$ the orthogonal projector of the approximation of the subspace at time t .

Eigenvector Streak We use the definition in (Gemp et al., 2021a), where we measure the number of consecutive principal components that are within some ϵ of the ground truth principal components. This metric is harsher than subspace error and measures whether the actual bottom- k eigenvectors are being approximated. Below we will show results on a variety of problems. For all figures except Fig. 1, the horizontal axis is plotted on a logarithmic scale to make it more convenient to interpret convergence rate differences in terms of orders of magnitude.

5.3 PROTO-VALUE FUNCTIONS FOR MDPs

We consider computing the proto-value functions, basis functions for value functions of a tabular Markov Decision Process (MDP) illustrated in Figure 1. Specifically, we consider an MDP with 3 consecutive rooms with the middle connected to each of the outer rooms by small doors. The grid world is $10s + 1$ cells tall and $30s + 1$ cells wide. The doorways take up $\frac{1}{h}$ of the available vertical space ($\frac{10s+1}{h}$ cells tall).

The top- k proto-value functions are computed by first constructing the graph of all possible state transitions in the MDP (states s are nodes and undirected edges indicate possible transitions $s \rightarrow s'$; note $s \leftarrow s'$ is also included in the graph) and then computing the bottom- k eigenvectors of the corresponding graph Laplacian. Proto-value functions have been shown to be useful as a basis for learning value functions in reinforcement learning tasks (Mahadevan, 2005).

Figure 2 demonstrates that certain transformations of the graph Laplacian do, in fact, accelerate convergence to the bottom- k eigenvectors. Even when using a power series approximation, the number of steps can be reduced by roughly an order of magnitude ($10\times$). The exact matrix logarithm reduces the number of steps by over two orders of magnitude ($100\times$). However, it is known that the Taylor series approximation of the logarithm is only convergent for $\rho(L) < 2$, and we are unable to find a series approximation that is accurate enough over the graph Laplacian’s full spectrum.

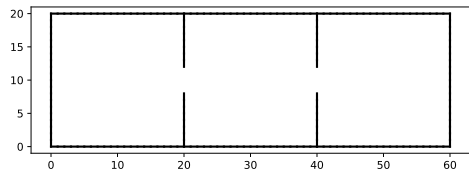


Figure 1: 3-Room MDP Grid World ($s = 2, h = 10$).

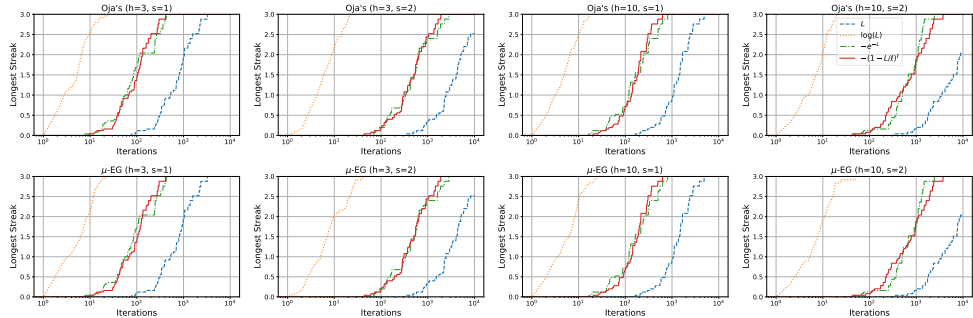


Figure 2: 3-Room MDP. Longest eigenvector streak (higher is better) is plotted over training for two different scalable SVD methods: μ -EG and Oja’s algorithm. Three nonlinear matrix transformations of matrix L are compared against the identity transformation. One transformation (in red) is a series approximation of degree $\ell = 251$ to the exact operation (in green).

Figure 3 reveals that both the nonlinear transformations and the series approximation accelerate the minimization of subspace error as well.

5.4 WELL-CLUSTERED GRAPHS

We generate graphs of n nodes split into k cliques. The k cliques are connected to each other by a random number between 0 and 25 of “short circuit” edges between the clusters.

Figure 4 reveals qualitatively similar results to Figure 2. The series approximation works in all cases in graphs with 1000 nodes and 5 or fewer clusters. However, we can see two settings where the series approximation to the exponential transformation fails to perform well. In both settings, the number of nodes in the graph is 2000 and the number of clusters is less than 5. With 5 clusters, the series approximation succeeds. Fewer clusters while keeping the total number of nodes constant implies the individual clusters have a much higher number of edges. In other words, the maximum degree in the graph blows up. The spectral radius, maximum eigenvalue of the graph Laplacian, is upper bounded by two times the max degree. Hence, with a higher max degree we expect a larger spectral radius. In order for a series approximation to perform similarly to its exact counterpart, the approximation must contain enough terms in its polynomial to be accurate over the range $[0, 2\text{deg}^*]$. Therefore, our hypothesis is that the series approximation requires more terms to be accurate enough to accelerate

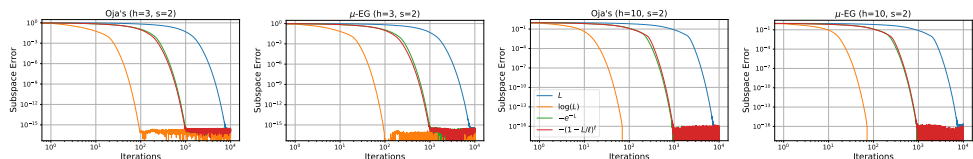


Figure 3: 3-Room MDP. Subspace error (lower is better) is plotted over training for two different scalable SVD methods: μ -EG and Oja’s algorithm. Three nonlinear matrix transformations of matrix L are compared against the identity transformation. One transformation (in red) is a series approximation of degree $\ell = 251$ to the exact operation (in green).

