

Phase Retrieval via Randomized Kaczmarz: Theoretical Guarantees

Yan Shuo Tan ^{*} Roman Vershynin [†]

January 11, 2018

Abstract

We consider the problem of phase retrieval, i.e. that of solving systems of quadratic equations. A simple variant of the randomized Kaczmarz method was recently proposed for phase retrieval, and it was shown numerically to have a computational edge over state-of-the-art Wirtinger flow methods. In this paper, we provide the first theoretical guarantee for the convergence of the randomized Kaczmarz method for phase retrieval. We show that it is sufficient to have as many Gaussian measurements as the dimension, up to a constant factor. Along the way, we introduce a sufficient condition on measurement sets for which the randomized Kaczmarz method is guaranteed to work. We show that Gaussian sampling vectors satisfy this property with high probability; this is proved using a chaining argument coupled with bounds on VC dimension and metric entropy.

1 Introduction

The phase retrieval problem is that of solving a system of quadratic equations

$$|\langle a_i, z \rangle|^2 = b_i^2, \quad i = 1, 2, \dots, m \quad (1.1)$$

where $a_i \in \mathbb{R}^n$ (or \mathbb{C}^n) are known sampling vectors, $b_i > 0$ are observed measurements, and $z \in \mathbb{R}^n$ (or \mathbb{C}^n) is the decision variable. This problem is well motivated by practical concerns [13] and has been a topic of study from at least the early 1980s.

Over the last half a decade, there has been great interest in constructing and analyzing algorithms with provable guarantees given certain classes of sampling vector sets. One line of research involves “lifting” the quadratic system to a linear system, which is then solved using convex relaxation (*PhaseLift*) [5]. A second method is to formulate and solve a linear program in the natural parameter space using an anchor vector (*PhaseMax*) [14, 2, 15]. Although both of these methods can be proved to have near optimal sample efficiency, the most empirically successful approach has been to directly optimize various naturally-formulated non-convex loss functions, the most notable of which are displayed in Table 1.

These loss functions enjoy nice properties which make them amenable to various optimization schemes [19, 12]. Those with provable guarantees include the prox-linear method of [11], and various gradient descent methods [4, 6, 24, 26, 9]. Some of these methods also involve adaptive measurement pruning to enhance performance.

^{*}Department of Mathematics, University of Michigan, yanshuo@umich.edu.

[†]Department of Mathematics, University of California, Irvine, rvershyn@uci.edu.

Loss function	Name	Papers
$f(z) = \sum_{i=1}^m (a_i, z ^2 - b_i^2)^2$	Wirtinger Flow objective	[4, 19]
$f(z) = \sum_{i=1}^m (a_i, z - b_i)^2$	Amplitude Flow objective	[24, 26]
$f(z) = \sum_{i=1}^m \left a_i, z ^2 - b_i^2 \right $	Robust Wirtinger Flow objective	[12, 11, 9]

Table 1: Non-convex loss functions for phase retrieval

In 2015, Wei [25] proposed adapting a family of randomized Kaczmarz methods for solving the phase retrieval problem. He was able to show using numerical experiments that these methods perform comparably with state-of-the-art Wirtinger flow (gradient descent) methods when the sampling vectors are real or complex Gaussian, or when they follow the coded diffraction pattern (CDP) model [4]. He also showed that randomized Kaczmarz methods outperform Wirtinger flow when the sampling vectors are the concatenation of a few unitary bases. Unfortunately, [25] was not able to provide adequate theoretical justification for the convergence of these methods (see Theorem 2.6 in [25]).

In this paper, we attempt to bridge this gap by showing that the basic randomized Kaczmarz scheme used in conjunction with truncated spectral initialization achieves *linear convergence* to the solution with high probability, whenever the sampling vectors are drawn uniformly from the sphere¹ S^{n-1} and the number of measurements m is larger than a constant times the dimension n .

It is also interesting to note that the basic randomized Kaczmarz scheme is exactly *stochastic gradient descent* for the Amplitude Flow objective, which suggests that other gradient descent schemes can also be accelerated using stochasticity.

1.1 Randomized Kaczmarz for solving linear systems

The Kaczmarz method is a fast iterative method for solving systems of overdetermined linear equations that works by iteratively satisfying one equation at a time. In 2009, Strohmer and Vershynin [18] were able to give a provable guarantee on its rate of convergence, provided that the equation to be satisfied at each step is selected using a prescribed randomized scheme.

Suppose our system to be solved is given by

$$Ax = b, \tag{1.2}$$

where A is an m by n matrix. Denoting the rows of A by a_1^T, \dots, a_m^T , we can write (1.2) as the system of linear equations

$$\langle a_i, x \rangle = b_i, \quad i = 1, \dots, m.$$

The solution set of each equation is a hyperplane. The randomized Kaczmarz method is a simple iterative algorithm in which we *project the running approximation onto the hyperplane of a randomly chosen equation*. More formally, at each step k we randomly choose an index $r(k)$ from $[m]$ such

¹This is essentially equivalent to being real Gaussian because of the concentration of norm phenomenon in high dimensions. Also, one may normalize vectors easily.

that the probability that $r(k) = i$ is proportional to $\|a_i\|_2^2$, and update the running approximation as follows:

$$x_k := x_{k-1} + \frac{b_{r(k)} - \langle a_{r(k)}, x_{k-1} \rangle}{\|a_{r(k)}\|_2^2} a_{r(k)}.$$

Strohmer and Vershynin [18] were able to prove the following theorem:

Theorem 1.1 (Linear convergence for linear systems). *Let $\kappa(A) = \|A\|_F / \sigma_{\min}(A)$. Then for any initialization x_0 to the equation (1.2), the estimates given to us by randomized Kaczmarz satisfy*

$$\mathbb{E}\|x_k - x\|_2^2 \leq (1 - \kappa(A)^{-2})^k \|x_0 - x\|_2^2.$$

Note that if A has bounded condition number, then $\kappa(A) \asymp 1/\sqrt{n}$.

1.2 Randomized Kaczmarz for phase retrieval

In the phase retrieval problem (1.1), each equation

$$|\langle a_i, x \rangle| = b_i$$

defines two hyperplanes, one corresponding to each of $\pm x$. A natural adaptation of the randomized Kaczmarz update for this situation is then to project the running approximation to the *closer hyperplane*. We restrict to the case where each measurement vector a_i has unit norm, so that in equations, this is given by

$$x_k := x_{k-1} + \eta_k a_{r(k)}, \tag{1.3}$$

where

$$\eta_k = \text{sign}(\langle a_{r(k)}, x_{k-1} \rangle) b_{r(k)} - \langle a_{r(k)}, x_{k-1} \rangle.$$

In order to obtain a convergence guarantee for this algorithm, we need to choose x_0 so that it is close enough to the signal vector x . This is unlike the case for linear systems where we could start with an arbitrary initial estimate $x_0 \in \mathbb{R}^n$, but the requirement is par for the course for phase retrieval algorithms. Unsurprisingly, there is a rich literature on how to obtain such estimates [5, 6, 26, 24]. The best methods are able to obtain a good initial estimate using $O(n)$ samples.

1.3 Contributions and main result

The main result of our paper guarantees the linear convergence of randomized Kaczmarz algorithm for phase retrieval for random measurements a_i that are drawn independently and uniformly from the unit sphere.

