RANDOMIZED KACZMARZ ALGORITHMS: EXACT MSE ANALYSIS AND OPTIMAL SAMPLING PROBABILITIES

Ameya Agaskar[1,2], Chuang Wang[3,4] and Yue M. Lu[1]

[1]Harvard University, Cambridge, MA 02138, USA
[2]MIT Lincoln Laboratory, Lexington, MA 02420, USA
[3] Institute of Theoretical Physics, Chinese Academy of Sciences, Beijing 100190, China
[4] University of the Chinese Academy of Sciences, Beijing 100049, China
Email: {aagaskar, yuelu}@seas.harvard.edu, wangchuang@itp.ac.cn

ABSTRACT

The Kaczmarz method, or the algebraic reconstruction technique (ART), is a popular method for solving large-scale overdetermined systems of equations. Recently, Strohmer *et al.* proposed the randomized Kaczmarz algorithm, an improvement that guarantees exponential convergence to the solution. This has spurred much interest in the algorithm and its extensions. We provide in this paper an exact formula for the mean squared error (MSE) in the value reconstructed by the algorithm. We also compute the exponential decay rate of the MSE, which we call the "annealed" error exponent. We show that the typical performance of the algorithm is far better than the average performance. We define the "quenched" error exponent to characterize the typical performance. This is far harder to compute than the annealed error exponent, but we provide an approximation that matches empirical results. We also explore optimizing the algorithm's row-selection probabilities to speed up the algorithm's convergence.

Index Terms— Overdetermined linear systems, Kaczmarz Algorithm, randomized Kaczmarz algorithm

1. INTRODUCTION

The Kaczmarz algorithm [1], also known under the name Algebraic Reconstruction Technique (ART) [2], is a popular method for solving a large-scale overdetermined system of linear equations. Let

$$y = Ax, \qquad (1)$$

where A is a full-rank $m \times n$ matrix with $m \geq n$. Given $y \in \mathbb{R}^m$, the algorithm proceeds to solve for x as follows: An initial guess $x^{(0)}$ is chosen arbitrarily. The iterations then start with the first row, proceed in succession to the last row, and then cycle back to the first row, and so on. When row r is chosen, the current estimate $x^{(k)}$ is projected onto the hyperplane $\{x \in \mathbb{R}^n : a_r^T x = y_r\}$ to obtain $x^{(k+1)}$. Here, a_r^T is the rth row of A.

Due to its simplicity, the Kaczmarz algorithm has been widely used in signal and image processing. It is also a special case of the

projection onto convex sets (POCS) algorithm [3] for finding an intersection of many convex sets: in our case, we are looking for the intersection of a set of $(n - 1)$-dimensional hyperplanes in $\mathbb{R}^n$.

It is well-known that the rate of convergence of the original Kaczmarz algorithm depends heavily on the exact ordering of the rows in A [4]. Recognizing this issue, Strohmer and Vershynin proposed in [5] a randomized Kaczmarz algorithm (RKA) that, instead of cycling sequentially through the rows in a deterministic fashion, chooses a row *at random* at each step. In their paper, they analyzed a specific probability distribution: choosing row i with probability proportional to its squared norm $\|a_i\|^2$. They then showed the following upper bound on the mean squared error (MSE) of the RKA:

$$\mathbb{E}\|x^{(N)} - x\|^2 \leq \left(1 - \kappa_A^{-2}\right)^N \|x^{(0)} - x\|^2, \qquad (2)$$

where $\kappa_A \stackrel{\text{def}}{=} \|A\|_F \|A^{-1}\|_2$ is the scaled condition number of A, and A^{-1} is its left-inverse. Since $\kappa_A \geq \sqrt{n}$, the above bound guarantees that the MSE decays *exponentially* as the RKA iterations proceed.

