

VARIANCE REDUCTION IN STOCHASTIC METHODS FOR LARGE-SCALE REGULARIZED LEAST-SQUARES PROBLEMS

1st Yusuf Yiğit Pilavcı
Univ. Grenoble Alpes
Grenoble INP, GIPSA-lab
Grenoble, France



2nd Pierre-Olivier Amblard
Univ. Grenoble Alpes, CNRS
Grenoble INP, GIPSA-lab
Grenoble, France



3rd Simon Barthelmé
Univ. Grenoble Alpes, CNRS
Grenoble INP, GIPSA-lab
Grenoble, France



4th Nicolas Tremblay
Univ. Grenoble Alpes, CNRS
Grenoble INP, GIPSA-lab
Grenoble, France



Abstract—Large dimensional least-squares and regularized least-squares problems are expensive to solve. There exist many approximate techniques, some deterministic (like conjugate gradient), some stochastic (like stochastic gradient descent). Among the latter, a new class of techniques uses Determinantal Point Processes (DPPs) to produce unbiased estimators of the solution. In particular, they can be used to perform Tikhonov regularization on graphs using random spanning forests, a specific DPP. While the unbiasedness of these algorithms is attractive, their variance can be high. We show here that variance can be reduced by combining the stochastic estimator with a deterministic gradient-descent step, while keeping the property of unbiasedness. We apply this technique to Tikhonov regularization on graphs, where the reduction in variance is found to be substantial at very small extra cost.

Index Terms—graph signal processing, smoothing, variance reduction, random spanning forests.

I. INTRODUCTION

In linear least-squares problems, given measurements $\mathbf{y} \in \mathbb{R}^n$ and predictors $\mathbf{A} \in \mathbb{R}^{n \times p}$, we seek the vector $\hat{\mathbf{x}}$ verifying:

$$\hat{\mathbf{x}} = \operatorname{argmin}_{\mathbf{x} \in \mathbb{R}^p} \|\mathbf{A}\mathbf{x} - \mathbf{y}\|_2^2 + \lambda \mathbf{x}^t \mathbf{P} \mathbf{x} \quad (1)$$

where $\lambda \mathbf{x}^t \mathbf{P} \mathbf{x}$ is a regularization term. $\hat{\mathbf{x}}$ can be computed exactly at cost $\mathcal{O}(np^2)$, but for very large n or p this is costly, and approximate methods may be used instead. There are dozens of such methods, either deterministic or stochastic in nature. The best known example of the former are the various gradient descent methods [1], while among the latter the most popular is certainly stochastic gradient descent [2]. An attractive alternative to stochastic gradient descent is to use methods based on Determinantal Point Processes, which, roughly, consist in finding a well-chosen, random subset of the rows of \mathbf{A} so that solving the least-squares problem for *just these rows* gives an *unbiased* estimator for the solution to the

This work was partly funded by the French National Research Agency in the framework of the "Investissements d'avenir" program (ANR-15-IDEX-02), the LabEx PERSYVAL (ANR-11-LABX-0025-01), the ANR GraVa (ANR-18-CE40-0005), the ANR GRANOLA (ANR-21-CE48-0009), the MIAI@Grenoble Alpes chairs "LargeDATA at UGA" and "Pollutants" (ANR-19-P3IA-0003).

full problem [3], [4]. This contrasts with stochastic gradient descent, where the estimator is generally biased after a finite number of steps.

One drawback of this new class of methods is that sampling a subset of rows with the right properties may be expensive, so that reducing the variance by simply increasing the number of estimates becomes rapidly too costly. We show below that a very simple way to reduce variance is to combine a stochastic estimator with (at least one) step of gradient descent. The result is another unbiased estimator, one that is guaranteed to have lower variance. We apply our technique to Tikhonov regularization on graphs, where DPP-based estimators are particularly well-adapted [5], [6]. As in standard applications of gradient descent, most of the difficulty consists in determining how large the gradient descent step should be. In graphs this problem turns out to be quite tractable, with good heuristics available. Numerical results confirm that the reduction in variance obtained is substantial, at small computational cost.