Theorem 1.2 (Convergence guarantee for algorithm). *Fix $\epsilon > 0$, $0 < \delta_1 \leq 1/2$, and $0 < \delta, \delta_2 \leq 1$. There are absolute constants $C, c > 0$ such that if*

$$m \geq C(n \log(m/n) + \log(1/\delta)),$$

then with probability at least $1 - \delta$, m sampling vectors selected uniformly and independently from the unit sphere S^{n-1} form a set such that the following holds: Let $x \in \mathbb{R}^n$ be a signal vector and let x_0 be an initial estimate satisfying $\|x_0 - x\|_2 \leq c\sqrt{\delta_1}\|x\|_2$. Then for any $\epsilon > 0$, if

$$K \geq 2(\log(1/\epsilon) + \log(2/\delta_2))n,$$

then the K -th step randomized Kaczmarz estimate x_K satisfies $\|x_K - x\|_2^2 \leq \epsilon \|x_0 - x\|_2^2$ with probability at least $1 - \delta_1 - \delta_2$.

Comparing this result with Theorem 1.1, we observe two key differences. First, there are now two sources of randomness: one is in the creation of the measurements a_i , and the other is in the selection of the equation at every iteration of the algorithm. The theorem gives a guarantee that holds with high probability over both sources of randomness. Theorem 1.2 also requires an initial estimate x_0 . This is not hard to obtain. Indeed, using the *truncated spectral initialization* method of [6], we may obtain such an estimate with high probability given $m \gtrsim n$. For more details, see Proposition B.1.

The proof of this theorem is more nontrivial than the Strohmer-Vershynin analysis of randomized Kaczmarz algorithm for linear systems [18]. We break down the argument in smaller steps, each of which may be of independent interest to researchers in this field.

First, we generalize the Kaczmarz update formula (1.3) and define what it means to take a randomized Kaczmarz step with respect to any probability measure on the sphere S^{n-1} : we choose a measurement vector at each step according to this measure. Using a simple geometric argument, we then provide a bound for the expected decrement in distance to the solution set in a single step, where the quality of the bound is given in terms of the properties of the measure we are using for the Kaczmarz update (Lemma 2.1).

Performing the generalized Kaczmarz update with respect to the uniform measure on the sphere corresponds to running the algorithm with unlimited measurements. We utilize the symmetry of the uniform measure to compute an explicit formula for the bound on the stepwise expected decrement in distance. This decrement is geometric whenever we make the update from a point making an angle of less than $\pi/8$ with the true solution, so we obtain linear convergence conditioned on no iterates escaping from the “basin of linear convergence”. We are able to bound the probability of this bad event using a supermartingale inequality (Theorem 3.1).

Next, we abstract out the property of the uniform measure that allows us to obtain local linear convergence. We call this property the *anti-concentration on wedges* property, calling it ACW for short. Using this convenient definition, we can easily generalize our previous proofs for the uniform measure to show that all ACW measures give rise to randomized Kaczmarz update schemes with local linear convergence (Theorem 4.3).

The usual Kaczmarz update corresponds running the generalized Kaczmarz update with respect to $\mu_A := \frac{1}{m} \sum_{i=1}^m \delta_{a_i}$. We are able to prove that when the a_i ’s are selected uniformly and independently from the sphere, then μ_A satisfies the ACW condition with high probability, so long as $m \gtrsim n$ (Theorem 5.9). The proof of this fact uses VC theory and a *chaining argument*, together with metric entropy estimates.

Finally, we are able to put everything together to prove a guarantee for the full algorithm in Section 6. In that section, we also discuss the failure probabilities δ , δ_1 and δ_2 , and how they can be controlled.

1.4 Related work

During the preparation of this manuscript, we became aware of independent simultaneous work done by Jeong and Güntürk. They also studied the randomized Kaczmarz method adapted to phase retrieval, and obtained almost the same result that we did (see [16] and Theorem 1.1 therein). In order to prove their guarantee, they use a stopping time argument similar to ours, but replace the ACW condition with a stronger condition called *admissibility*. They prove that measurement systems comprising vectors drawn independently and uniformly from the sphere satisfy this property

with high probability, and the main tools they use in their proof are hyperplane tessellations and a net argument together with Lipschitz relaxation of indicator functions.

After submitting the first version of this manuscript, we also became aware of independent work done by Zhang, Zhou, Liang, and Chi [26]. Their work examines stochastic schemes in more generality (see Section 3 in their paper), and they claim to prove linear convergence for both the randomized Kaczmarz method as well as what they called *Incremental Reshaped Wirtinger Flow*. However, they only prove that the distance to the solution decreases in expectation under a single Kaczmarz update (an analogue of our Lemma 2.1 specialized to real Gaussian measurements). As we will see in our paper, this bound cannot be naively iterated.

1.5 Notation

Throughout the paper, C and c are absolute constants that can change from line to line.

2 Computations for a single step

In this section, we will compute what happens in expectation for a single update step of the randomized Kaczmarz method. It will be convenient to generalize our sampling scheme slightly as follows. When we work with a fixed matrix A , we may view our selection of a random row $a_{r(k)}$ as drawing a random vector according to the measure $\mu_A := \frac{1}{m} \sum_{i=1}^m \delta_{a_i}$. We need not restrict ourselves to sums of Diracs. For any probability measure μ on the sphere S^{n-1} , we define the random map $P = P_\mu$ on vectors $z \in \mathbb{R}^n$ by setting

$$Pz := z + \eta a, \quad (2.1)$$

where

$$\eta = \text{sign}(\langle a, z \rangle) |\langle a, x \rangle| - \langle a, z \rangle \quad \text{and} \quad a \sim \mu. \quad (2.2)$$

Note that as before, x is a fixed vector in \mathbb{R}^n (think of x as the actual solution of the phase retrieval problem). We call P_μ the *generalized Kaczmarz projection with respect to μ* . Using this update rule over independent realizations of P, P_1, P_2, \dots , together with an initial estimate x_0 , gives rise to a *generalized randomized Kaczmarz algorithm* for finding x : set the k -th step estimate to be

$$x_k := P_k P_{k-1} \cdots P_1 x_0. \quad (2.3)$$

Fix a vector $z \in \mathbb{R}^n$ that is closer to x than to $-x$, i.e. so that $\langle x, z \rangle > 0$, and suppose that we are trying to find x . Examining the formula in (2.2), we see that P projects z onto the right hyperplane (i.e., the one passing through x instead of the one passing through $-x$) if and only if $\langle a, z \rangle$ and $\langle a, x \rangle$ have the same sign. In other words, this occurs if and only if the random vector a does *not* fall into the region of the sphere defined by

$$W_{x,z} := \{v \in S^{n-1} \mid \text{sign}(\langle v, x \rangle) \neq \text{sign}(\langle v, z \rangle)\}. \quad (2.4)$$

This is the region lying between the two hemispheres with normal vectors x and z . We call such a region a *spherical wedge*, since in three dimensions it has the shape depicted in Figure 1.