The work of Strohmer and Vershynin spurred a great deal of interest in RKA and its various extensions (see, *e.g.*, [6]–[12]). The original analysis in [5] assumes that the linear inverse problem is consistent (*i.e.*, noise-free). The noisy case was studied in [7]. A more general algorithm, involving random projections onto blocks of rows, was analyzed in [10]. Recently, Zouzias and Freris [9] proposed a randomized extended Kaczmarz algorithm which converges to the least squares estimate of an inconsistent system of linear equations.

We provide three contributions in this paper:

1. *An exact MSE formula:* All previous works on analyzing the performance of RKA provide strict upper bounds on the MSE. In this paper, we present an *exact* closed-form formula for the MSE of RKA after N iterations, for any N. Due to space constraint, we only present the noise-free case in this paper. However, the technique we use can be extended to more general settings as studied in [7, 9, 10].

2. *Annealed and quenched error exponents:* We provide an exact formula for the annealed error exponent, which measures the asymptotic rate of decay of the MSE, and we provide a good approximation for the quenched error exponent, which measures the asymptotic rate of decay of the squared error during a typical realization of the algorithm.

3. *Optimal sampling probabilities:* Our exact MSE formula allows us to pose a simple semidefinite programming (SDP) problem, the solution of which leads to optimal row-selection probabilities to minimize the MSE of the RKA.

Randomized Kaczmarz Algorithm [5]

Input: $A \in \mathbb{R}^{m \times n}$ with rows $a_1^T, a_2^T, \ldots, a_m^T$; $y \in \mathbb{R}^m$; selection probabilities $p_1, \ldots, p_m$ with $\sum_i p_i = 1$; iteration count N.
Output: $\widehat{x} \in \mathbb{R}^n$, an estimate for $x \in \mathbb{R}^n$ solving $y = Ax$.
   Initialize $x^{(0)}$ arbitrarily.
   for $k = 1$ to N **do**
      $r \leftarrow i$ with probability p_i.
      $x^{(k)} \leftarrow x^{(k-1)} + \frac{y_r - a_r^T x^{(k-1)}}{\|a_r\|^2} a_r$
   end for
   $\widehat{x} \leftarrow x^{(N)}$.

---

2. EXACT PERFORMANCE ANALYSIS

2.1. Overview of RKA

Given a matrix $A \in \mathbb{R}^{m \times n}$ and vector $y \in \mathbb{R}^m$, the randomized Kaczmarz algorithm attempts to find a solution $x \in \mathbb{R}^n$ to (1) as follows[1] The iterand $x^{(0)} \in \mathbb{R}^n$ is initialized arbitrarily. At each step k, a row r_k is chosen at random. The probability of choosing row i is p_i; the row-selection probabilities $p_1, \ldots, p_m$ are tunable parameters of the algorithm. The iterand is then updated according to the formula

$$x^{(k)} = x^{(k-1)} + \frac{y_{r_k} - a_{r_k}^T x^{(k-1)}}{\|a_{r_k}\|^2} a_{r_k}. \tag{3}$$

The algorithm is listed above. The intuition behind the algorithm is simple. Each row of A and its corresponding entry in y defines a hyperplane on which the solution x must lie; at each time step in the RKA algorithm we randomly select one of these hyperplanes and project the iterand onto it, getting closer to the true solution with each step.

2.2. An Exact MSE Formula

Originally, Strohmer *et al.* proposed a specific probability distribution: $p_i = \frac{\|a_i\|^2}{\|A\|_F^2}$, where $\|\cdot\|_F$ is the Frobenius norm, and analyzed the behavior of the algorithm in terms of the propeties of A. However, the solution to (1) is invariant to arbitrary and independent scalings of the rows. Thus, by looking at the properties of a rescaled version of A, their analysis can be applied to arbitrary row-selection probabilities. Indeed, their results show that

$$\left(1 - 2N/\kappa_A(p)^2\right) \leq \frac{\mathbb{E}\|x_N - x\|^2}{\|x_0 - x\|^2} \leq \left(1 - \kappa_A(p)^{-2}\right)^N, \tag{4}$$

where $\kappa_A(p) = \left\| \widetilde{A}^{-1} D_p^{-1/2} \right\|$, and we have defined $\widetilde{A}$ as the row-normalized version of A, and D_p as the diagonal matrix with $p_1, p_2, \ldots, p_m$ on the diagonal. $\widetilde{A}^{-1}$ is the left-inverse, which is guaranteed to exist because A is a tall, full-rank matrix.