Main idea The main idea of the paper is extremely simple. Solving a least-squares problem of the form of Eq.(1) is equivalent to minimizing a quadratic form

$$f(\mathbf{x}) = \frac{1}{2} \mathbf{x}^t \mathbf{Q} \mathbf{x} - \mathbf{r}^t \mathbf{x}. \quad (2)$$

Minimizing f by gradient descent consists in taking steps of the form

$$\mathbf{x}_t = \mathbf{x}_{t-1} - \alpha \nabla f(\mathbf{x}_{t-1}) = \mathbf{x}_{t-1} - \alpha (\mathbf{Q} \mathbf{x}_{t-1} - \mathbf{r}),$$

where α is the step size. Since the gradient is zero at the solution, $\hat{\mathbf{x}} = \operatorname{argmin} f(\mathbf{x})$ is a fixed point of the iteration. Now suppose that $\bar{\mathbf{x}}$ is an unbiased estimator of $\hat{\mathbf{x}}$, *i.e.*, a random variable such that $\mathbb{E}(\bar{\mathbf{x}}) = \hat{\mathbf{x}}$. Then, it is easy to check that

$$\mathbb{E}(\bar{\mathbf{x}} - \alpha(\mathbf{Q}\bar{\mathbf{x}} - \mathbf{r})) = \hat{\mathbf{x}},$$

meaning that $\bar{\mathbf{x}}$ stays unbiased after one step of gradient descent (or indeed several). We show below that setting the step size α correctly guarantees a reduction in variance, and

give practical solutions for finding an appropriate step size in a graph signal processing setting.

II. BACKGROUND

Tikhonov regularization in graphs. The regularised least-squares estimator we are interested in is graph Tikhonov regularisation (GTR), a method for denoising graph signals. A “graph signal” is a vector of measurements associated with the nodes of a graph, for instance brain activity in n brain regions, where the graph models neural connectivity across regions. GTR also occurs as a subproblem in other methods, like semi-supervised learning [6], which is why finding an efficient approximation algorithm is of high interest.

Let us first set notation. We denote a graph by the set $\mathcal{G} = (\mathcal{V}, \mathcal{E}, w)$ with $|\mathcal{V}| = n$ vertices and $|\mathcal{E}| = m$ edges. $w : \mathcal{V} \times \mathcal{V} \mapsto \mathbb{R}^+$ is called the weight function and is non-zero only for $(i, j) \in \mathcal{E}$. In this work, we only consider connected (*i.e.*, there exists a path between any pair of nodes) and undirected graphs (*i.e.*, verifying $\forall (i, j) \in \mathcal{E}, w(i, j) = w(j, i)$). In addition, we use certain matrices to depict the algebraic properties of graphs. These are the (weighted) adjacency matrix $W = [w(i, j)]_{i, j} \in \mathbb{R}^{n \times n}$, the diagonal degree matrix whose diagonal entries are $D_{i, i} = \sum_{j=1}^n w(i, j)$ and the graph Laplacian, $L = D - W$. It is well-known that $L \in \mathbb{R}^{n \times n}$ is a symmetric positive semi-definite matrix with n eigenvalues $\lambda_1 = 0 < \lambda_2 \leq \dots \leq \lambda_n$ for any undirected, connected graph [7].

Given a noisy version $\mathbf{y} \in \mathbb{R}^n$ of an underlying graph signal $\mathbf{x} \in \mathbb{R}^n$ that one wishes to recover, GTR consists in solving the following minimization problem to estimate \mathbf{x} :

$$\hat{\mathbf{x}} = \underset{\mathbf{z} \in \mathbb{R}^n}{\operatorname{argmin}} q \|\mathbf{z} - \mathbf{y}\|_2^2 + \mathbf{z}^\top \mathbf{L} \mathbf{z}, \quad (3)$$

where $q \in \mathbb{R}^+$ is a parameter that adjusts the balance between the data fidelity term $\|\mathbf{z} - \mathbf{y}\|_2^2$ and the regularization term $\mathbf{z}^\top \mathbf{L} \mathbf{z}$, which forces the solution to be smooth over the graph. Note that this is a special case of eq. (1) with $\mathbf{A} = \mathbf{I}$, $\mathbf{P} = \mathbf{L}$ and $\lambda = 1/q$. The explicit solution to this problem reads:

$$\hat{\mathbf{x}} = \mathbf{K} \mathbf{y} \quad \text{with} \quad \mathbf{K} = (q\mathbf{I} + \mathbf{L})^{-1}$$

The direct computation of $\hat{\mathbf{x}}$ requires (a worst-case) $\mathcal{O}(n^3)$ elementary operations due to the inversion of $(q\mathbf{I} + \mathbf{L})$. As n increases, this computation becomes prohibitive. In the next section, we list state-of-the-art methods which avoid this expensive computation.

