

ALTERNATING PROJECTIONS WITH VOLUME SAMPLING

Anonymous authors

Paper under double-blind review

ABSTRACT

The method of Alternating Projections (AP) is a fundamental iterative technique with applications to problems in machine learning, optimization and signal processing. Examples include the Gauss-Seidel algorithm which is used to solve large-scale regression problems and the Kaczmarz and projections onto convex sets (POCS) algorithms that are fundamental to iterative reconstruction. Progress has been made with regards to the questions of efficiency and rate of convergence in the randomized setting of the AP method. Here, we extend these results with volume sampling to block (batch) sizes greater than 1 and provide explicit formulas that relate the convergence rate bounds to the spectrum of the underlying system. These results, together with a trace formula and associated volume sampling, prove that convergence rates monotonically improve with larger block sizes, a feature that can not be guaranteed in general with uniform sampling (e.g., in SGD).

1 INTRODUCTION

Solving a system of linear equations is one of the most fundamental computational problems. It is core to many problems in optimization, machine learning and signal processing that require solutions to large-scale linear systems. Advances in computational efficiency of solving large-scale linear systems directly translate to more efficient algorithms for non-linear optimization (Wilson et al., 2021). Randomization of classical algorithms such as coordinate descent and gradient descent has been instrumental in solving large-scale optimization problems. Characterizing the convergence of these randomized algorithms for solving linear systems is key to optimizing their performance and developing acceleration techniques that are used for more general classes (e.g., strongly convex) of objective functions. Broadly speaking, these randomized descent algorithms are row-space and column-space methods for solving a consistent linear system $\mathbf{Ax} = \mathbf{b}$, that in the randomized setting can be viewed as a sequence of alternating projections onto certain subspaces (see Section 2) specified by \mathbf{A} . When a single column/row is selected at each iterate (e.g., coordinate descent or SGD) these are rank-1 projections onto hyperplanes, and more generally when a subset of columns/rows are chosen (e.g., block coordinate descent with size n) the resulting rank- n projections are onto subspaces with codimension n .

Due to their sequential stochastic nature, these algorithms lend themselves to a Markovian view that facilitates the use of ergodic theory of (continuous state space) Markov chains to determine convergence criteria as well as rates of convergence. Analyzed as a time-homogeneous Markov chain, the existence and uniqueness of the stationary measure concentrated at the solution of the linear system is guaranteed by a stochastic version of the contraction mapping theorem. Likewise, the rate of convergence can be bounded using the *spectral gap* of a certain operator (see Section 2) that is constructed, for each algorithm, from \mathbf{A} . While these results are well understood in Markov chain theory, establishing the relationship between the spectral gap—and hence the rate—to the spectrum of \mathbf{A} has been a fundamental theoretical challenge. Understanding this relationship is of practical importance since performance of competing iterative methods such as Krylov subspace methods (e.g., conjugate gradients) are well understood in terms of the spectrum of \mathbf{A} (Saad, 2003).

For the base case of rank-1 projections (e.g., coordinate descent or SGD), when rows/columns of \mathbf{A} are sampled, i.i.d., with probabilities according to their lengths (Strohmer & Vershynin, 2009; Leventhal & Lewis, 2010), the spectral gap is simply determined from the smallest singular value

of \mathbf{A} . We demonstrate that for rank- n projectors (e.g., block coordinate descent), the spectral gap is entirely determined from all of the singular values of \mathbf{A} when subsets are sampled, i.i.d., according to their volumes. Establishing the relationship between the spectral gap and the singular values of \mathbf{A} is a significant departure from standard results on (deterministic) block methods. In block methods the rate depends on the condition of the worst block in a partition of \mathbf{A} – a quantity that depends on the performance of the partitioning method and is decoupled from the spectrum of \mathbf{A} .

Our theory shows the exact process by which the spectral gap improves with n . We demonstrate that the singular values of \mathbf{A} , or equivalently eigenvalues of $\mathbf{A}^T \mathbf{A}$ are nonlinearly transformed with an attraction towards their mean as n increases. This evolution, we show, is the n^{th} step in a recursive formulation of the Cayley-Hamilton theorem known as the Faddeev–LeVerrier algorithm.

To sample with probabilities according to volumes when n is relatively small (as large as $n = 15$, as our experiments show) one can employ rejection sampling; for larger n efficient volume sampling is made possible by more sophisticated techniques that were developed by the pioneering work of (Deshpande & Rademacher, 2010; Deshpande et al., 2006).

2 RANDOMIZED DESCENT ALGORITHMS

Randomized algorithms for solving a linear system of equations $\mathbf{A}\mathbf{x} = \mathbf{b}$, by and large, descend on an objective function that guarantees almost sure convergence. The specifics of the objective function as well as the descent varies from technique to technique. In this section we show that the method of alternating projections presents a unifying perspective for the analysis of these techniques.

Given an $N \times N$ positive definite matrix \mathbf{A} , the randomized Gauss-Seidel algorithm updates the iterate, \mathbf{x}_k , a single coordinate at a time which is chosen at random. The objective function being minimized here is $f(\mathbf{x}) = \|\mathbf{A}^{1/2}(\mathbf{x} - \mathbf{x}_*)\|^2$, with \mathbf{x}_* being the solution to the linear system. A descent along a direction \mathbf{d} with exact line search updates the iterate \mathbf{x}_k according to (Leventhal & Lewis, 2010):

$$\mathbf{x}_{k+1} = \mathbf{x}_k + \frac{\langle \mathbf{d}, \mathbf{b} - \mathbf{A}\mathbf{x}_k \rangle}{\langle \mathbf{d}, \mathbf{A}\mathbf{d} \rangle} \mathbf{d} = \mathbf{x}_k + \mathbf{d} \left(\mathbf{d}^T \mathbf{A} \mathbf{d} \right)^{-1} \mathbf{d}^T (\mathbf{b} - \mathbf{A}\mathbf{x}_k). \quad (1)$$

Randomized Gauss-Seidel, chooses \mathbf{d} to be a randomly-chosen coordinate vector \mathbf{e}_n (i.e., all zeros except along the n^{th} coordinate axis of \mathbf{x}) resulting in a coordinate descent in iterations $\mathbf{x}_k \rightarrow \mathbf{x}_*$.