When $a \notin W_{x,z}$, we can use the Pythagorean theorem to write

$$\|z - x\|_2^2 = \|Pz - x\|_2^2 + \langle z - x, a \rangle^2. \quad (2.5)$$




Figure 1: Geometry of $W_{x,z}$




Figure 2: Orientation of x , z , and Pz when $a \in W_{x,z}$ and when $a \notin W_{x,z}$. H_+ and H_- denote respectively the hyperplanes defined by the equations $\langle y, a \rangle = b$ and $\langle y, a \rangle = -b$. H_0 denotes the hyperplane defined by the equation $\langle y, a \rangle = 0$. The left diagram demonstrates the situation when $a \in W_{x,z}$, thereby justifying (2.5). The right diagram demonstrates the situation when $a \notin W_{x,z}$, thereby justifying (2.8).

Rearranging gives

$$\|Pz - x\|_2^2 = \|z - x\|_2^2(1 - \langle \tilde{z}, a \rangle^2), \quad (2.6)$$

where $\tilde{z} = (z - x)/\|z - x\|_2$.

In the complement of this event, we get

$$Pz = z + \langle a, (-x) - z \rangle a = z - \langle a, z - x \rangle + \langle a, -2x \rangle,$$

and using orthogonality,

$$\|Pz - x\|_2^2 = \|z - x\|_2^2 - \langle a, z - x \rangle^2 + \langle a, 2x \rangle^2. \quad (2.7)$$

Since z gets projected to the hyperplane containing $-x$, it may move further away from x . However, we can bound how far away it can move. Because $\langle a, x \rangle$ has the opposite sign as $\langle a, z \rangle$, we have

$$|\langle a, z + x \rangle| < |\langle a, z - x \rangle|,$$

and so

$$|\langle a, 2x \rangle| = |\langle a, (z - x) - (z + x) \rangle| < 2|\langle a, z - x \rangle|.$$

Substituting this into (2.7), we get the bound

$$\|Pz - x\|_2^2 \leq \|z - x\|_2^2 + 3\langle a, z - x \rangle^2 = \|z - x\|_2^2(1 + 3\langle \tilde{z}, a \rangle^2), \quad (2.8)$$

where \tilde{z} is as before.

We can combine (2.6) and (2.8) into a single inequality by writing

$$\begin{aligned} \|Pz - x\|_2^2 &\leq \|z - x\|_2^2(1 - \langle \tilde{z}, a \rangle^2)1_{W_{x,z}^c}(a) + \|z - x\|_2^2(1 + 3\langle \tilde{z}, a \rangle^2)1_{W_{x,z}}(a) \\ &= \|z - x\|_2^2(1 - (1 - 4 \cdot 1_{W_{x,z}}(a))\langle \tilde{z}, a \rangle^2) \\ &= \|z - x\|_2^2(1 - \langle \tilde{z}, (1 - 4 \cdot 1_{W_{x,z}}(a))aa^T \tilde{z} \rangle). \end{aligned}$$

Taking expectations, we can remove the role that \tilde{z} plays by bounding this as follows.

$$\begin{aligned} \mathbb{E}[\|z - x\|_2^2(1 - \langle \tilde{z}, (1 - 4 \cdot 1_{W_{x,z}}(a))aa^T \tilde{z} \rangle)] &= \|z - x\|_2^2(1 - \langle \tilde{z}, \mathbb{E}[(1 - 4 \cdot 1_{W_{x,z}}(a))aa^T] \tilde{z} \rangle) \\ &\leq \|z - x\|_2^2[1 - \lambda_{\min}(\mathbb{E}aa^T - 4\mathbb{E}aa^T 1_{W_{x,z}}(a))]. \end{aligned}$$

We may thus summarize what we have obtained in the following lemma.

Lemma 2.1 (Expected decrement). *Fix vectors $x, z \in \mathbb{R}^n$, a probability measure μ on S^{n-1} , and let $P = P_\mu$, $W_{x,z}$ be defined as in (2.1) and (2.4) respectively. Then*

$$\mathbb{E}\|Pz - x\|_2^2 \leq [1 - \lambda_{\min}(\mathbb{E}aa^T - 4\mathbb{E}aa^T 1_{W_{x,z}}(a))]\|z - x\|_2^2.$$

Let us next compute what happens for $\mu = \sigma$, the uniform measure on the sphere. It is easy to see that $\mathbb{E}aa^T = \frac{1}{n}I_n$, so it remains to compute $\mathbb{E}aa^T 1_{W_{x,z}}(a)$. To do this, we make a convenient choice of coordinates: Let θ be the angle between z and x . We assume that both points lie in the plane spanned by e_1 and e_2 , the first two basis vectors, and that the angle between z and x is bisected by e_1 , as illustrated in Figure 3.




Figure 3: Choice of coordinates

For convenience, denote $M := \mathbb{E}aa^T 1_{W_{x,z}}(a)$. Let Q denote the orthogonal projection operator onto the span of e_1 and e_2 . Then $Q(W_{x,z})$ is the union of two sectors of angle θ , which are respectively bisected by e_2 and $-e_2$. Recall that all coordinate projections of the uniform random vector a are uncorrelated. It is clear that from the symmetry in Figure 3 that they remain uncorrelated even when conditioning on the event that $a \in W_{x,z}$. As such, M is a diagonal matrix.

Let ϕ denote the anti-clockwise angle of Qa from e_2 (see Figure 3). We may write

$$\langle a, e_1 \rangle^2 = \|Qa\|_2^2 \langle Qa / \|Qa\|_2, e_1 \rangle^2 = \|Qa\|_2^2 \sin^2 \phi.$$

Note that the magnitude and direction of Qa are independent, and $a \in W_{x,z}$ if either ϕ or $\phi - \pi$ lies between $-\theta/2$ and $\theta/2$. We therefore have

$$M_{11} = \mathbb{E}[\langle a, e_1 \rangle^2 1_{W_{x,z}}(a)] = \mathbb{E}[\|Qa\|_2^2 \mathbb{E} \sin^2 \phi 1_{(-\theta/2, \theta/2)}(\phi \text{ or } \phi - \pi)].$$

By a standard calculation using symmetry, we have $\mathbb{E}\|Qa\|_2^2 = 2/n$. Since ϕ is distributed uniformly on the circle, we can compute

$$\mathbb{E} \sin^2 \phi 1_{(-\theta/2, \theta/2)}(\phi \text{ or } \phi - \pi) = \frac{1}{\pi} \int_{-\theta/2}^{\theta/2} \sin^2 t dt = \frac{1}{\pi} \int_{-\theta/2}^{\theta/2} \frac{1 - \cos(2t)}{2} dt = \frac{\theta - \sin \theta}{2\pi}.$$

As such, we have $M_{11} = (\theta - \sin \theta)/n\pi$, and by a similar calculation, $M_{22} = (\theta + \sin \theta)/n\pi$. Meanwhile, for $i \geq 3$ we have

$$\begin{aligned} M_{ii} &= \frac{\text{Tr}(M) - M_{11} - M_{22}}{n - 2} \\ &= \frac{\mathbb{E}[\|(I - Q)a\|_2^2 1_{W_{x,z}}(a)]}{n - 2} \\ &= \frac{\mathbb{E}\|(I - Q)a\|_2^2 \mathbb{E} 1_{(-\theta/2, \theta/2)}(\phi \text{ or } \phi - \pi)}{n - 2} \\ &= \frac{(n - 2)/n \cdot \theta/\pi}{n} = \frac{\theta}{n\pi}. \end{aligned}$$

This implies that

$$\lambda_{\max}(M_\theta) = \frac{\theta + \sin \theta}{n\pi}. \quad (2.9)$$

We have now completed proving the following lemma.