DPP-based estimators for Graph Tikhonov Regularization.

In the special case of GTR, the most popular methods are deterministic. For large n , state-of-the-art methods can be divided roughly into two groups, namely iterative methods (*e.g.* conjugate gradient [8]) and polynomial approximations (*e.g.* Chebyshev polynomials [9]). Both classes of methods run in linear time with the number of edges, m .

As an alternative stochastic method for GTR, we have proposed in previous works unbiased Monte Carlo estimators for approximating $\hat{\mathbf{x}}$ which also scale linearly with m [5], [6].

Our method can be viewed as a DPP-based subsampling of the rows of a matrix, but the case of GTR is particularly favorable. The appropriate DPP is the random forest process [10], which is easy to sample from, and its analytic properties enable variance reduction via conditional Monte Carlo. The resulting estimator, called $\bar{\mathbf{x}}$ in [5], [6], is unbiased ($\mathbb{E}(\bar{\mathbf{x}}) = \hat{\mathbf{x}}$), can be obtained in $\mathcal{O}(m)$ elementary operations, and takes a very simple form: based on the random forest process, it samples a random partitioning of the vertices (a random set of disjoint subsets of the vertices whose union gives \mathcal{V}), and then averages the signal in each part of the partition. We repeat this process N times and average to get a denoised signal. The details of the estimator can be found in [6], but are irrelevant for what we discuss here.

III. PROPOSED METHOD

In this section, we propose an improved estimator for estimating $\mathbf{K} \mathbf{y}$. As described in the introduction, we combine our unbiased estimator with a gradient descent step to reduce variance. This variance reduction technique can also be viewed as an instance of the control variate method [11] in the Monte Carlo literature.

A. The gradient descent update

Note that the solution to (3) also minimizes the cost function:

$$F(\mathbf{z}) = \frac{1}{2} \mathbf{z}^\top \mathbf{K}^{-1} \mathbf{z} - \mathbf{z}^\top \mathbf{y}.$$

As stated in the introduction, applying the gradient step to $\bar{\mathbf{x}}$ yields a new estimator:

$$\bar{\mathbf{z}} := \bar{\mathbf{x}} - \alpha (\mathbf{K}^{-1} \bar{\mathbf{x}} - \mathbf{y}). \quad (4)$$

Proposition 1. $\bar{\mathbf{z}}$ is an unbiased estimator for $\hat{\mathbf{x}}$. Moreover, assuming¹ $\operatorname{tr}(\operatorname{Cov}(\mathbf{K}^{-1} \bar{\mathbf{x}})) > 0$, the MSE of $\bar{\mathbf{z}}$ is a quadratic function of α which is minimized for:

$$\alpha^* = \frac{\operatorname{tr}(\operatorname{Cov}(\mathbf{K}^{-1} \bar{\mathbf{x}}, \bar{\mathbf{x}}))}{\operatorname{tr}(\operatorname{Cov}(\mathbf{K}^{-1} \bar{\mathbf{x}}))}. \quad (5)$$

Proof. Since $\bar{\mathbf{x}}$ is an unbiased estimator, the expectation of $\bar{\mathbf{z}}$ immediately reads $\mathbb{E}[\bar{\mathbf{z}}] = \mathbb{E}[\bar{\mathbf{x}}] - \alpha (\mathbf{K}^{-1} \mathbb{E}[\bar{\mathbf{x}}] - \mathbf{y}) = \hat{\mathbf{x}}$. Now, let us focus on the variance of this new estimator:

$$\begin{aligned} \operatorname{tr}(\operatorname{Cov}(\bar{\mathbf{z}})) &= \operatorname{tr}(\operatorname{Cov}(\bar{\mathbf{x}})) + \alpha^2 \operatorname{tr}(\operatorname{Cov}(\mathbf{K}^{-1} \bar{\mathbf{x}})) \\ &\quad - 2\alpha \operatorname{tr}(\operatorname{Cov}(\mathbf{K}^{-1} \bar{\mathbf{x}}, \bar{\mathbf{x}})). \end{aligned} \quad (6)$$