Similarly the randomized Kaczmarz algorithm considers the objective function $f(\mathbf{x}) = \|\mathbf{x} - \mathbf{x}_*\|^2$ for a consistent system of equations. Given a matrix $\mathbf{A} \in \mathbb{R}^{M \times N}$ with $M \geq N$ rows, the descent direction \mathbf{d} is chosen randomly from rows of \mathbf{A} . To facilitate accessing subsets of rows, we consider the collection of rows in \mathbf{A} as a set and denote by $\mathbf{a} \in \mathbf{A}$ a vector whose transpose, \mathbf{a}^T , is some row of \mathbf{A} , and by b_a the corresponding element of $\mathbf{b} \in \mathbb{R}^M$ on the right hand side of the system $\mathbf{A}\mathbf{x} = \mathbf{b}$. The one-step optimal descent along this direction is accomplished by:

$$\mathbf{x}_{k+1} = \mathbf{x}_k + \frac{b_a - \langle \mathbf{a}, \mathbf{x}_k \rangle}{\langle \mathbf{a}, \mathbf{a} \rangle} \mathbf{a}. \quad (2)$$

This results in a descent in iterations $\mathbf{x}_k \rightarrow \mathbf{x}_*$ that lies at the intersection of all hyperplanes represented by rows of \mathbf{A} .

Common to these descent algorithms is a notion of projection onto a subspace described by the underlying linear system $\mathbf{A}\mathbf{x} = \mathbf{b}$. For Gauss-Seidel these subspaces are the ones spanned by the rows $\mathbf{a} \in \mathbf{A}^{1/2}$ (with the n^{th} row corresponding to the choice of \mathbf{e}_n and $\|\mathbf{a}\|^2 = \mathbf{A}_{n,n}$) which is well defined for a symmetric positive definite matrix. In contrast, in Kaczmarz, the subspaces are simply spanned by rows $\mathbf{a} \in \mathbf{A}$. Let $\mathbf{P}_1^{\mathbf{a}} := \mathbf{a} \mathbf{a}^T / \|\mathbf{a}\|^2$ denote the rank-1 orthogonal projector onto the space spanned by \mathbf{a} . The dynamics of these iterations can be analyzed by the contraction they introduce in each step to an error vector \mathbf{z}_k by alternating projections:

$$\mathbf{z}_{k+1} = (\mathbf{I} - \mathbf{P}_1^{\mathbf{a}}) \mathbf{z}_k. \quad (3)$$

For the Gauss-Seidel method the notion of error vector is a residual measure $\mathbf{z}_k := \mathbf{A}^{1/2}(\mathbf{x}_k - \mathbf{x}_*)$ and in the Kaczmarz iterations the error is simply $\mathbf{z}_k := \mathbf{x}_k - \mathbf{x}_*$. When subspace projections $\mathbf{P}_1^{\mathbf{a}}$ are chosen at random, as discussed in the next section, the randomized iterations introduce a contraction in the error vector \mathbf{z}_k . This contraction can be quantified by analyzing the expected norm of the error vector, $\mathbb{E}\|\mathbf{z}_k\|^2$, that can be used to bound the rate of convergence.

2.1 CONVERGENCE RATE

Existence and uniqueness of the invariant measure for the Markov chain—the point mass at the solution of the linear system $\mathbf{A}\mathbf{x} = \mathbf{b}$ —can be demonstrated under very general conditions (Hairer, 2018) for these randomized alternating projection algorithms. It is easy to show that randomized alternating projections are Feller Markov processes, which upon the application of the Krylov-Bogoliubov theorem guarantees the existence of an invariant measure. Conversely, *deterministic contraction*, a property of orthogonal projections, guarantees the uniqueness of this invariant measure (see (Hairer, 2018) for details).

Given a probability distribution for randomly (and independently) choosing a coordinate \mathbf{e}_n in Gauss-Seidel, or choosing a row $\mathbf{a} \in \mathbf{A}$ in Kaczmarz, the geometric rate of convergence is dependent on the *mean* of the projections \mathbf{P}_1^α that this distribution engenders. In particular, a lower bound on the rate of convergence is obtained from the spectral radius of the average contraction in (3): $\lambda_{\max} \mathbb{E}[\mathbf{I} - \mathbf{P}_1^\alpha] = \lambda_{\max}(\mathbf{I} - \mathbb{E}[\mathbf{P}_1^\alpha]) = 1 - \lambda_{\min}(\mathbb{E}[\mathbf{P}_1^\alpha])$. The key quantity in this bound is the *spectral gap* (i.e., the smallest eigenvalue) of the operator representing the average of all the noted rank-1 projectors:

$$\tau_1 := \lambda_{\min}(\mathbb{E}[\mathbf{P}_1^\alpha]). \quad (4)$$

To elaborate, given a starting point \mathbf{x}_0 , the randomized Gauss-Seidel’s residual decays geometrically as: $\mathbb{E}\|\mathbf{r}_k\|^2 \leq (1 - \tau_1)^k \|\mathbf{r}_0\|^2$ and Kaczmarz’s error as $\mathbb{E}\|\mathbf{x}_k - \mathbf{x}_*\|^2 \leq (1 - \tau_1)^k \|\mathbf{x}_0 - \mathbf{x}_*\|^2$.