Lemma 2.2 (Expected decrement for uniform measure). *Fix vectors $x, z \in \mathbb{R}^n$ such that $\langle z, x \rangle > 0$, and let $P = P_\sigma$ denote the generalized Kaczmarz projection with respect to σ , the uniform measure on the sphere. Let θ be the angle between z and x . Then*

$$\mathbb{E}\|Pz - x\|_2^2 \leq \left[1 - \frac{1 - 4(\theta + \sin \theta)/\pi}{n}\right] \|z - x\|_2^2.$$

Remark 2.3. By being more careful, one may compute an *exact* formula for the expected decrement rather than a bound as is the case in previous lemma. This is not necessary for our purposes and does not give better guarantees in our analysis, so the computation is omitted.

3 Local linear convergence using unlimited uniform measurements

In this section, we will show that if we start with an initial estimate that is close enough to the ground truth x , then repeatedly applying generalized Kaczmarz projections with respect to the uniform measure σ gives linear convergence in expectation. This is exactly the situation we would

be in if we were to run randomized Kaczmarz given an unlimited supply of independent sampling vector a_1, a_2, \dots drawn uniformly from the sphere.

We would like to imitate the proof for linear convergence of randomized Kaczmarz for linear systems (Theorem 1.1) given in [18]. We denote by X_k the estimate after k steps, using capital letters to emphasize the fact that it is a random variable. If we know that X_k takes the value $x_k \in \mathbb{R}^n$, and the angle θ_k that z makes with x_k is smaller than $\pi/8$, then, Lemma 2.2 tells us

$$\mathbb{E}[\|X_{k+1} - x\|_2^2 \mid X_k = x_k] \leq (1 - \alpha_\sigma/n)\|x_k - x\|_2^2, \quad (3.1)$$

where $\alpha_\sigma := 1/2 - 4\sin(\pi/8)/\pi > 0$.

The proof for Theorem 1.1 proceeds by unconditioning and iterating a bound similar to (3.1). Unfortunately, our bound depends on x_k being in a specific region in \mathbb{R}^n and does not hold arbitrarily. Nonetheless, by using some basic concepts from stochastic process theory, we may derive a *conditional* linear convergence estimate. The details are as follows.

For each k , let \mathcal{F}_k denote the σ -algebra generated by a_1, a_2, \dots, a_k , where a_k is the sampling vector used in step k . Let $B \subset \mathbb{R}^n$ be the region comprising all points making an angle less than or equal to $\pi/8$ with x . This is our *basin of linear convergence*. Let us assume a fixed initial estimate $x_0 \in B$. Now define a stopping time τ via

$$\tau := \min\{k : X_k \notin B\}. \quad (3.2)$$

For each k , and $x_k \in B$, we have

$$\begin{aligned} \mathbb{E}[\|X_{k+1} - x\|_2^2 1_{\tau > k+1} \mid X_k = x_k] &\leq \mathbb{E}[\|X_{k+1} - x\|_2^2 1_{\tau > k} \mid X_k = x_k] \\ &= \mathbb{E}[\|X_{k+1} - x\|_2^2 1_{\tau > k} \mid X_k = x_k, \mathcal{F}_k] \\ &= \mathbb{E}[\|X_{k+1} - x\|_2^2 \mid X_k = x_k, \mathcal{F}_k] 1_{\tau > k} \\ &\leq (1 - \alpha_\sigma/n)\|x_k - x\|_2^2 1_{\tau > k}. \end{aligned}$$

Here, the first inequality follows from the inclusion $\{\tau > k+1\} \subset \{\tau > k\}$, the first equality statement from the Markov nature of the process (X_k) , the second equality statement from the fact that τ is a stopping time, while the second inequality is simply (3.1). Taking expectations with respect to X_k then gives

$$\begin{aligned} \mathbb{E}[\|X_{k+1} - x\|_2^2 1_{\tau > k+1}] &= \mathbb{E}[\mathbb{E}[\|X_{k+1} - x\|_2^2 1_{\tau > k+1} \mid X_k]] \\ &\leq (1 - \alpha_\sigma/n)\mathbb{E}[\|X_k - x\|_2^2 1_{\tau > k}]. \end{aligned}$$

By induction, we therefore obtain

$$\mathbb{E}[\|X_k - x\|_2^2 1_{\tau > k}] \leq (1 - \alpha_\sigma/n)^k \|x_0 - x\|_2^2.$$

We have thus proven the first part of the following convergence theorem.

Theorem 3.1 (Linear convergence from unlimited measurements). *Let x be a vector in \mathbb{R}^n , let $\delta > 0$, and let x_0 be an initial estimate to x such that $\|x_0 - x\|_2 \leq \delta\|x\|_2$. Suppose that our measurements a_1, a_2, \dots are fully independent random vectors distributed uniformly on the sphere*

S^{n-1} . Let X_k be the estimate given by the randomized Kaczmarz update formula (2.3) at step k , and let τ be the stopping time defined via (3.2). Then for every $k \in \mathbb{Z}_+$,

$$\mathbb{E}[\|X_k - x\|_2^2 1_{\tau=\infty}] \leq (1 - \alpha_\sigma/n)^k \|x_0 - x\|_2^2, \quad (3.3)$$

where $\alpha_\sigma = 1/2 - 4 \sin(\pi/8)/\pi > 0$. Furthermore, $\mathbb{P}(\tau < \infty) \leq (\delta/\sin(\pi/8))^2$.

Proof. In order to prove the second statement, we combine a stopping time argument with a supermartingale maximal inequality. Set $Y_k := \|X_{\tau \wedge k} - x\|_2^2$. We claim that Y_k is a supermartingale. To see this, we break up its conditional expectation as follows:

$$\begin{aligned} \mathbb{E}[Y_{k+1} \mid \mathcal{F}_k] &= \mathbb{E}[\|X_{\tau \wedge (k+1)} - x\|_2^2 1_{\tau \leq k} \mid \mathcal{F}_k] + \mathbb{E}[\|X_{\tau \wedge (k+1)} - x\|_2^2 1_{\tau > k} \mid \mathcal{F}_k] \\ &= \mathbb{E}[\|X_{\tau \wedge k} - x\|_2^2 1_{\tau \leq k} \mid \mathcal{F}_k] + \mathbb{E}[\|X_{k+1} - x\|_2^2 1_{\tau > k} \mid \mathcal{F}_k]. \end{aligned}$$

Since $\|X_{\tau \wedge k} - x\|_2^2$ is measurable with respect to \mathcal{F}_k , we get

$$\mathbb{E}[\|X_{\tau \wedge k} - x\|_2^2 1_{\tau \leq k} \mid \mathcal{F}_k] = \|X_{\tau \wedge k} - x\|_2^2 1_{\tau \leq k} = Y_k 1_{\tau \leq k}.$$