The upper bound in (2) is sufficient to show that the error decays exponentially as the RKA iterations proceed. However, we show that it is possible to compute the *exact* error after N iterations of RKA, for any $N \geq 1$, given the initial error. This will allow us to precisely characterize the rate of decay of the error.

Proposition 1. *After N iterations of the randomized Kaczmarz algorithm with initial iterand $x^{(0)}$, the average error is given by*

$$\mathbb{E}\left\| x^{(N)} - x \right\|^2 \tag{5}$$
$$= \text{vec}(I_n)^T R_A(p)^N \text{vec}\left(\left(x^{(0)} - x\right)\left(x^{(0)} - x\right)^T \right),$$

[1]The extension of the analysis in this paper to the complex case is simple, but complicates the notation enough that we analyze only the real case here.

where $\text{vec}(\cdot)$ *is the vectorization operator stacking a matrix's columns into a vector, and we have defined*

$$R_A(p) = \sum_{i=1}^m p_i \left(P_{a_i}^\perp \otimes P_{a_i}^\perp \right). \tag{6}$$

Here, $P_{a_i}^\perp \overset{def}{=} \left(I_n - \frac{a_i a_i^T}{\|a_i\|^2} \right)$ *is the orthogonal projection onto the hyperplane orthogonal to* a_i, *and* $\otimes$ *is the Kronecker matrix product, so* $R_A(p)$ *is an* $n^2 \times n^2$ *matrix.*

Proof. To simplify the expressions, we define $z^{(k)} = x^{(k)} - x$ as the error vector after k iterations of the algorithm, with $z^{(0)}$ being the error in the initial guess. Then at each iteration, the error updates according to

$$z^{(k)} = z^{(k-1)} - \frac{a_{r_k}^T z^{(k-1)}}{\|a_{r_k}\|^2} a_{r_k} = P_{a_{r_k}}^\perp z^{(k-1)}, \tag{7}$$

where the r_k are i.i.d. indices chosen according to the probabilities $p_1, \ldots, p_m$.

To simplify the notation, we define $Q_k \overset{def}{=} P_{a_{r_k}}^\perp$. The error after N steps is then related to the initial error as

$$z^{(N)} = Q_N Q_{N-1} \cdots Q_1 z^{(0)}, \tag{8}$$

where $Q_1, \ldots, Q_N$ are i.i.d. random matrices. In particular, $Q_k = P_{a_i}^\perp = \left(I - \frac{a_i a_i^T}{\|a_i\|^2} \right)$ with probability p_i. The MSE after N steps can be computed as follows

$$\begin{aligned} \mathbb{E}\left\| z^{(N)} \right\|^2 &= \mathbb{E}\left\| Q_N Q_{N-1} \cdots Q_1 z^{(0)} \right\|^2 \\ &= \mathbb{E}\, z^{(0)T} Q_1 Q_2 \cdots Q_N I Q_N \cdots Q_2 Q_1 z^{(0)} \\ &= \mathbb{E}\, \text{trace}\left(Q_1 Q_2 \cdots Q_N I Q_N \cdots Q_2 Q_1 z^{(0)} z^{(0)T} \right) \\ &= \left[\mathbb{E}\, \text{vec}\left(Q_1 Q_2 \cdots Q_N I Q_N \cdots Q_2 Q_1 \right) \right]^T \\ &\qquad \text{vec}\left(z^{(0)} z^{(0)T} \right), \end{aligned} \tag{9}$$