A *fundamental* challenge is to relate this spectral gap, τ_1 , to quantities that can be measured and computed from \mathbf{A} . While in the deterministic realm of alternating projections this quantity depends on a series of subspace angles that are difficult to compute (Galántai, 2005; Deutsch, 1995; Nelson & Neumann, 1987), in the randomized regime, the spectral gap (and hence the rate) is simply determined from the smallest singular value of \mathbf{A} if the probability distribution is set to select a projector \mathbf{P}_1^α with a probability proportional to length of \mathbf{a} squared: $pr(\mathbf{a}) \sim \|\mathbf{a}\|^2$. These probabilities are simply the diagonal elements of \mathbf{A} in Gauss-Seidel and lengths of rows of \mathbf{A} in Kaczmarz. This was first observed for the randomized Kaczmarz algorithm in (Strohmer & Vershynin, 2009) where $\mathbb{E}[\mathbf{P}_1] = \mathbf{A}^T \mathbf{A} / \text{Tr}(\mathbf{A} \mathbf{A}^T)$ and $\tau_1 = \lambda_{\min}(\mathbf{A}^T \mathbf{A}) / \text{Tr}(\mathbf{A} \mathbf{A}^T)$ and then leveraged in the randomized Gauss-Seidel (coordinate descent) case (Leventhal & Lewis, 2010) where $\mathbb{E}[\mathbf{P}_1] = \mathbf{A} / \text{Tr} \mathbf{A}$ and $\tau_1 = \lambda_{\min}(\mathbf{A}) / \text{Tr} \mathbf{A}$.

2.2 VOLUME SAMPLING

To facilitate faster convergence in practical applications, there has been a long line of research investigating block methods for iterative solvers such as block Gauss-Seidel and block Kaczmarz (Saad, 2003; Elfving, 1980) that have also been explored in randomized settings (Needell & Tropp, 2014; Liu & Wright, 2016; Gower & Richtárik, 2015; Tu et al., 2017). The critical difficulty in block methods is that the convergence rate depends on the worst condition number among all blocks in a given partition of \mathbf{A} – a quantity that is disconnected from the spectrum of \mathbf{A} .

Generalizing (1) to coordinate descent on a block of n coordinates, a matrix \mathbf{D} containing n coordinate vectors replaces \mathbf{d} , giving:

$$\mathbf{x}_{k+1} = \mathbf{x}_k + \mathbf{D} (\mathbf{D}^T \mathbf{A} \mathbf{D})^{-1} \mathbf{D}^T (\mathbf{b} - \mathbf{A} \mathbf{x}_k). \quad (5)$$

$\mathbf{D}^T \mathbf{A} \mathbf{D}$ selects a principal minor of \mathbf{A} that is chosen according to the descent coordinates. The subspaces for projections in this case are spanned by subsets of size n chosen from rows of $\mathbf{A}^{1/2}$, that we denote by $\mathbf{A}_n \subset \mathbf{A}^{1/2}$ with $\mathbf{A}_n \mathbf{A}_n^T = \mathbf{D}^T \mathbf{A} \mathbf{D}$ (for Kaczmarz $\mathbf{A}_n \subset \mathbf{A}$). The rank- n projectors that introduce contraction in each step of (3) are given by $\mathbf{P}_n := \mathbf{A}_n^T (\mathbf{A}_n \mathbf{A}_n^T)^{-1} \mathbf{A}_n$. Then the convergence rate is similarly bounded by $1 - \tau_n$ with the spectral gap of the expected projector:

$$\tau_n := \lambda_{\min}(\mathbb{E}[\mathbf{P}_n]). \quad (6)$$

We establish that the spectral gap is determined from an evolution of singular values of \mathbf{A} towards their mean when subsets are chosen with probabilities proportional to the square of the volumes they subtend. Specifically the expected projector $\mathbb{E}[\mathbf{P}_n] = \sum_{\mathbf{A}_n} pr(\mathbf{A}_n) \mathbf{P}_n$ is formed with probabilities according to their volumes: $pr(\mathbf{A}_n) \sim \det(\mathbf{A}_n \mathbf{A}_n^T)$ which is simply the determinant of the corresponding minor of \mathbf{A} . As we will see the normalization constant that is the sum of all squared

volumes $\text{vol}_n := \sum_{\mathbf{A}_n} \det(\mathbf{A}_n \mathbf{A}_n^T)$, can be calculated efficiently for any n using a trace formula despite the combinatorial nature of all subset of size n . This means that for small n (in fact, as large as $n = 15$, as our experiments demonstrate) one can employ simple rejection sampling techniques for volume sampling and for large n more sophisticated Markov chain sampling techniques (Deshpande & Rademacher, 2010; Deshpande et al., 2006) provide efficient volume sampling.

3 RESULTS

We demonstrate that under volume sampling the spectrum of $\mathbb{E}[\mathbf{P}_n]$ as n increases evolves from the spectrum of $\mathbb{E}[\mathbf{P}_1] = \mathbf{A}/N$ towards its mean (see Fig. 1 for an example). We prove that this evolution is described recursively by the Faddeev-LeVerrier algorithm traditionally used for computing coefficients of the characteristic polynomial of \mathbf{A} . For the Gauss-Seidel method with a block size n , the spectrum of \mathbf{A} is transformed according to the spectrum of the matrix Φ_n , with $\Phi_1 := \mathbf{A}$, defined recursively for $n > 1$ as:

$$\Phi_n = \Phi_1 \left(\frac{\text{Tr } \Phi_{n-1}}{n-1} \mathbf{I} - \Phi_{n-1} \right). \quad (7)$$