Meanwhile, on the event $\tau > k$, we have $X_k \in B$, so we may use (3.1) to obtain

$$\mathbb{E}[\|X_{k+1} - x\|_2^2 1_{\tau > k} \mid \mathcal{F}_k] = \mathbb{E}[\|X_{k+1} - x\|_2^2 \mid \mathcal{F}_k] 1_{\tau > k} \leq (1 - \alpha_\sigma/n) \|X_k - x\|_2^2 1_{\tau > k}.$$

Next, notice that

$$\|X_k - x\|_2^2 1_{\tau > k} = \|X_{\tau \wedge k} - x\|_2^2 1_{\tau > k} = Y_k 1_{\tau > k}.$$

Combining these calculations gives

$$\mathbb{E}[Y_{k+1} \mid \mathcal{F}_k] \leq Y_k 1_{\tau \leq k} + (1 - \alpha_\sigma/n) Y_k 1_{\tau > k} \leq Y_k.$$

Now define a second stopping time T to be the earliest time k such that $\|X_k - x\|_2 \geq \sin(\pi/8) \cdot \|x\|_2$. A simple geometric argument tells us that $T \leq \tau$, and that T also satisfies

$$T = \inf\{k \mid Y_k \geq \sin^2(\pi/8) \|x\|_2^2\}.$$

As such, we have

$$\mathbb{P}(\tau < \infty) \leq \mathbb{P}(T < \infty) = \mathbb{P}\left(\sup_{1 \leq k < \infty} Y_k \geq \sin^2(\pi/8) \|x\|_2^2\right).$$

Since (Y_k) is a non-negative supermartingale, we may apply the supermartingale maximal inequality to obtain a bound on the right hand side:

$$\mathbb{P}\left(\sup_{1 \leq k < \infty} Y_k \geq \sin^2(\pi/8) \|x\|_2^2\right) \leq \frac{\mathbb{E}Y_0}{\sin^2(\pi/8) \|x\|_2^2} \leq (\delta/\sin(\pi/8))^2.$$

This completes the proof of the theorem. \square

Corollary 3.2. Fix $\epsilon > 0$, $0 < \delta_1 \leq 1/2$, and $0 < \delta_2 \leq 1$. In the setting of Theorem 3.1, suppose that $\|x_0 - x\|_2 \leq \sqrt{\delta_1} \sin(\pi/8) \|x\|_2$. Then with probability at least $1 - \delta_1 - \delta_2$, if $k \geq (\log(2/\epsilon) + \log(1/\delta_2))n/\alpha_\sigma$ then $\|X_k - x\|_2^2 \leq \epsilon \|x_0 - x\|_2^2$.

Proof. First observe that

$$\mathbb{P}(\tau < \infty) \leq \left(\frac{\sqrt{\delta_1} \sin(\pi/8)}{\sin(\pi/8)} \right)^2 = \delta_1 \leq 1/2.$$

Next, since

$$\begin{aligned} \mathbb{E}[\|X_k - x\|_2^2 1_{\tau=\infty}] &= \mathbb{E}[\|X_k - x\|_2^2 \mid \tau = \infty] \mathbb{P}(\tau = \infty) + 0 \cdot \mathbb{P}(\tau < \infty) \\ &\geq \frac{1}{2} \mathbb{E}[\|X_k - x\|_2^2 \mid \tau = \infty], \end{aligned}$$

applying Theorem 3.1 gives

$$\mathbb{E}[\|X_k - x\|_2^2 \mid \tau = \infty] \leq 2(1 - \alpha_\sigma/n)^k \|x_0 - x\|_2^2.$$

Applying Markov's inequality then gives

$$\begin{aligned} \mathbb{P}(\|X_k - x\|_2^2 > \epsilon \|x_0 - x\|_2^2 \mid \tau = \infty) &\leq \frac{\mathbb{E}[\|X_k - x\|_2^2 \mid \tau = \infty]}{\epsilon \|x_0 - x\|_2^2} \\ &\leq \frac{2(1 - \alpha_\sigma/n)^k}{\epsilon}. \end{aligned}$$

Plugging our choice of k into this last bound shows that it is in turn bounded by δ_2 . We therefore have

$$\begin{aligned} \mathbb{P}(\|X_k - x\|_2^2 \leq \epsilon \|x_0 - x\|_2^2) &= \mathbb{P}(\|X_k - x\|_2^2 \leq \epsilon \|x_0 - x\|_2^2 \mid \tau = \infty) \mathbb{P}(\tau = \infty) \\ &\geq (1 - \delta_2)(1 - \delta_1) \\ &\geq 1 - \delta_1 - \delta_2 \end{aligned}$$

as we wanted. \square

4 Local linear convergence for $\text{ACW}(\theta, \alpha)$ measures

We would like to extend the analysis in the previous section to the setting where we only have access to finitely many uniform measurements, i.e. when we are back in the situation of (1.1). When we sample uniformly from the rows of A , this can be seen as running the generalized randomized Kaczmarz algorithm using the measure $\mu_A = \frac{1}{m} \sum_{i=1}^m \delta_{a_i}$ as opposed to $\mu = \sigma$.

If we retrace our steps, we will see that the key property of the uniform measure σ that we used was that if $W \subset S^{n-1}$ is a wedge² of angle θ , then we could make $\lambda_{\max}(\mathbb{E}_\sigma aa^T 1_W(a))$ arbitrarily small by taking θ small enough (see equation (2.9)). We do not actually need such a strong statement. It suffices for there to be an absolute constant α such that

$$\lambda_{\min}(\mathbb{E} aa^T - 4\mathbb{E} aa^T 1_W(a)) \geq \frac{\alpha}{n} \quad (4.1)$$

holds for θ small enough.

²Recall that a wedge of angle θ is the region of the sphere between two hemispheres with normal vectors making an angle of θ .

Definition 4.1 (Anti-concentration). If a probability measure μ on S^{n-1} satisfies (4.1) for all wedges W of angle less than θ , we say that it is *anti-concentrated on wedges of angle θ at level α* , or for short, that it satisfies the $\text{ACW}(\theta, \alpha)$ condition.

Abusing notation, we say that a measurement matrix A is $\text{ACW}(\theta, \alpha)$ if the uniform measure on its rows is $\text{ACW}(\theta, \alpha)$. Plugging in this definition into Lemma 2.1, we immediately get the following statement.

Lemma 4.2 (Expected decrement for ACW measure). *Let μ be a probability measure on the sphere S^{n-1} satisfying the $\text{ACW}(\theta, \alpha)$ condition for some $\alpha > 0$ and some acute angle $\theta > 0$. Let $P = P_\mu$ denote the generalized Kaczmarz projection with respect to μ . Then for any $x, z \in \mathbb{R}^n$ such that the angle between them is less than θ , we have*

$$\mathbb{E}\|Pz - x\|_2^2 \leq (1 - \alpha/n)\|z - x\|_2^2. \quad (4.2)$$

We may now imitate the arguments in the previous section to obtain a guarantee for local linear convergence for the generalized randomized Kaczmarz algorithm using such a measure μ .