where we have used two elementary matrix identities: $\text{trace}(AB) = \text{trace}(BA)$ and $\text{trace}(A^T B) = \text{vec}(A)^T \text{vec}(B)$ for any matrices A and B. We also used the fact that $Q_k = Q_k^T$. The expectation of the product of matrices can be explicitly computed as follows:

$$\begin{aligned} &\mathbb{E}\, \text{vec}\left(Q_1 Q_2 \cdots Q_N I Q_N \cdots Q_2 Q_1 \right) \\ &= \mathbb{E}(Q_1 \otimes Q_1)\, \text{vec}\left(Q_2 Q_3 \cdots Q_N I Q_N \cdots Q_3 Q_2 \right) \tag{10} \\ &= \mathbb{E}(Q_1 \otimes Q_1)\, \mathbb{E}\, \text{vec}\left(Q_2 Q_3 \cdots Q_N I Q_N \cdots Q_3 Q_2 \right) \tag{11} \\ &= \left[\mathbb{E} Q_1 \otimes Q_1 \right]^N \text{vec}(I), \tag{12} \end{aligned}$$

where (10) is due to the identity $\text{vec}(ABC^T) = (C \otimes A)\,\text{vec}(B)$ and the fact that Q_1 is symmetric, (11) is due to the independence of the random matrices, and (12) is the result of repeating the preceding two steps N times. Combining (9) and (12) completes the proof. $\square$

Remark 1. $R_A(p)$ *is an* $n^2 \times n^2$ *matrix; however, due to its structure, it can be multiplied by a vector in* $\mathbb{R}^{n^2}$ *using* $O(mn^2)$ *operations rather than the naive* $O(n^4)$.
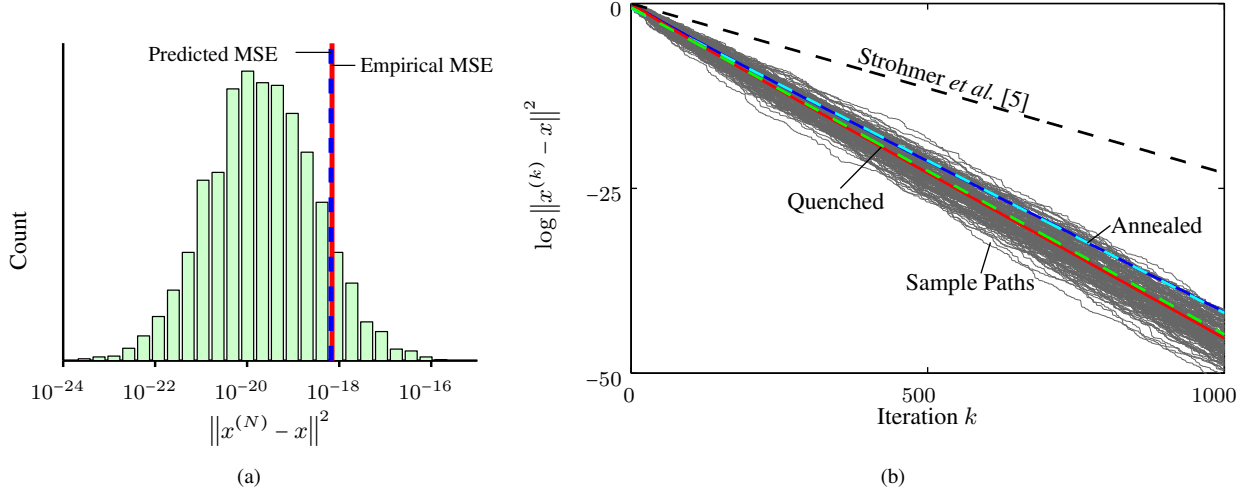
Fig. 1: (a) Histogram of squared errors after the simulation described in Section 2.3. The errors are plotted on a logarithmic scale to show the full range of errors; on a linear scale, the histogram is an L-shaped distribution with a spike at the origin and a long, thin tail. The location of the empirical MSE is overlaid on the histogram (red solid line), as is the exact MSE as given in Proposition 1 (blue dashed line). (b) Of the 3007 simulation trials, the "error trajectories" of 150 randomly-selected trials are plotted here (gray lines). On a logarithmic scale, there is a clear linear trend. Overlaid on these trajectories is the (annealed) average error trajectory (blue solid line) of all 3007 trials, and the prediction based on the annealed error exponent (cyan dashed line). We have also plotted the quenched average error trajectory, i.e. the average of the log of the error (red solid line), and the prediction based on the quenched error exponent (green dashed line) as given in (16). These are much more representative of the typical behavior of the algorithm. The upper bound of Strohmer *et al.* [5] is also shown (black dashed line).