We show that $\mathbb{E}[\mathbf{P}_n] = \Phi_n / \text{Tr } \Phi_n$. This shows that the spectral gap $\tau_n = \lambda_{\min}(\Phi_n) / \text{Tr } \Phi_n$ is a polynomial of degree n over the spectrum of \mathbf{A} . For $n = 1$ the results of (Leventhal & Lewis, 2010) follows. For $n = 2$ this implies the spectrum of \mathbf{A} is transformed by the quadratic polynomial $\Phi_2(x) = (\text{Tr } \mathbf{A})x - x^2$. Specifically the spectral gap $\tau_2 = \lambda_{\min}(\Phi_2) / \text{Tr } \Phi_2$ is the smallest eigenvalue after this quadratic transformation and $\text{Tr } \Phi_2$ is the sum of those transformed eigenvalues. More generally the spectrum of \mathbf{A} is transformed according to the degree- n polynomial:

$$\Phi_n(x) := \sum_{p=1}^n (-1)^{p-1} \text{vol}_{n-p} x^p, \quad (8)$$

where $\text{vol}_n := \sum_{\mathbf{A}_n} \det \mathbf{A}_n \mathbf{A}_n^T$ is the sum of squared volumes of all subsets of size n that constitutes the normalization factor discussed in Section 2.2. We show that for any n , the trace formula $\text{vol}_n = \text{Tr } \Phi_n$ provides an efficient computation of the normalization factor that avoids the combinatorially large computation over all subset of size n . This process is the same for Kaczmarz except for the starting point: $\Phi_1 := \mathbf{A}^T \mathbf{A}$ which coincides with (Strohmer & Vershynin, 2009) for $n = 1$ and provides a similar evolution of eigenvalues of $\mathbf{A}^T \mathbf{A}$ for block sizes $n > 1$.

The evolution of spectrum in (7) introduces an attraction to their mean as they progress towards equalization as $\Phi_N = \mathbf{I}$ according to the Cayley-Hamilton theorem. To illustrate the nature of this evolution, as n increases, Fig. 1 shows the process on two examples of spectra with a linear decay and an exponential decay.

4 PROJECTIONS WITH VOLUME SAMPLING

As discussed in Section 3 when subsets of n rows are selected $\mathbf{A}_n \subset \mathbf{A}^{1/2}$ according to their volumes, the expected projector $\mathbb{E}[\mathbf{P}_n] = \Phi_n / \text{Tr } \Phi_n$ evolves in a recursive fashion described by (7). In this section we prove this result by establishing a combinatorial analysis of the set of projectors corresponding to all subsets \mathbf{A}_n together with a recursive expansion for rank- n orthogonal projectors that to the best of our knowledge is new.

4.1 RECURSIVE EXPANSION OF PROJECTORS

Lemma 1. *Let U denote the span of n linearly independent vectors $\{\mathbf{a}_1, \dots, \mathbf{a}_n\}$. For each $1 \leq s \leq n$, let $\mathbf{P}_1^s := \mathbf{a}_s \mathbf{a}_s^T / \|\mathbf{a}_s\|^2$ be the orthogonal projector into the subspace spanned by \mathbf{a}_s , and $\mathbf{P}_{n-1}^{\bar{s}}$ be the orthogonal projector into the subspace $U_{\bar{s}}$ of U , spanned by all but \mathbf{a}_s . Furthermore, let the angle between \mathbf{a}_s and $U_{\bar{s}}$ be denoted by θ_s . Then the orthogonal projector into U has the expansion:*

$$\mathbf{P}_n = \sum_{s=1}^n \frac{1}{\sin^2 \theta_s} \mathbf{P}_1^s (\mathbf{I} - \mathbf{P}_{n-1}^{\bar{s}}). \quad (9)$$

216
217
218
219
220
221
222
223
224
225
226
227
228
229
230
231
232
233
234
235
236
237
238
239
240
241
242
243
244
245
246
247
248
249
250
251
252
253
254
255
256
257
258
259
260
261
262
263
264
265
266
267
268
269




Figure 1: Evolution of spectrum towards its mean with exemplar linear and exponential decay with $N = 8$ according to volume sampling with (7).

Proof. Let \mathbf{L} denote the expansion on the right hand side of (9). We prove $\mathbf{P}_n = \mathbf{L}$. First, we note that for any $t \neq s$, $\mathbf{P}_{n-1}^s \mathbf{a}_t = \mathbf{a}_t$, which implies that $(\mathbf{I} - \mathbf{P}_{n-1}^s) \mathbf{a}_t = \mathbf{0}$. We decompose $\mathbf{a}_s = \mathbf{a}_s^\parallel + \mathbf{a}_s^\perp$ into its components in the subspace U_s and orthogonal to that subspace: $\mathbf{a}_s^\parallel = \mathbf{P}_{n-1}^s \mathbf{a}_s$ and $\mathbf{a}_s^\perp = (\mathbf{I} - \mathbf{P}_{n-1}^s) \mathbf{a}_s$. The length of the orthogonal component is therefore $\|\mathbf{a}_s^\perp\| = \|\mathbf{a}_s\| \sin \theta_s$.

We now use these observations to show that

$$\mathbf{L} \mathbf{a}_t = \mathbf{a}_t \quad 1 \leq t \leq n.$$

Based on the above, the only term in the summation corresponding to $\mathbf{L} \mathbf{a}_t$ that is non-zero is when $s = t$, so that

$$\mathbf{L} \mathbf{a}_t = \frac{1}{\sin^2 \theta_t} \mathbf{P}_1^t \mathbf{a}_t^\perp = \frac{\langle \mathbf{a}_t, \mathbf{a}_t^\perp \rangle}{\sin^2 \theta_t \|\mathbf{a}_t\|^2} \mathbf{a}_t.$$

Now $\langle \mathbf{a}_t, \mathbf{a}_t^\perp \rangle = \|\mathbf{a}_t\| \|\mathbf{a}_t^\perp\| \cos(\angle(\mathbf{a}_t^\perp, \mathbf{a}_t))$, and since $\angle(\mathbf{a}_t^\perp, \mathbf{a}_t)$ is complementary to $\theta_t = \angle(\mathbf{a}_t^\parallel, \mathbf{a}_t)$, we get $\langle \mathbf{a}_t, \mathbf{a}_t^\perp \rangle = \|\mathbf{a}_t\| \|\mathbf{a}_t^\perp\| \sin \theta_t$. Plugging in $\|\mathbf{a}_t^\perp\| = \|\mathbf{a}_t\| \sin \theta_t$ then gives the result.