Eq. (6) is a quadratic function in α and is minimized at α^* given in Eq. (5). \square

Connection with control variates. The proposed method is a particular instance of the control variate method. This technique leverages an additional random variable, called control variate, with a known expectation to improve the Monte Carlo

¹ Cov denotes the covariance matrix for a single input. Looking closely at the definition of $\bar{\mathbf{x}}$ in [6] as well as its variance property, one can show that $\operatorname{tr}(\operatorname{Cov}(\mathbf{K}^{-1} \bar{\mathbf{x}})) = 0$ only arises when \mathbf{y} is a constant vector, in which case $\bar{\mathbf{x}}$ is always equal to \mathbf{y} whatever the sampled forest, such that by Eq. (4), $\bar{\mathbf{z}} = \bar{\mathbf{x}}$ and there is no point in using $\bar{\mathbf{z}}$.

performance. In our case, the control variate is $K^{-1}\bar{\mathbf{x}}$ whose expectation is the input signal \mathbf{y} .

Implementation. $\bar{\mathbf{z}}$ requires the computation of the control variate per sample in addition to $\bar{\mathbf{x}}$. The nontrivial part of this computation includes the matrix-vector product $L\bar{\mathbf{x}}$. However, we can avoid to repeat this product per sample. Instead, we can calculate the sample mean $\frac{1}{N}\sum_{i=1}^N \bar{\mathbf{x}}^{(i)}$, then do the product once to directly compute the sample mean for $\bar{\mathbf{z}}$.

B. How to select α ?

Unfortunately, calculating the optimal value α^* requires information that is not readily available. However, as the MSE (Eq. (6)) is a quadratic function in α minimized in α^* , we know that any choice of $\alpha \in (0, 2\alpha^*)$ will necessarily decrease the variance (see an illustration of this in Fig. 1). We now discuss several options to find such appropriate –yet not optimal– α . Our examination can be separated into two parts, each proposing a choice of α with different motivations. First, we look for a constant for α that is cheap to compute and ensures variance reduction. Second, we consider how to approximate α^* within the Monte Carlo simulations.

A first option. To ensure variance reduction, $\bar{\mathbf{z}}$ must verify:

$$\mathbb{E}[\|\bar{\mathbf{z}} - \hat{\mathbf{x}}\|_2] \leq \mathbb{E}[\|\bar{\mathbf{x}} - \hat{\mathbf{x}}\|_2].$$

A much more restrictive constraint can be written as:

$$\|\bar{\mathbf{z}} - \hat{\mathbf{x}}\|_2 \leq \|\bar{\mathbf{x}} - \hat{\mathbf{x}}\|_2. \quad (7)$$

This constraint is classical in the gradient descent literature, where it is used to find a safe range for α . This particular choice costs $\mathcal{O}(1)$ to calculate for a given graph.

Proposition 2. Let d_{max} be the maximum degree of the graph. Then, setting $\alpha = \frac{2q}{q+2d_{max}}$ ensures that:

$$\|\bar{\mathbf{z}} - \hat{\mathbf{x}}\|_2 \leq \|\bar{\mathbf{x}} - \hat{\mathbf{x}}\|_2.$$

Proof. The following must hold to satisfy Eq. (7)

$$\begin{aligned} \|\bar{\mathbf{x}} - \hat{\mathbf{x}} - \alpha(K^{-1}\bar{\mathbf{x}} - \mathbf{y})\|_2 &\leq \|\bar{\mathbf{x}} - \hat{\mathbf{x}}\|_2 \\ \|(I - \alpha K^{-1})(\bar{\mathbf{x}} - \hat{\mathbf{x}})\|_2 &\leq \|\bar{\mathbf{x}} - \hat{\mathbf{x}}\|_2. \end{aligned} \quad (8)$$

This inequality holds for all $\bar{\mathbf{x}}, \hat{\mathbf{x}}$ if and only if $(I - \alpha K^{-1})$ is a contraction mapping which means:

$$|\mu_i| \leq 1, \quad \forall i \in \{1, \dots, n\},$$

where μ_i 's are the eigenvalues of the matrix $(I - \alpha K^{-1})$. Rewriting this constraint in terms of the eigenvalues of the graph Laplacian, we have: $\forall i, |1 - \alpha \frac{\lambda_i}{q} - \alpha| \leq 1$, i.e.: $\forall i, 0 \leq \alpha \leq \frac{2q}{\lambda_i + q}$. The tightest upper bound is given by $\lambda_i = \lambda_n$ and recalling $\lambda_n \leq 2d_{max}$ [7] finishes the proof. \square