2.3. Error Exponents: Annealed vs. Quenched

Proposition 1 confirms earlier bounds showing that the error decays exponentially. In fact, for generic values of the initial error vector, we have $\mathbb{E}\left\|z^{(N)}\right\|^2 = \exp(-\gamma_a N + o(N))$, where γ_a is the *annealed* error exponent, defined by

$$\gamma_a \overset{\text{def}}{=} \lim_{N\to\infty} -\frac{1}{N} \log \mathbb{E}\left\|z^{(N)}\right\|^2. \tag{13}$$

It is not hard to see that $\gamma_a = -\log \lambda_{\max}(R_A(p))$, where $\lambda_{\max}(\cdot)$ is the largest eigenvalue of a matrix.

To test our result, we simulated 3007 trials of the Kaczmarz algorithm for solving a linear system of dimension 150×20. The same system was used for each run, as well as the same initial vector. The matrix A was chosen to have independent standard normal entries (note that none of our analysis depends on A being drawn in this way, and similar results can be obtained with other matrices). We tracked the error after every iteration for each run. The row was chosen uniformly at random for each iteration. Figure 1(a) shows a histogram of the errors after 1000 iterations. The histogram was computed and is plotted on a logarithmic scale because of the wide range of resulting errors. The empirical MSE is overlaid on the histogram, as well as our prediction based on Proposition 1.

It is clear that our prediction matches the empirical value quite well. However, it is also clear that there is more to the story. Over 90% of the realizations have an error smaller than the mean, which is more than 10^2 times smaller than the worst realization. It appears that the average error is not necessarily a great representation of the *typical* error; in reality, there are occasional, rare, extreme failures that cause the average error to be much higher than the "typical" error.

A more representative measure of the error's decay rate is the *quenched* error exponent:

$$\gamma_q \overset{\text{def}}{=} \lim_{N\to\infty} -\frac{1}{N} \mathbb{E} \log \left\|z^{(N)}\right\|^2. \tag{14}$$

Here, the logarithm of the error is taken *before* the expectation. The annealed and quenched error exponents we have defined are formally similar to Lyapunov exponents of products of random matrices, a problem well-studied by statistical physicists for use in modeling dynamical systems [13]. The terms "annealed" and "quenched" are borrowed from their analysis and have certain physical meanings, but to us they are just convenient names for two interesting quantities.

The quenched error exponent is far more difficult to analyze than the annealed one, a fact well known to the physicists [13, 14]. Jensen's inequality tells us that $\gamma_q \geq \gamma_a$. To obtain more information, physicists often rely on non-rigorous heuristics that are verified numerically or experimentally. One such heuristic is the replica method, which provides an approximation for the quenched Lyapunov exponent [13]. The physicists have their own intuition for this approximation, but our engineer's intuition is quite simple. The quintessential heavy-tailed distribution is the log-normal distribution. So let us assume that the error distribution is $\left\|z^{(N)}\right\|^2 \sim \log\text{-}\mathcal{N}(N\mu, N\sigma^2)$. Then $\log\left\|z^{(N)}\right\|^2 \sim \mathcal{N}(N\mu, N\sigma^2)$. The log-normal assumption is supported by the histogram in Figure 1(a): the logarithm of the squared errors appear to follow a Gaussian distribution. The quenched error exponent is seen to be simply $\gamma_q = -\mu$. Now we need to compute the parameters of the distribution. Under these assumptions, $\mathbb{E}\left\|z^{(N)}\right\|^2 = \exp(N[\mu + \frac{1}{2}\sigma^2])$ and $\mathbb{E}\left\|z^{(N)}\right\|^4 = \exp(N[2\mu + 2\sigma^2])$. Solving this system of equations, we obtain:

$$\mu = \frac{1}{N}\left[2\log\mathbb{E}\left\|z^{(N)}\right\|^2 - \frac{1}{2}\log\mathbb{E}\left\|z^{(N)}\right\|^4\right]. \tag{15}$$

Thus, our approximation for the quenched error exponent is

$$\gamma_q \approx 2\gamma_a - \frac{1}{2}\gamma_a^{(2)}, \tag{16}$$
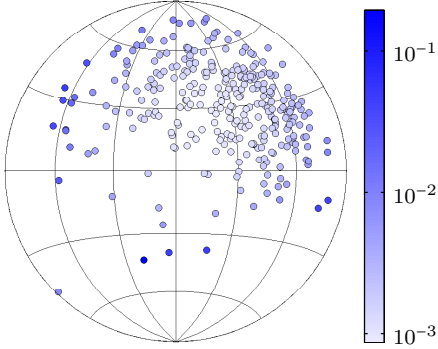
Fig. 2: Optimal selection probabilities for a non-uniform matrix. The plot is an equal-area projection of the entire unit hemisphere in R^3. Each row in the matrix is represented by a point on the plot; the color represents the optimal selection probability computed using `cvx`.

where

$$\gamma_a^{(2)} = \lim_{N \to \infty} -\frac{1}{N} \log \mathbb{E} \left\| \boldsymbol{z}^{(N)} \right\|^4. \qquad (17)$$

To compute $\gamma_a^{(2)}$, we define

$$R_{\boldsymbol{A}}^{(2)}(\boldsymbol{p}) = \sum_{i=1}^{m} p_i \left(\boldsymbol{P}_{\boldsymbol{a}_i}^{\perp} \otimes \boldsymbol{P}_{\boldsymbol{a}_i}^{\perp} \otimes \boldsymbol{P}_{\boldsymbol{a}_i}^{\perp} \otimes \boldsymbol{P}_{\boldsymbol{a}_i}^{\perp} \right), \qquad (18)$$

and have

$$\gamma_a^{(2)} = -\log \lambda_{\max} \left(R_{\boldsymbol{A}}^{(2)}(\boldsymbol{p}) \right). \qquad (19)$$

$R_{\boldsymbol{A}}^{(2)}(\boldsymbol{p})$ is an $n^4 \times n^4$ matrix, but it can be applied in time $O(mn^4)$ instead of the naive $O(n^8)$. So finding the largest eigenvalue is not as complex as one might naively expect.

Figure 1(b) illustrates our argument and shows just how good the replica method approximation is. We have plotted, on a logarithmic scale, the error trajectory of many trials as the iterations proceeded. (Only 150 randomly-selected trials are shown to prevent the figure from getting too cluttered). We have also plotted the logarithm of the average error, which matches the linear trendline predicted by the annealed error exponent γ_a, and the average of the logarithm of the error trajectories, which matches the linear trendline predicted by our approximation for the quenched error exponent γ_q. The quenched values are clearly more representative of the typical performance of the algorithm than the annealed ones. The close match indicates that our approximation is valid. For comparison purposes, we have also plotted the upper bound provided by Strohmer *et al.* [5].