For any vector \mathbf{z} orthogonal to U , we have: $\mathbf{P}_{n-1}^s \mathbf{z} = \mathbf{0}$ and $\mathbf{P}_1^s \mathbf{z} = \mathbf{0}$ since \mathbf{z} is orthogonal to all the n vectors $\mathbf{a}_1, \dots, \mathbf{a}_n$; therefore, $\mathbf{L} \mathbf{z} = \mathbf{0}$.

This shows $\mathbf{L} = \mathbf{P}_n$ is the orthogonal projector into U , which is unique. \square

We expand on the notion of orthogonal projectors and introduce a quasi projector that is well-defined even when the vectors are not linearly independent:

$$\mathbf{Q}_n := \mathbf{A}_n^T \text{adj}(\mathbf{G}_n) \mathbf{A}_n. \quad (10)$$

The adjugate matrix, $\text{adj}(\mathbf{G}_n)$, is also the cofactor matrix of $\mathbf{G}_n := \mathbf{A}_n \mathbf{A}_n^T$ due to its symmetry. When the vectors are linearly dependent $\mathbf{Q}_n = \mathbf{0}$ (as we will see), otherwise the quasi projector is a scaled version of the orthogonal projector:

$$\mathbf{Q}_n = v_n^2 \mathbf{P}_n \quad \text{where} \quad v_n^2 := \det \mathbf{G}_n. \quad (11)$$

Here v_n^2 represents the (square of) n -volume of the parallelepiped formed by the rows chosen in \mathbf{A}_n . Using this volume definition, we can present a corollary to Theorem 1:

Corollary 1 (Recursive Quasi Projector). *Under the assumptions of the lemma, let $v_1^s := \|\mathbf{a}_s\|$ and $v_{n-1}^{\bar{s}}$ denote the volume of the parallelepiped formed by all but \mathbf{a}_s . Moreover, let $\mathbf{Q}_1^s = \mathbf{a}_s \mathbf{a}_s^T$ and $\mathbf{Q}_{n-1}^{\bar{s}}$ denote the quasi projectors corresponding to \mathbf{P}_1^s and $\mathbf{P}_{n-1}^{\bar{s}}$, respectively. Then we have:*

$$\mathbf{Q}_n = \sum_{s=1}^n \mathbf{Q}_1^s \left((v_{n-1}^{\bar{s}})^2 \mathbf{I} - \mathbf{Q}_{n-1}^{\bar{s}} \right). \quad (12)$$

Proof. The volume of the parallelepiped can be computed from the volume of any facet, $v_{n-1}^{\bar{s}}$, and the corresponding height, $v_1^s \sin \theta_s$: $v_n = v_{n-1}^{\bar{s}} v_1^s \sin \theta_s$. \square

Lemma 2. *For a linearly dependent set of vectors, $\{\mathbf{a}_1, \dots, \mathbf{a}_n\}$, the quasi projector $\mathbf{Q}_n = \mathbf{0}$, the zero matrix.*

Proof. Based on (10), $\mathbf{Q}_n^T \mathbf{Q}_n = \mathbf{0}$ since $\mathbf{G}_n \text{adj}(\mathbf{G}_n) = (\det \mathbf{G}_n) \mathbf{I} = \mathbf{0}$. This implies that all columns of \mathbf{Q}_n have zero norm. Hence $\mathbf{Q}_n = \mathbf{0}$. On the other hand, since Corollary 1 is under the linearly independent assumption, we show, using a continuity argument, that for a linearly dependent set Corollary 1 still holds.

The quasi projector, defined by (10) $\mathbf{Q}_n := \mathbf{A}_n^T \text{adj}(\mathbf{G}_n) \mathbf{A}_n$, has elements that are each a continuous function of vectors $\mathbf{a}_1, \dots, \mathbf{a}_n$ that constitute rows of \mathbf{A}_n . This follows from the elements of the adjugate matrix (minors) being signed volumes of subsets of \mathbf{a} 's. Volume (determinant) is a continuous function of its set of vectors. Corollary 1 can therefore be expanded to include linearly dependent sets since any linearly dependent set of $n \leq N$ vectors in \mathbb{R}^N can be perturbed to become a linearly independent set. \square

4.2 COMBINATORIAL ANALYSIS

We can now go back and perform an expectation analysis for $\mathbb{E}[\mathbf{P}_n]$ over all subset \mathbf{A}_n of size n . As a reminder, for Gauss-Seidel with a symmetric positive definite $\mathbf{A} \in \mathbb{R}^{N \times N}$ there are the $\binom{N}{n}$ subsets of rows of $\mathbf{A}^{1/2}$ with each $\mathbf{A}_n \mathbf{A}_n^T$ being a principal minor of \mathbf{A} . In case of Kaczmarz with $\mathbf{A} \in \mathbb{R}^{M \times N}$ and $M \geq N$, there are $\binom{M}{n}$ subsets \mathbf{A}_n that are simply the subsets of rows. The combinatorial arguments are identical for Gauss-Seidel and Kaczmarz and we present the more general argument based on M rows that in the case of Gauss-Seidel simplifies to $M = N$.

We introduce a choice function that indexes these possible choices: $(i) \mapsto \{\mathbf{a}_1, \dots, \mathbf{a}_n\}$ denotes the set of n rows corresponding to the i^{th} choice, $1 \leq i \leq \binom{M}{n}$. For example, $v_n(i)$ denotes the volume of the parallelepiped formed by the vectors from the rows selected for the i^{th} choice, and likewise, $\mathbf{Q}_n(i)$ and $\mathbf{P}_n(i)$ denote the corresponding quasi projector and orthogonal projector.