A second option. In this section, we approximate α^* from the samples. Let us define $\bar{\mathbf{y}} := K^{-1}\bar{\mathbf{x}}$. Then, Eq. (5) may be rewritten as $\alpha^* = \frac{\text{tr}(\text{Cov}(\bar{\mathbf{x}}, \bar{\mathbf{y}}))}{\text{tr}(\text{Cov}(\bar{\mathbf{y}}))}$. Using N samples $\bar{\mathbf{x}}^{(1)}, \dots, \bar{\mathbf{x}}^{(N)}$ and the corresponding $\bar{\mathbf{y}}^{(1)}, \dots, \bar{\mathbf{y}}^{(N)}$, we can empirically estimate α^* with the sample covariance matrices:

$$\hat{\alpha} = \frac{\text{tr}(\widehat{\text{Cov}}(\bar{\mathbf{x}}, \bar{\mathbf{y}}))}{\text{tr}(\widehat{\text{Cov}}(\bar{\mathbf{y}}))}. \quad (9)$$




Fig. 1: Empirical squared error of $\bar{\mathbf{x}}$ (orange horizontal line) and $\bar{\mathbf{z}}$ (blue parabola) w.r.t α on two graphs generated by random models. On the left is a random regular graph with $n = 1000$ and $m = 10000$. On the right is a Barabasi-Albert model with parameter $k = 10$ resulting in $n = 1000$ and $m = 9900$. The green (resp. red) vertical dashed line shows $\alpha = \frac{2q}{q+2d_{max}}$ (resp. the estimated $\hat{\alpha}$ from the samples). The blue dot represents the best possible variance reduction obtained for $\alpha = \alpha^*$. The number of Monte Carlo samples is set to $N = 10$ and the error results are averaged over 200 realizations. The signal is a random vector generated from $\mathcal{N}(\mathbf{0}, I)$.

The analysis in [12, Ch 8.9] shows that along with $\hat{\alpha}$, $\bar{\mathbf{z}}$ may yield a biased estimator. However, there are heuristics indicating that the bias diminishes with rate $\mathcal{O}(N^{-1})$ which is often negligible for large N w.r.t. the estimation error.

An empirical comparison of these two options for α is given in Fig. 1. In a graph with a regular degree distribution, fixing α to $\frac{2q}{q+2d_{max}}$ gives a performance close to optimum. However, in the case of a graph with a broad degree distribution, $\frac{2q}{q+2d_{max}}$ fails to provide a good approximation for α^* , and one needs to resort to $\hat{\alpha}$ in order to leverage the full variance reduction potential offered by this gradient descent step.

IV. EXPERIMENTS

We illustrate the improved estimator $\bar{\mathbf{z}}$ over real data sets and compare it with $\bar{\mathbf{x}}$.

Tikhonov denoising. The first data set consists of the trip records of taxis in New York City² over 260 zones. From their geographical positions, we build a $k = 5$ -nearest neighbour graph. The signal, in our case, is the median of the total amount charged to passengers at 260 drop-off zones. In the experiments, we add artificial Gaussian noise to the signal and reconstruct it via Tikhonov denoising and its estimates $\bar{\mathbf{x}}$ and $\bar{\mathbf{z}}$. Fig. 2 summarizes the results. In Fig. 2d, we can clearly see the improvement via $\bar{\mathbf{z}}$ w.r.t. $\bar{\mathbf{x}}$ in denoising performance. We observe that $\bar{\mathbf{z}}$ with a constant α slightly outperforms the case with $\hat{\alpha}$. This is probably due to the poor estimation of $\hat{\alpha}$ with $N = 2$.