3. OPTIMAL ROW-SELECTION PROBABILITIES

Given a matrix $\boldsymbol{A}$, we may wish to choose the row selection probabilities $p_1, p_2, \ldots, p_m$ that provide the fastest convergence. A tractable way to do this is to optimize the annealed error exponent γ_a, which measures the decay rate of the MSE. This is equivalent to the following optimization problem:

$$(p_1, \ldots, p_m) = \underset{\boldsymbol{p} \in \Delta^{n-1}}{\arg\min} \, \lambda_{\max}(R_{\boldsymbol{A}}(\boldsymbol{p})), \qquad (20)$$

where Δ^{n-1} is the unit simplex in $\mathbb{R}^n$. The function $\lambda_{\max}(R_{\boldsymbol{A}}(\boldsymbol{p}))$ is convex [15], as is the set Δ^{n-1}, so (20) is a convex optimization problem (more specifically, it is a semidefinite programming problem). Thus, finding the optimal probability distribution $\boldsymbol{p}$ is quite tractable. Note that Dai *et al.* recently considered an optimized randomized Kaczmarz algorithm [11], in which the row-selection probabilities were chosen to optimize a bound on the MSE's decay rate. However, we optimize the *exact* decay rate of the MSE.
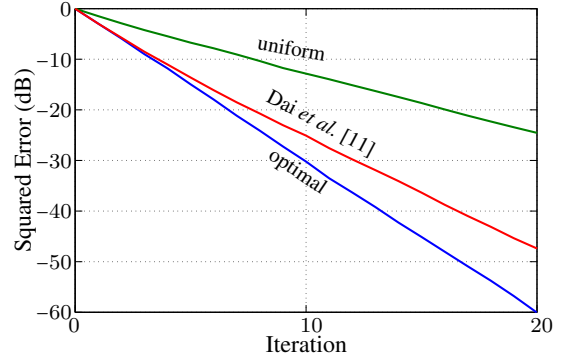
Fig. 3: Quenched average squared errors versus RKA iteration under the uniform, Dai *et al.*'s approximate optimal, and optimal row selection probabilities, for the 1000 x 3 matrix described in the text. The average is taken over 1007 trials.

To illustrate the kind of improvement possible by optimizing the row selection probabilities, and develop some intuition on the optimum choice, we computed the optimal values for a matrix of size 300×3. The elements of the matrix were chosen as independent Gaussian random variables with a variance of 0.5; the columns had means 0.5, 1, and 2, respectively. We used the `cvx` convex optimization software package to compute the optimal row selection probabilities for this matrix [16, 17].

Since the problem is invariant to the scale and sign of each rows, each row in the matrix $\boldsymbol{A}$ can be represented as a point on the unit hemisphere. Thus, the matrix and row probabilities can be illustrated as in Figure 2 by plotting each row as a point on a 2D projection of a unit hemisphere. We used the Lambert equal-area projection, which is measure-preserving and therefore allows us to accurately visualize the sampling density everywhere in the space. The darker points represent rows that are selected with high probability in the optimal selection scheme; the lighter ones are selected with lower probability. We would expect an optimal scheme to choose rows that are far from any other rows with higher probability than rows that are in close proximity to many other rows, in order to reduce redundancy and cover the whole space. The figure conforms to this intuition.

Figure 3 illustrates the improvement of the optimal randomization scheme over simply choosing rows uniformly at random. After 20 iterations, the optimal scheme has an error 36 dB lower than the uniform scheme, and 12 dB lower than the sub-optimal scheme of Dai *et al.*

Of course, in practice, there is a tradeoff between the computation time saved by needing fewer iterations and the computation time spent determining the optimal row selection probabilities in advance. The main purpose of the exact optimization proposed in this work is to develop intuition and validate sub-optimal heuristics. A fast or on-line method for approximating the optimal probabilities would be very beneficial for large-scale problems.

4. CONCLUSIONS

We provided a complete characterization of the randomized Kaczmarz algorithm. This included an exact formula for the MSE and the annealed error exponent characterizing its decay rate, plus an approximation for the quenched error exponent that captures the typical error decay rate. We also explored choosing the row-selection probabilities to achieve the best convergence properties for the algorithm.