We establish $\mathbb{E}[\mathbf{P}_n]$, as a polynomial in the matrix \mathbf{A} , when expectation is taken according to volume probability for the i^{th} choice set proportional to $v_n^2(i)$. Recalling the definition of vol_n , sum of squared volumes $\text{vol}_n = \sum_i v_n^2(i)$, we have: $pr(i) = v_n^2(i)/\text{vol}_n$. Then the average projector is $\mathbb{E}[\mathbf{P}_n] = \sum_i pr(i) \mathbf{P}_n(i) = 1/\text{vol}_n \sum_i \mathbf{Q}_n(i)$. To state our key result, we define a total quasi projector for a matrix \mathbf{A} :

$$\Phi_n := \sum_{i=1}^{\binom{M}{n}} \mathbf{Q}_n(i).$$

We now fully characterize Φ_n , specifically its spectrum, in our main result. When $n = 1$ the total quasi projector is $\Phi_1 = \mathbf{A} = \sum_i \mathbf{Q}_1(i) = \sum_{\mathbf{a} \in \mathbf{A}^{1/2}} \mathbf{a} \mathbf{a}^T$ for Gauss-Seidel (and for Kaczmarz the Gram matrix: $\Phi_1 = \mathbf{A}^T \mathbf{A} = \sum_i \mathbf{Q}_1(i) = \sum_{\mathbf{a} \in \mathbf{A}} \mathbf{a} \mathbf{a}^T$). For a larger set of rows $n > 1$ we show that Φ_n is a degree- n polynomial of the Gram matrix.

Theorem 1 (Total Quasi Projector). *For $n > 1$ rows we have:*

$$\Phi_n = \Phi_1 (\text{vol}_{n-1} \mathbf{I} - \Phi_{n-1}) \quad (13)$$

where $\text{vol}_n = \text{vol}_n(\mathbf{A}) := \sum_{i=1}^{\binom{M}{n}} v_n^2(i)$.

Proof. Before arguing over the $\binom{M}{n}$ choices, we first expand the set of choices and define an operator that sums quasi projectors over all n -ordered choices of M rows with replacement for a total of M^n choices:

$$\tilde{\Phi}_n := \sum_{j=1}^{M^n} \mathbf{Q}_n(j) = \sum_{\mathbf{a}_1 \in \mathbf{A}} \sum_{\mathbf{a}_2 \in \mathbf{A}} \cdots \sum_{\mathbf{a}_n \in \mathbf{A}} \mathbf{Q}_n.$$

We denote the sum of volumes in the expanded setting as $\tilde{\text{vol}}_n := \sum_{j=1}^{M^n} v_n^2(j)$. The expanded choices allow for summation over individual vectors that can sift through the recursion in (12). For example, the first term in (12) for $s = 1$ shows:

$$\begin{aligned} & \sum_{\mathbf{a}_1 \in \mathbf{A}} \sum_{\mathbf{a}_2 \in \mathbf{A}} \cdots \sum_{\mathbf{a}_n \in \mathbf{A}} \mathbf{Q}_1^1 \left((v_{n-1}^{\bar{1}})^2 \mathbf{I} - \mathbf{Q}_{n-1}^{\bar{1}} \right) \\ &= \sum_{\mathbf{a}_1 \in \mathbf{A}} \mathbf{Q}_1^1 \sum_{\mathbf{a}_2, \dots, \mathbf{a}_n \in \mathbf{A}} \left((v_{n-1}^{\bar{1}})^2 \mathbf{I} - \mathbf{Q}_{n-1}^{\bar{1}} \right) \\ &= \sum_{\mathbf{a} \in \mathbf{A}} \mathbf{Q}_1 \sum_{\mathbf{a}_1, \dots, \mathbf{a}_{n-1} \in \mathbf{A}} \left((v_{n-1})^2 \mathbf{I} - \mathbf{Q}_{n-1} \right) \\ &= \Phi_1 \left(\tilde{\text{vol}}_{n-1} \mathbf{I} - \tilde{\Phi}_{n-1} \right). \end{aligned}$$

Observing that the result of this summation is independent of s , allows us to establish:

$$\tilde{\Phi}_n = n \Phi_1 \left(\tilde{\text{vol}}_{n-1} \mathbf{I} - \tilde{\Phi}_{n-1} \right).$$

Now we observe that in a particular choice of n -rows with replacement, if any row of \mathbf{A} is selected more than once $v_n^2(j) = 0$ and $\mathbf{Q}_n = \mathbf{0}$ for any n . This means we can shrink the space of choices to n -permutations *without* replacement, with a total of ${}^M P_n := M!/(M-n)!$ choices, and still obtain the same $\tilde{\Phi}_n$:

$$\sum_{j=1}{{}^M P_n} \mathbf{Q}_n(j) = \tilde{\Phi}_n = n \Phi_1 \left(\tilde{\text{vol}}_{n-1} \mathbf{I} - \tilde{\Phi}_{n-1} \right).$$

To further shrink the space of choices to n -combinations, we note that permuting the order of the rows in a particular choice does not change the [squared volume](#) of the parallelepiped they form. This means $\text{vol}_n = \tilde{\text{vol}}_n/n!$ for any n . Moreover, permuting the rows in a particular choice does not change the orthogonal projector \mathbf{P}_n and consequently \mathbf{Q}_n . This means $\Phi_n = \tilde{\Phi}_n/n!$ for any n :

$$\begin{aligned} \Phi_n &= \frac{n}{n!} \Phi_1 \left((n-1)! \text{vol}_{n-1} \mathbf{I} - \tilde{\Phi}_{n-1} \right) = \frac{1}{(n-1)!} \Phi_1 \left((n-1)! \text{vol}_{n-1} \mathbf{I} - (n-1)! \tilde{\Phi}_{n-1} \right) \\ &= \Phi_1 \left(\text{vol}_{n-1} \mathbf{I} - \Phi_{n-1} \right). \end{aligned}$$

□

As a consequence, unwinding the recursion, we have:

Corollary 2 (Polynomial Form).

$$\Phi_n = \sum_{p=1}^n (-1)^{p-1} \text{vol}_{n-p} \Phi_1^p. \quad (14)$$

Since singular values of \mathbf{A} , when squared, match the eigenvalues of Φ_1 , this establishes (8).