Node classification. In the second illustration, we use $\bar{\mathbf{z}}$ for solving the semi-supervised node classification problem. The purpose is to classify each vertex by leveraging the underlying graph while the class information is available over only a few vertices. We define a labeling $\mathbf{Y} = [\mathbf{y}_1, \dots, \mathbf{y}_k] \in \mathbb{R}^{n \times k}$ for

²The dataset is available at <https://www1.nyc.gov/site/tlc/about/tlc-trip-record-data.page>

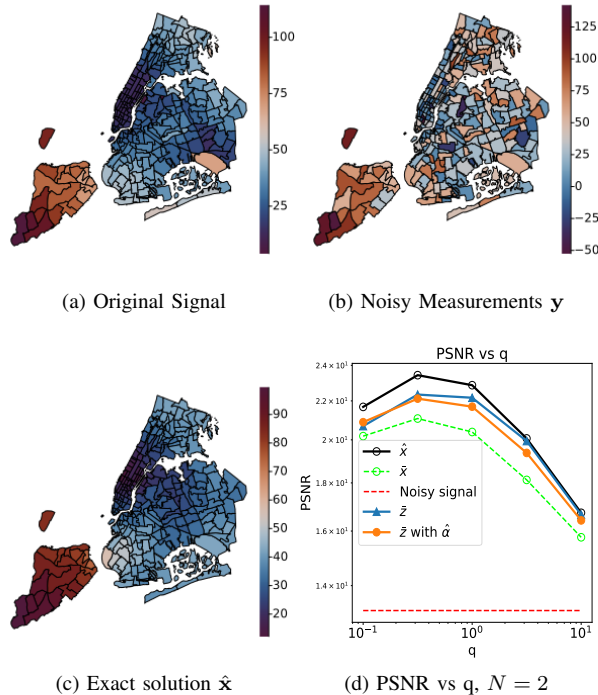


Fig. 2: Tikhonov denoising on New York City’s taxi network. A $k = 5$ - nn graph with $|\mathcal{V}| = 260$, $|\mathcal{E}| = 777$ is built based on the proximity of the zones. The original signal in (a) is the median of the total amount charged to passengers in drop-off locations. The noisy signal (b) is generated by adding noise drawn from a Gaussian distribution $\mathcal{N}(\mathbf{x}, 25)$. (c) is the exact solution $\mathbf{K}\mathbf{y}$ with $q = 0.32$. In (d), peak signal to noise ratio is plotted for \mathbf{y} , $\bar{\mathbf{x}}$, $\bar{\mathbf{z}}$ with $\alpha = \frac{2q}{q+2d_{max}}$ and $\bar{\mathbf{z}}$ with $\hat{\alpha}$ as in Eq. (9).

k classes where $\mathbf{y}_l(i)$ is 1 whenever node i is known to be in class j , otherwise it is 0. Given the set of labeled vertices ℓ , typically $|\ell| \ll |\mathcal{V}|$, the goal is to find a classification function $\mathbf{F} = [\mathbf{f}_1, \dots, \mathbf{f}_k] \in \mathbb{R}^{n \times k}$. One approach is generalized semi supervised learning framework by Avrachevko *et al.* [13] which calculates the following solution:

$$\mathbf{F} = \mathbf{D}^{1-\sigma} \mathbf{K} \mathbf{D}^{\sigma-1} \mathbf{Y} \quad \text{with } \mathbf{K} = \left(\mathbf{D} + \frac{2}{\mu} \mathbf{L}\right)^{-1} \mathbf{D} \quad (10)$$

where $\mu \in \mathbb{R}^+$ and $\sigma \in [0, 1]$ are the hyperparameters of the algorithm. The cumbersome operation in this solution occurs in the calculations involving \mathbf{K} due to the matrix inversion. In [5], we have already showed that $\bar{\mathbf{x}}$ can approximate these calculations with a certain setting:

$$\bar{\mathbf{z}} = \mathbf{D}^{1-\sigma} (\bar{\mathbf{x}} - \alpha (\mathbf{K}^{-1} \bar{\mathbf{x}} - \mathbf{y}_l)) \mathbf{D}^{\sigma-1}.$$

Adapting the calculations in Prop. 2, one obtains the safe option as $\alpha = \frac{2\mu}{\mu+4}$. In the experiments, we examine the classification accuracy of the exact solution and RSF estimates on the benchmark data sets Cora and Citeseer citation network³. Fig. 3 presents the classification accuracy given by $\hat{\mathbf{x}}$, $\bar{\mathbf{x}}$ and $\bar{\mathbf{z}}$. We observe that $\bar{\mathbf{z}}$ with the constant α performs quite similar




Fig. 3: Classification performance on the Cora and Citeseer datasets. m denotes the number of labeled vertices per class. In the experiments, we use $\mu = 1.0$, $\sigma = 0$ and $\alpha = \frac{2\mu}{\mu+4}$. The number of Monte Carlo samples is 50 and the experiments are repeated 100 times for each m .

to $\bar{\mathbf{z}}$ with $\hat{\alpha}$ while both outperforming $\bar{\mathbf{x}}$: for this particular example $\frac{2\mu}{\mu+4}$ gives a very good approximation for α^* .