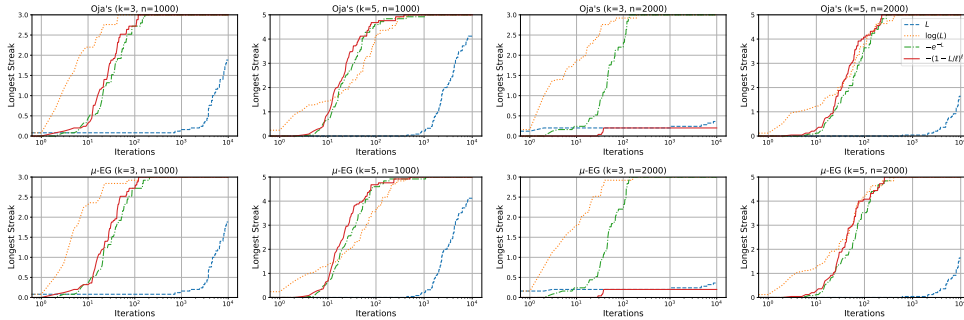


Figure 4: Cliques. Longest eigenvector streak (higher is better) is plotted over training for two different scalable SVD methods: μ -EG and Oja’s algorithm. Three nonlinear matrix transformations of matrix L are compared against the identity transformation. One transformation (in red) is a series approximation of degree $\ell = 251$ to the exact operation (in green).

convergence. We perform an additional experiment in Section A.2 of the appendix³ that supports this hypothesis.

6 CONCLUSION

In this work, we propose a general approach to dilating eigengaps of symmetric matrices ($X^\top X$) to accelerate convergence under a stochastic optimization model (samples of X are fed to the algorithm in minibatches). We discuss this approach in the context of eigendecomposition of the graph Laplacian, the SVD of a sparse matrix with a specific structure, and show how to parallelize polynomials that approximately dilate the spectrum of the matrix. We also show how this can be accomplished with samples in an unbiased way. We apply this approach to spectral-clustering and other related settings such as approximating proto-value functions and clustering a graph completed with link prediction.

In future work, we are interested in *meta-learning* the coefficients of the polynomials to accelerate convergence and improving upon the simple rejection sampling scheme for generating unbiased random walks. We also aim to run stochastic, parallelized experiments at scale on much larger graphs and explore variance reduction techniques.

ACKNOWLEDGEMENTS

We thank Daniel A. Spielman and Daniele Calandriello for sharing their expertise and pointing us to related work.

REFERENCES

Zeyuan Allen-Zhu and Yuanzhi Li. First efficient convergence for streaming k-PCA: a global, gap-free, and near-optimal rate. In *2017 IEEE 58th Annual Symposium on Foundations of Computer Science (FOCS)*, pp. 487–492. IEEE, 2017.

Maria-Florina Balcan, Simon Shaolei Du, Yining Wang, and Adams Wei Yu. An improved gap-dependency analysis of the noisy power method. In *Conference on Learning Theory*, pp. 284–309. PMLR, 2016.

Fan RK Chung and Fan Chung Graham. *Spectral graph theory*. Number 92. American Mathematical Soc., 1997.

Guy Barrett Coleman and Harry C Andrews. Image segmentation by clustering. *Proceedings of the IEEE*, 67(5):773–785, 1979.

³For the appendix, see the ArXiv version of this paper: <https://arxiv.org/abs/2207.14589>.

- Santo Fortunato. Community detection in graphs. *Physics reports*, 486(3-5):75–174, 2010.
- Ian Gemp, Brian McWilliams, Claire Vernade, and Thore Graepel. Eigengame: PCA as a Nash equilibrium. *International Conference on Learning Representations*, 2021a.
- Ian Gemp, Brian McWilliams, Claire Vernade, and Thore Graepel. Eigengame Unloaded: When playing games is better than optimizing. *arXiv preprint arXiv:2102.04152*, 2021b.
- Michael Kapralov, Silvio Lattanzi, Navid Nouri, and Jakab Tardos. Efficient and local parallel random walks. *Advances in Neural Information Processing Systems*, 34, 2021.
- Dheeraj Kumar and James C Bezdek. Visual approaches for exploratory data analysis: A survey of the visual assessment of clustering tendency (VAT) family of algorithms. *IEEE Systems, Man, and Cybernetics Magazine*, 6(2):10–48, 2020.
- Jakub Lacki, Slobodan Mitrović, Krzysztof Onak, and Piotr Sankowski. Walking randomly, massively, and efficiently. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, pp. 364–377, 2020.
- James R Lee, Shayan Oveis Gharan, and Luca Trevisan. Multiway spectral partitioning and higher-order Cheeger inequalities. *Journal of the ACM (JACM)*, 61(6):1–30, 2014.
- Sridhar Mahadevan. Proto-value functions: Developmental reinforcement learning. In *Proceedings of the 22nd international conference on Machine learning*, pp. 553–560, 2005.
- Richard Peng, He Sun, and Luca Zanetti. Partitioning well-clustered graphs: Spectral clustering works! In *Conference on Learning Theory*, pp. 1423–1455. PMLR, 2015.
- Ohad Shamir. A stochastic PCA and SVD algorithm with an exponential convergence rate. In *International Conference on Machine Learning*, pp. 144–152. PMLR, 2015.
- Yangqiu Song, Chenguang Wang, Ming Zhang, Hailong Sun, and Qiang Yang. Spectral label refinement for noisy and missing text labels. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*, 2015.
- Cheng Tang. Exponentially convergent stochastic k-PCA without variance reduction. *Advances in Neural Information Processing Systems*, 2019.
- Ulrike Von Luxburg. A tutorial on spectral clustering. *Statistics and computing*, 17(4):395–416, 2007.