An important consequence of the theorem is that the volume measures, vol_n associated with a set of vectors $\mathbf{a} \in \mathbf{A}$, can be recursively computed efficiently, which is of independent interest (McMullen, 1984; Dyer et al., 1998; Gover & Krikorian, 2010):

Theorem 2 (Volume Computation of all n -subsets). *Given a set of M vectors, arranged in rows of \mathbf{A} , the sum of squared volumes of parallelepipeds formed by size- n subsets is:*

$$\text{vol}_n = \sum_{i=1}^{\binom{M}{n}} v_n^2(i) = \frac{\text{Tr } \Phi_n}{n}.$$

The trace formula avoids the combinatorially-large computation over all subsets.

378 *Proof.* Using the fact that $\text{Tr } \mathbf{P}_n = \text{rank}(\mathbf{P}_n) = n$, and linearity of trace, we have $\text{Tr } \mathbf{Q}_n = nv_n^2$.
 379 Summing over all choices gives us:

$$381 \text{Tr } \Phi_n = \sum_{i=1}^{\binom{M}{n}} \text{Tr } \mathbf{Q}_n(i) = n \sum_{i=1}^{\binom{M}{n}} v_n^2(i) = n \text{vol}_n.$$

□


386 This shows the recursion in (13) can be written entirely in terms of Φ_1 :

$$388 \Phi_n = \Phi_1 \left(\frac{\text{Tr } \Phi_{n-1}}{n-1} \mathbf{I} - \Phi_{n-1} \right).$$

390 This provides a recursive form of a Cayley-Hamilton expansion for Φ_n , in terms of powers of \mathbf{A} for
 391 Gauss-Seidel (and powers of $\mathbf{A}^T \mathbf{A}$ for Kaczmarz) and their traces known as the Faddeev–LeVerrier
 392 algorithm.

394 5 EXPERIMENTS

395 We ran multiple instances of kernel ridge regression problems solved using randomized Gauss-
 396 Seidel for various values of block size n . We found that the convergence rates in each instance was
 397 patterned according to our theoretical predictions. We therefore present results from a prototypical
 398 numerical experiment.



424 Figure 2: Convergence results for randomized Gauss-Seidel with block sizes of $n = 1, 10, \text{ and } 15$.
 425 Dashed lines correspond to the rate bounds predicted by theory. The case $n = 1$ corresponds to the
 426 results of (Leventhal & Lewis, 2010). Each faint curve corresponds to a different run of the iterative
 427 algorithm. The solid lines correspond to ensemble averages across these trials. 30 trials were used
 428 to compute the ensemble average.

429 The details of the experimental setup was as follows. We generated multiple datasets $\{\langle \mathbf{a}_i, b_i \rangle\}_{i=1}^N$
 430 where the dimensionality of \mathbf{a}_i was set to be either 10, 25, or 100. Datasets ranged in size from
 431 $N = 25$ to 40. We limited the size of the dataset to a range for which a simple rejection sampler

could be used to sample rows according to their subtended volume. This was done to validate the theoretical results of the paper without the influence of any confounding effects introduced by a more sophisticated volume sampler (Deshpande & Rademacher, 2010) that while efficient, is influenced by burn-in and variance of stationary distribution effects. A shift invariant Gaussian kernel was then applied to the data to create one, two, or three clusters: $K_{ij} = \exp(-\gamma\|\mathbf{a}_i - \mathbf{a}_j + C_{i,j}\|_2^2)$, where $C_{i,j}$ was set to 0 if \mathbf{a}_i and \mathbf{a}_j belonged to the same cluster, and was set to randomly chosen and preset quantities for every pairing of disparate cluster memberships. This allowed us to experiment with varying levels of block dominance in \mathbf{K} and concomitant effects on its spectrum. We then solved the linear system $(\mathbf{K} + \lambda\mathbf{I})\mathbf{x} = \mathbf{b}$ using randomized Gauss-Seidel where at each iteration $n = 1, 10$ or 15 coordinates were simultaneously updated. The tuning parameters γ, λ were set to various fixed values. Note that $(\mathbf{K} + \lambda\mathbf{I})$ is always symmetric positive definite for appropriately chosen λ and therefore randomized Gauss-Seidel converges.

In each experiment 30 independent trials of randomized Gauss-Seidel were run using i.i.d. volume sampled rows of block size 1, 10 and 15. The modified residual $\mathbf{r}_k = (\mathbf{K} + \lambda\mathbf{I})^{1/2}(\mathbf{x}_k - \mathbf{x}_*)$ was computed at each iteration and $\|\mathbf{r}_k\|^2$ was recorded. These residuals were then averaged across the trials to create an ensemble average, and was plotted in addition to the specific randomized trial of Gauss-Seidel. Also was plotted the rate bound as predicted by our theory. The results are presented in Figure 2.

The main implication of our results is establishing the explicit relationship between the rate of convergence in randomized AP (e.g., Gauss-Seidel) for solving $\mathbf{A}\mathbf{x} = \mathbf{b}$ to the spectrum of \mathbf{A} . As discussed in Section 2, $1 - \tau_n$ provides a bound on the rate of convergence and the spectral gap τ_n is simply the spectral gap of $\mathbb{E}[\mathbf{P}_n]$ that is related to the spectrum of \mathbf{A} through (7). This recursive process stipulates the evolution of the spectrum of $\mathbb{E}[\mathbf{P}_n]$ as n increases from 1 where the spectrum is that of $\mathbb{E}[\mathbf{P}_1] = \mathbf{A}/N$, towards its mean. We present an empirical validation of these theoretical results on the evolution of the spectrum of \mathbf{A} , and as a result show how the spectral gap τ_n increases as block size n grows, in Figure 3. Note in particular that due to the recursive nature of the formula 8, the spectral gap can be computed very efficiently for problems of very large sizes. An appropriate choice of n can then be made based on balancing the larger computational cost of volume sampling n rows vis-a-vis the rate gains provided by a larger n .