V. CONCLUSION

The main idea in this work is that, given an unbiased estimator of the solution to a large-scale least-squares problem (such as in the recent lines of works based on DPPs [3], [4]), then one (or in fact several) step of classical deterministic gradient descent leaves the estimator unbiased while potentially significantly reducing its variance, for a very small extra cost.

We illustrate this idea on graph Tikhonov regularization via random spanning forests, where we show that the variance of the estimator may easily be divided by several factors if one chooses correctly the gradient descent step α . The simplest (and free to compute) choice for α is by far $2q/(q+2d_{max})$. Even though this choice is shown to necessarily decrease the variance, its performance is hindered when the degree distribution of the graph is broad and d_{max} becomes a crude upper bound of the distribution’s mode.

A natural workaround that we will explore in future work is to precondition the gradient step by \mathbf{D}^{-1} yielding (still unbiased) estimators of the form $\bar{\mathbf{z}} := \bar{\mathbf{x}} - \alpha \mathbf{D}^{-1} (\mathbf{K}^{-1} \bar{\mathbf{x}} - \mathbf{y})$.

³These datasets are available at <https://linqs.soc.ucsc.edu/data>

REFERENCES

- [1] Yurii Nesterov, *Introductory lectures on convex optimization: A basic course*, vol. 87, Springer Science & Business Media, 2003.
- [2] Léon Bottou, “Stochastic learning,” in *Advanced Lectures on Machine Learning*, Olivier Bousquet and Ulrike von Luxburg, Eds., Lecture Notes in Artificial Intelligence, LNAI 3176, pp. 146–168. Springer Verlag, Berlin, 2004.
- [3] Michał Dereziński and Michael W Mahoney, “Determinantal point processes in randomized numerical linear algebra,” *Notices of the American Mathematical Society*, vol. 68, no. 1, pp. 34–45, 2021.
- [4] Michaël Fanuel, Joachim Schreurs, and Johan Suykens, “Diversity sampling is an implicit regularization for kernel methods,” *SIAM Journal on Mathematics of Data Science*, vol. 3, no. 1, pp. 280–297, 2021.
- [5] Yusuf Y Yiğit Pilavcı, Pierre-Olivier Amblard, Simon Barthelme, and Nicolas Tremblay, “Smoothing graph signals via random spanning forests,” in *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*. IEEE, 2020, pp. 5630–5634.
- [6] Yusuf Yiğit Pilavcı, Pierre-Olivier Amblard, Simon Barthelme, and Nicolas Tremblay, “Graph tikhonov regularization and interpolation via random spanning forests,” *IEEE Transactions on Signal and Information Processing over Networks*, vol. 7, pp. 359–374, 2021.
- [7] Piet Van Mieghem, *Graph spectra for complex networks*, Cambridge University Press, 2010.
- [8] Yousef Saad, *Iterative methods for sparse linear systems*, SIAM, 2003.
- [9] David I Shuman, Pierre Vandergheynst, and Pascal Frossard, “Chebyshev polynomial approximation for distributed signal processing,” in *2011 International Conference on Distributed Computing in Sensor Systems and Workshops (DCOSS)*. IEEE, 2011, pp. 1–8.
- [10] Luca Avena and Alexandre Gaudillière, “Two applications of random spanning forests,” *Journal of Theoretical Probability*, vol. 31, no. 4, pp. 1975–2004, 2018.
- [11] Zdravko Botev and Ad Ridder, “Variance reduction,” *Wiley StatsRef: Statistics Reference Online*, pp. 1–6, 2017.
- [12] Art B. Owen, *Monte Carlo theory, methods and examples*, 2013.
- [13] Konstantin Avrachenkov, Alexey Mishenin, Paulo Gonçalves, and Marina Sokol, “Generalized optimization framework for graph-based semi-supervised learning,” in *Proceedings of the 2012 SIAM International Conference on Data Mining*. SIAM, 2012, pp. 966–974.