Theorem 4.3 (Linear convergence for ACW measure). *Suppose μ is an $\text{ACW}(\theta, \alpha)$ measure. Let x be a vector in \mathbb{R}^n , let $\delta > 0$, and let x_0 be an initial estimate to x such that $\|x_0 - x\|_2 \leq \delta\|x\|_2$. Let X_k denote the k -th step of the generalized randomized Kaczmarz method with respect to the measure μ , defined as in (2.3). Let Ω be the event that for every $k \in \mathbb{Z}_+$, X_k makes an angle less than θ with x . Then for every $k \in \mathbb{Z}_+$,*

$$\mathbb{E}[\|X_k - x\|_2^2 1_\Omega] \leq (1 - \alpha/n)^k \|x_0 - x\|_2^2. \quad (4.3)$$

Furthermore, $\mathbb{P}(\Omega^c) \leq (\delta/\sin\theta)^2$.

Proof. We repeat the proof of Theorem 3.1. Let $B_\mu \subset S^{n-1}$ be the region on the sphere comprising all points making an angle less than or equal to $\pi/8$ with x . Define stopping times τ_μ and T_μ as the earliest times that $X_k \notin B_\mu$ and $\|X_k - x_0\|_2 \geq \sin(\theta)\|x\|_2$ respectively. Again, $Y_k := X_{k \wedge \tau_\mu}$ is a supermartingale, so we may use the supermartingale inequality to bound the probability of Ω^c . Conditioned on the event Ω , we may iterate the bound given by Lemma 4.2 to obtain (4.3). \square

Corollary 4.4. *Fix $\epsilon > 0$, $0 < \delta_1 \leq 1/2$, and $0 < \delta_2 \leq 1$. In the setting of Theorem 4.3, suppose that $\|x_0 - x\|_2 \leq \sqrt{\delta_1} \sin(\theta)\|x\|_2$. Then with probability at least $1 - \delta_1 - \delta_2$, if $k \geq (\log(2/\epsilon) + \log(1/\delta_2))n/\alpha$ then $\|X_k - x\|_2^2 \leq \epsilon\|x_0 - x\|_2^2$.*

5 $\text{ACW}(\theta, \alpha)$ condition for finitely many uniform measurements

Following the theory in the previous section, we see that to prove linear convergence from finitely many uniform measurements, it suffices to show that the measurement matrix A is $\text{ACW}(\theta, \alpha)$ for some θ and α .

For a *fixed* wedge W , we can easily achieve (4.1) by using a standard matrix concentration theorem. By taking a union bound, we can guarantee that it holds over exponentially many wedges with high probability. However, the function $W \mapsto \lambda_{\max}(\mathbb{E}aa^T 1_W(a))$ is not Lipschitz with respect to any natural parametrization of wedges in S^{n-1} , so a naive net argument fails. To get around this, we use VC theory, metric entropy, and a chaining theorem from [10].

First, we will use the theory of VC dimension and growth functions to argue that all wedges contain approximately the right fraction of points. This is the content of lemma 5.1. In order to prove this, a fair number of standard definitions and results are required. These are all provided in Appendix A.

Lemma 5.1 (Uniform concentration of empirical measure over wedges). *Fix an acute angle $\theta > 0$. Let \mathcal{W}_θ denote the collection of all wedges of S^{n-1} of angle less than θ . Suppose A is an m by n matrix with rows a_i that are independent uniform random vectors on S^{n-1} , and let $\mu_A = \frac{1}{m} \sum_{i=1}^m \delta_{a_i}$. Then if $m \geq (4\pi/\theta)^2(2n \log(2em/n) + \log(2/\delta))$, with probability at least $1 - \delta$, we have*

$$\sup_{W \in \mathcal{W}} \mu_A(W) \leq 2\theta/\pi.$$

Proof. Using VC theory [21], we have

$$\mathbb{P}(\sup_{W \in \mathcal{W}} |\mu_A(W) - \sigma(W)| \geq u) \leq 4\Pi_{\mathcal{W}_\theta}(2m) \exp(-mu^2/16) \quad (5.1)$$

whenever $m \geq 2/u^2$. Let \mathcal{S} be the collection of all sectors of any angle, and let \mathcal{H} denote the collection of all hemispheres. By Claim A.3 and the Sauer-Shelah lemma (Lemma A.1) relating VC dimension to growth functions, we have $\Pi_{\mathcal{H}}(2m) \leq (2em/n)^n$.

Next, notice that using the notation in (A.2), we have $\mathcal{W} = \mathcal{H} \triangle \mathcal{H}$. As such, we may apply Claim A.4 to get

$$\Pi_{\mathcal{W}}(2m) \leq (2em/n)^{2n}.$$

We now plug this bound into the right hand side of (5.1), set $u = \theta/\pi$, and simplify to get

$$\mathbb{P}(\sup_{W \in \mathcal{W}} |\mu_A(W) - \sigma(W)| \geq \theta/\pi) \leq 4 \exp(2n \log(2em/n) - m(\theta/\pi)^2/16).$$

Our assumption implies that $m \geq 2/(\theta/\pi)^2$ so the bound holds, and also that the bound is less than δ . Finally, since $\mathcal{W}_\theta \subset \mathcal{W}$, on the complement of this event, any $W \in \mathcal{W}_\theta$ satisfies

$$\mu_A(W) \leq \sigma(W) + \theta/\pi \leq 2\theta/\pi$$

as we wanted. □

For every wedge $W \in \mathcal{W}_\theta$, we may associate the configuration vector

$$s_{W,A} := (1_W(a_1), 1_W(a_2), \dots, 1_W(a_m)).$$

We can write

$$\lambda_{\max}(\mathbb{E}_{\mu_A} a a^T 1_W(a)) = \frac{1}{m} \lambda_{\max}(A^T S_{W,A} A), \quad (5.2)$$

where $S_{W,A} = \text{diag}(s_{W,A})$. $S_{W,A}$ is thus a selector matrix, and if we condition on the good event given to us by the previous theorem, it selects at most a $2\theta/\pi$ fraction of the rows of A . This means that $s_{W,A} \in \mathcal{S}_{2\theta/\pi}$, where we define

$$\mathcal{S}_\tau := \{d \in \{0, 1\}^m \mid \langle d, 1 \rangle \leq \tau \cdot m\}.$$

We would like to majorize the quantity in (5.2) uniformly over all wedges W by the quantity $\frac{1}{4}\lambda_{\min}(\mathbb{E}_{\mu_A}aa^T)$. In order to do this, we define a stochastic process $(Y_{s,v})$ indexed by $s \in \mathcal{S}_{2\theta/\pi}$ and $v \in B_2^n$, setting

$$Y_{s,v} := nv^T A^T \text{diag}(s) A = \sum_{i=1}^m s_i \langle \sqrt{n}a_i, v \rangle^2. \quad (5.3)$$

If we condition on the good set in Lemma 5.1, it is clear that

$$\sup_{W \in \mathcal{W}_\theta} \frac{1}{m} \lambda_{\max}(A^T S_{W,A} A) \leq \frac{1}{nm} \sup_{s \in \mathcal{S}_{2\theta/\pi}, v \in B_2^n} Y_{s,v},$$

so it suffices to bound the quantity on the right. We will do this using a slightly sophisticated form of chaining, which requires us to make a few definitions.