6 CONCLUSION

This paper generalizes the results of (Strohmer & Vershynin, 2009; Leventhal & Lewis, 2010) that establish the relationship between the performance of randomized Gauss-Seidel and Kaczmarz algorithms to the spectrum of \mathbf{A} when single coordinates or rows (i.e., $n = 1$) are sampled according to their lengths. We establish that when $n > 1$ coordinates (or rows) are selected according to their volumes, the spectral gap that bounds the convergence rate is similarly determined from the spectrum of \mathbf{A} that is nonlinearly transformed according to the Faddeev-LaVerrier algorithm. We derive a recursive formulation of this evolution of spectrum towards its mean and establish efficient volume computation results that avoids combinatorially large computations over subsets. These results establish the convergence of the method of alternating projections under volume sampling for which efficient algorithms have been developed in theoretical computer science.

7 REPRODUCIBILITY STATEMENT

The essential result of this paper is the evolution of the spectrum of a matrix \mathbf{A} under the Faddeev-Laverrier algorithm is described recursively in (7) and in explicit polynomial form given in (8). Given the spectrum of \mathbf{A} its evolution at any n can be efficiently computed using the recursive Φ_n and the volume computation formula $\text{vol}_n = \text{Tr } \Phi_n$ that is established in Theorem 3. Hence the evolution process in Fig. 1 and Fig. 3 are reproducible. The convergence results in Fig. 2 are implementing (5) with randomness implemented by a simple rejection sampler. This experiment is also reproducible with minimal effort.

486
487
488
489
490
491
492
493
494
495
496
497
498
499
500
501
502
503
504
505
506
507
508
509
510
511
512
513
514
515
516
517
518
519
520
521
522
523
524
525
526
527
528
529
530
531
532
533
534
535
536
537
538
539




Figure 3: Log-scale view of spectrum of $\mathbb{E}[\mathbf{P}_1] = \mathbf{A}/N$ ($n = 1$), to $\mathbb{E}[\mathbf{P}_n]$ for $n = 10$ and $n = 15$ and the corresponding spectral gap, τ_n , from the experiment in the previous figure and its evolution according to (7).

REFERENCES

- Amit Deshpande and Luis Rademacher. Efficient volume sampling for row/column subset selection. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pp. 329–338, 2010.
- Amit Deshpande, Luis Rademacher, Santosh S Vempala, and Grant Wang. Matrix approximation and projective clustering via volume sampling. *Theory of Computing*, 2(1):225–247, 2006.
- Frank Deutsch. The angle between subspaces of a hilbert space. In *Approximation theory, wavelets and applications*, pp. 107–130. Springer, 1995.
- Martin Dyer, Peter Gritzmann, and Alexander Hufnagel. On the complexity of computing mixed volumes. *SIAM Journal on Computing*, 27(2):356–400, 1998.
- Tommy Elfving. Block-iterative methods for consistent and inconsistent linear equations. *Numerische Mathematik*, 35:1–12, 1980.
- Aurél Galántai. On the rate of convergence of the alternating projection method in finite dimensional spaces. *Journal of mathematical analysis and applications*, 310(1):30–44, 2005.
- Eugene Gover and Nishan Krikorian. Determinants and the volumes of parallelotopes and zonotopes. *Linear Algebra and its Applications*, 433(1):28–40, 2010.
- Robert M. Gower and Peter Richtárik. Randomized iterative methods for linear systems. *SIAM Journal on Matrix Analysis and Applications*, 36(4):1660–1690, 2015.
- Martin Hairer. Ergodic properties of markov processes. 2018. URL <https://www.hairer.org/notes/Markov.pdf>.
- Dennis Leventhal and Adrian S Lewis. Randomized methods for linear constraints: convergence rates and conditioning. *Mathematics of Operations Research*, 35(3):641–654, 2010.

540 Ji Liu and Stephen Wright. An accelerated randomized kaczmarz algorithm. *Mathematics of Com-*
541 *putation*, 85(297):153–178, 2016.

542 Peter McMullen. Volumes of projections of unit cubes. *Bulletin of the London Mathematical Society*,
543 16(3):278–280, 1984.

544 Deanna Needell and Joel A Tropp. Paved with good intentions: analysis of a randomized block
545 kaczmarz method. *Linear Algebra and its Applications*, 441:199–221, 2014.

546 Stuart Nelson and Michael Neumann. Generalizations of the projection method with applications
547 to sor theory for hermitian positive semidefinite linear systems. *Numerische Mathematik*, 51:
548 123–141, 1987.

549 Yousef Saad. *Iterative methods for sparse linear systems*. SIAM, 2003.

550 Thomas Strohmer and Roman Vershynin. A randomized kaczmarz algorithm with exponential con-
551 vergence. *Journal of Fourier Analysis and Applications*, 15(2):262–278, 2009.

552 Stephen Tu, Shivaram Venkataraman, Ashia C. Wilson, Alex Gittens, Michael I. Jordan, and Ben-
553 jamin Recht. Breaking locality accelerates block gauss-seidel. In *Proceedings of the 34th Inter-*
554 *national Conference on Machine Learning - Volume 70, ICML’17*, pp. 3482–3491. JMLR.org,
555 2017.

556 Ashia C Wilson, Ben Recht, and Michael I Jordan. A Lyapunov analysis of accelerated methods in
557 optimization. *Journal of Machine Learning Research*, 22(113):1–34, 2021.

558
559
560
561
562
563
564
565
566
567
568
569
570
571
572
573
574
575
576
577
578
579
580
581
582
583
584
585
586
587
588
589
590
591
592
593