5. REFERENCES

[1] S. Kaczmarz, "Angenäherte auflösung von systemen linearer gleichungen," *Bull. Internat. Acad. Polon. Sci. Lettres A*, pp. 335–357, 1937.

[2] R. Gordon, R. Bender, and G. T. Herman, "Algebraic reconstruction techniques (ART) for three-dimensional electron microscopy and X-ray photography," *Journal of theoretical Biology*, vol. 29, no. 3, p. 471–481, 1970.

[3] H. Trussell and M. Civanlar, "Signal deconvolution by projection onto convex sets," in *Acoustics, Speech, and Signal Processing, IEEE International Conference on ICASSP '84.*, vol. 9, Mar. 1984, pp. 496–499.

[4] G. T. Herman and L. B. Meyer, "Algebraic reconstruction techniques can be made computationally efficient [positron emission tomography application]," *Medical Imaging, IEEE Transactions on*, vol. 12, no. 3, p. 600–609, 1993.

[5] T. Strohmer and R. Vershynin, "A randomized Kaczmarz algorithm with exponential convergence," *Journal of Fourier Analysis and Applications*, vol. 15, no. 2, p. 262–278, 2009.

[6] Y. Censor, G. T. Herman, and M. Jiang, "A note on the behavior of the randomized Kaczmarz algorithm of Strohmer and Vershynin," *Journal of Fourier Analysis and Applications*, vol. 15, no. 4, pp. 431–436, Aug. 2009.

[7] D. Needell, "Randomized Kaczmarz solver for noisy linear systems," *BIT Numerical Mathematics*, vol. 50, no. 2, pp. 395—403, 2010.

[8] X. Chen and A. M. Powell, "Almost sure convergence of the Kaczmarz algorithm with random measurements," *Journal of Fourier Analysis and Applications*, vol. 18, no. 6, pp. 1195—1214, 2012.

[9] A. Zouzias and N. M. Freris, "Randomized extended Kaczmarz for solving least squares," *SIAM Journal on Matrix Analysis and Applications*, vol. 34, no. 2, pp. 773–793, 2013.

[10] D. Needell and J. A. Tropp, "Paved with good intentions: Analysis of a randomized block Kaczmarz method," *Linear Algebra and its Applications*, vol. 441, pp. 199—221, 2014.

[11] L. Dai, M. Soltanalian, and K. Pelckmans, "On the randomized Kaczmarz algorithm," *IEEE Signal Process. Lett.*, vol. 21, no. 3, pp. 330–333, Mar. 2014.

[12] B. Recht and C. Ré, "Toward a noncommutative arithmetic-geometric mean inequality: Conjectures, case-studies, and consequences," in *Conference on Learning Theory*, 2012.

[13] A. Crisanti, G. Paladin, and A. Vulpiani, *Products of Random Matrices in Statistical Physics*, ser. Springer Series in Solid-State Sciences, M. Cardona, P. Fulde, K. Klitzing, H.-J. Queisser, and H. K. V. Lotsch, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 1993, vol. 104, 00378.

[14] J. N. Tsitsiklis and V. D. Blondel, "The lyapunov exponent and joint spectral radius of pairs of matrices are hard—when not impossible—to compute and to approximate," *Mathematics of Control, Signals, and Systems (MCSS)*, vol. 10, no. 1, pp. 31–40, 1997.

[15] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge University Press, 2003.

[16] M. Grant and S. Boyd, "CVX: Matlab software for disciplined convex programming, version 2.1," http://cvxr.com/cvx, Mar. 2014.

[17] ——, "Graph implementations for nonsmooth convex programs," in *Recent Advances in Learning and Control*, ser. Lecture Notes in Control and Information Sciences, V. Blondel, S. Boyd, and H. Kimura, Eds. Springer-Verlag Limited, 2008, pp. 95–110.