Let (T, d) be a metric space. A sequence $\mathcal{T} = (T_k)_{k \in \mathbb{Z}_+}$ of subsets of T is called *admissible* if $|T_0| = 1$, and $|T_k| \leq 2^{2^k}$ for all $k \geq 1$. For any $0 < \alpha < \infty$, we define the γ_α *functional* of (T, d) to be

$$\gamma_\alpha(T, d) := \inf_{\mathcal{T}} \sup_{t \in T} \sum_{k=0}^{\infty} 2^{k/\alpha} d(t, T_k).$$

Let d_1 and d_2 be two metrics on T . We say that a process (Y_t) has *mixed tail increments* with respect to (d_1, d_2) if there are constants c and C such that for all $s, t \in T$, we have the bound

$$\mathbb{P}(|Y_s - Y_t| \geq c(\sqrt{u}d_2(s, t) + ud_1(s, t))) \leq Ce^{-u}. \quad (5.4)$$

Remark 5.2. In [10], processes with mixed tail increments are defined as above but with the further restriction that $c = 1$ and $C = 2$. This is not necessary for the result that we need (Lemma 5.3) to hold. The indeterminacy of c and C gets absorbed into the final constant in the bound.

Lemma 5.3 (Theorem 5, [10]). *If $(Y_t)_{t \in T}$ has mixed tail increments, then there is a constant C such that for any $u \geq 1$, with probability at least $1 - e^{-u}$,*

$$\sup_{t \in T} |Y_t - Y_{t_0}| \leq C(\gamma_2(T, d_2) + \gamma_1(T, d_1) + \sqrt{u} \text{diam}(T, d_2) + u \text{diam}(T, d_1)).$$

At first glance, the γ_2 and γ_1 quantities seem mysterious and intractable. We will show however, that they can be bounded by more familiar quantities that are easily computable in our situation. Let us postpone this for the moment, and first show that our process $(Y_{s,v})$ has mixed tail increments.

Lemma 5.4 ($(Y_{s,v})$ has mixed tail increments). *Let $(Y_{s,v})$ be the process defined in (5.3). Define the metrics d_1 and d_2 on $\mathcal{S}_{2\theta/\pi} \times B_2^n$ using the norms $\| (w, v) \|_1 = \max\{\|w\|_\infty, \|v\|_2\}$ and $\| (w, v) \|_2 = \max\{\|w\|_2, \sqrt{2m\theta/\pi}\|v\|_2\}$. Then the process has mixed tail increments with respect to (d_1, d_2) .*

Proof. The main tool that we use is Bernstein's inequality [23] for sums of subexponential random variables. Observe that each $\sqrt{n}a_i$ is a subgaussian random vector with bounded subgaussian norm $\|\sqrt{n}a_i\|_{\psi_2} \leq C$, where C by an absolute constant. As such, for any $v \in B_2^n$, $\langle \sqrt{n}a_i, v \rangle^2$ is a subexponential random variable with bounded subexponential norm $\|\langle \sqrt{n}a_i, v \rangle^2\|_{\psi_1} \leq C^2$ [23].

Now fix v and let $s, s' \in \mathcal{S}_{2\theta/\pi}$. Then

$$Y_{s,v} - Y_{s',v} = \sum_{i=1}^m (s_i - s'_i) \langle \sqrt{n}a_i, v \rangle^2.$$

Using Bernstein, we have

$$\mathbb{P}(|Y_{s,v} - Y_{s',v}| \geq u) \leq 2 \exp(-c \min\{u^2/\|s - s'\|_2^2, u/\|s - s'\|_\infty\}). \quad (5.5)$$

Similarly, if we fix $s \in \mathcal{S}_{2\theta/\pi}$ and let $v, v' \in B_2^n$, then

$$\begin{aligned} Y_{s,v} - Y_{s,v'} &= \sum_{i=1}^m s_i (\langle \sqrt{n}a_i, v \rangle^2 - \langle \sqrt{n}a_i, v' \rangle^2) \\ &= \sum_{i=1}^m s_i \langle \sqrt{n}a_i, v - v' \rangle \langle \sqrt{n}a_i, v + v' \rangle. \end{aligned}$$

We can bound the subexponential norm of each summand via

$$\begin{aligned} \|s_i \langle \sqrt{n}a_i, v - v' \rangle \langle \sqrt{n}a_i, v + v' \rangle\|_{\psi_1} &\leq s_i \|\langle \sqrt{n}a_i, v - v' \rangle\|_{\psi_2} \cdot \|\langle \sqrt{n}a_i, v + v' \rangle\|_{\psi_1} \\ &\leq C s_i \|v - v'\|_2. \end{aligned}$$

As such,

$$\sum_{i=1}^m \|s_i \langle \sqrt{n}a_i, v - v' \rangle \langle \sqrt{n}a_i, v + v' \rangle\|_{\psi_1}^2 \leq C \|v - v'\|_2^2 \sum_{i=1}^m s_i^2 \leq C(2\theta/\pi)m \|v - v'\|_2^2.$$

Applying Bernstein as before, we get

$$\mathbb{P}(|Y_{s,v} - Y_{s,v'}| \geq u) \leq 2 \exp(-c \min\{u^2/(2\theta/\pi)m \|v - v'\|_2^2, u/\|v - v'\|_2\}). \quad (5.6)$$

Now, recall the simple observation that for any numbers $a, b \in \mathbb{R}$, we have

$$\max\{|a|, |b|\} \leq |a| + |b| \leq 2 \max\{|a|, |b|\}.$$

As such, for any $u > 0$, given $s, s' \in \mathcal{S}_{2\theta/\pi}$, $v, v' \in B_2^n$, we have

$$\begin{aligned} \sqrt{u} \|(s, v) - (s', v')\|_2 + u \|(s, v) - (s', v')\|_1 &\geq \frac{1}{2} (\sqrt{u} \|s - s'\|_2 + \sqrt{u} \sqrt{2m\theta/\pi} \|v - v'\|_2 + u \|s - s'\|_\infty \\ &\quad + u \|v - v'\|_2) \\ &\geq \frac{1}{2} \max\{\sqrt{u} \|s - s'\|_2 + u \|s - s'\|_\infty, \\ &\quad \sqrt{u} \sqrt{2m\theta/\pi} \|s - s'\|_2 + u \|v - v'\|_2\}. \end{aligned}$$

Since

$$|Y_{s,v} - Y_{s',v'}| \leq |Y_{s,v} - Y_{s',v}| + |Y_{s',v} - Y_{s',v'}|,$$

we have that if

$$|Y_{s,v} - Y_{s',v'}| \geq c(\sqrt{u} \|(s, v) - (s', v')\|_2 + u \|(s, v) - (s', v')\|_1),$$

then either

$$|Y_{s,v} - Y_{s',v}| \geq \frac{c}{4}(\sqrt{u}\|s - s'\|_2 + u\|s - s'\|_\infty)$$

or

$$|Y_{s',v} - Y_{s',v'}| \geq \frac{c}{4}\left(\sqrt{u}\sqrt{2m\theta/\pi}\|v - v'\|_2 + u\|v - v'\|_2\right).$$

We can then combine the bounds (5.6) and (5.5) to get

$$\mathbb{P}(|Y_{s,v} - Y_{s',v'}| \geq c(\sqrt{u}\|(w,v) - (w',v')\|_2 + u\|(w,v) - (w',v')\|_1)) \leq 4e^{-u}.$$

Hence, the process $(Y_{s,v})$ satisfies the definition (5.4) for having mixed tail increments. \square

We next bound the γ_1 and γ_2 functions for $\mathcal{S}_{2\theta/\pi} \times B_2^n$.

Lemma 5.5. *We may bound the γ_1 functional of $\mathcal{S}_{2\theta/\pi} \times B_2^n$ by*

$$\gamma_1(\mathcal{S}_{2\theta/\pi} \times B_2^n, \|\cdot\|_1) \leq C((2\theta/\pi) \log(\pi/2\theta)m + n).$$

Proof. The proof of the bound uses metric entropy and a version of Dudley's inequality. Let (T, d) be a metric space, and for any $u > 0$, let $N(T, d, u)$ denote the covering number of T at scale u , i.e. the smallest number of radius u balls needed to cover T . Dudley's inequality (see [20]) states that there is an absolute constant C for which

$$\gamma_1(T, d) \leq C \int_0^\infty \log N(T, d, u) du. \quad (5.7)$$

Recall that $\mathcal{S}_{2\theta/\pi}$ is the set of all $\{0, 1\}$ vectors with fewer than $2\theta/\pi$ ones. For convenience, let us assume that $2m\theta/\pi$ is an integer. We then have the inclusion

$$\mathcal{S}_{2\theta/\pi} \subset \bigcup_{I \in \mathcal{I}} [0, 1]^I,$$

where \mathcal{I} is the collection of all subsets of $[m]$ of size $2m\theta/\pi$, and $[0, 1]^I$ denotes the unit cube in the coordinate set I . We may then also write

$$\mathcal{S}_{2\theta/\pi} \times B_2^n \subset \bigcup_{I \in \mathcal{I}} ([0, 1]^I \times B_2^n).$$

Note that a union of covers for each $[0, 1]^I \times B_2^n$ gives a cover for $\mathcal{S}_{2\theta/\pi} \times B_2^n$. This, together with the symmetry of $\|\cdot\|_\infty$ with respect to permutation of the coordinates gives

$$N(\mathcal{S}_{2\theta/\pi} \times B_2^n, \|\cdot\|_1, u) \leq |\mathcal{I}| \cdot N([0, 1]^I \times B_2^n, \|\cdot\|_1, u)$$

for some fixed index set I .

We next generalize the notion of covering numbers slightly. Given two sets T and K , we let $N(T, K)$ denote the number of translates of K needed to cover the set T . It is easy to see that we have $N(T, d, u) = N(T, uB_d)$, where B_d is the unit ball with respect to the metric d . Since the unit ball for $\|\cdot\|_1$ is $B_\infty^m \times B_2^n$, we therefore have

$$\begin{aligned} N([0, 1]^I \times B_2^n, \|\cdot\|_1, u) &= N([0, 1]^I \times B_2^n, u(B_\infty^m \times B_2^n)) \\ &\leq N(B_\infty^{(2\theta/\pi)m} \times B_2^n, u(B_\infty^{(2\theta/\pi)m} \times B_2^n)). \end{aligned}$$

Such a quantity can be bounded using a volumetric argument. Generally, for any centrally symmetric convex body K in \mathbb{R}^n , we have (see Corollary 4.1.15 in [1])

$$N(K, uK) \leq (3/u)^n. \quad (5.8)$$

This implies that

$$\log N([0, 1]^I \times S^{n-1}, \|\cdot\|_1, u) \leq \log(3/u)((2\theta/\pi)m + n).$$

Finally, observe that

$$\log |\mathcal{I}| = \log \binom{m}{(2\theta/\pi)m} \leq (2\theta/\pi)m \log(e\pi/2\theta).$$

We can thus plug these last two bounds into (5.7), noting that the integrand is zero for $u \geq 1$ to get

$$\begin{aligned} \gamma_1(\mathcal{S}_{2\theta/\pi} \times B_2^n, \|\cdot\|_1) &\leq C \int_0^1 (2\theta/\pi)m \log(e\pi/2\theta) + \log(3/u)((2\theta/\pi)m + n) du \\ &\leq C((2\theta/\pi) \log(\pi/2\theta)m + n) \end{aligned}$$

as was to be shown. \square

Lemma 5.6. *We may bound the γ_2 functional of $\mathcal{S}_{2\theta/\pi} \times B_2^n$ by*

$$\gamma_2(\mathcal{S}_{2\theta/\pi} \times B_2^n, \|\cdot\|_2) \leq C\sqrt{2\theta/\pi}(m + \sqrt{mn}).$$

Proof. Since $\alpha = 2$, we may appeal directly to the theory of Gaussian complexity [22]. However, since we have already introduced some of the theory of metric entropy in the previous lemma, we might as well continue down this path. In this case, we have the Dudley bound

$$\gamma_2(T, d) \leq C \int_0^\infty \sqrt{\log N(T, d, u)} du \quad (5.9)$$

for any metric space (T, d) .

Observe that the unit ball for $\|\cdot\|_2$ is $B_2^m \times (2m\theta/\pi)^{-1/2} B_2^n$. On the other hand, we conveniently have

$$\mathcal{S}_{2\theta/\pi} \times B_2^n \subset \sqrt{2m\theta/\pi} B_2^m \times B_2^n.$$

As such, we have

$$\begin{aligned} N(\mathcal{S}_{2\theta/\pi} \times B_2^n, \|\cdot\|_2, u) &\leq N(\sqrt{2m\theta/\pi} B_2^m \times B_2^n, \|\cdot\|_2, u) \\ &= N(\sqrt{2m\theta/\pi} B_2^m \times B_2^n, u(B_2^m \times (2m\theta/\pi)^{-1/2} B_2^n)) \\ &= N(T, (2m\theta/\pi)^{-1/2} uT), \end{aligned}$$

where $T = \sqrt{2m\theta/\pi} B_2^m \times B_2^n$.

Plugging this into (5.9) and subsequently using the volumetric bound (5.8), we get

$$\begin{aligned} \gamma_2(\mathcal{S}_{2\theta/\pi} \times B_2^n, \|\cdot\|_2) &\leq C \int_0^\infty \sqrt{\log N(T, (2m\theta/\pi)^{-1/2} uT)} du \\ &= C\sqrt{2m\theta/\pi} \int_0^\infty \sqrt{\log N(T, uT)} du \\ &\leq C\sqrt{2m\theta/\pi} \sqrt{m+n}, \end{aligned}$$

which is clearly equivalent to the bound that we want. \square

At this stage, we can put everything together to bound the supremum of our stochastic process.

Theorem 5.7 (Bound on supremum of $(Y_{s,v})$). *Let $(Y_{s,v})$ be the process defined in (5.3). Let $0 < \delta < 1/e$, let θ be an acute angle, and suppose $m \geq \max\{n, \log(1/\delta)\pi/2\theta\}$. Then with probability at least $1 - \delta$, the supremum of the process satisfies*

$$\sup_{s \in \mathcal{S}_{2\theta/\pi}, v \in B_2^n} Y_{s,v} \leq C\sqrt{2\theta/\pi} \cdot m \quad (5.10)$$

Proof. It is easy to see that we have

$$\text{diam}(\mathcal{S}_{2\theta/\pi} \times B_2^n, \|\cdot\|_1) = 2,$$

and

$$\text{diam}(\mathcal{S}_{2\theta/\pi} \times B_2^n, \|\cdot\|_2) = 2\sqrt{2m\theta/\pi}.$$

Also observe that we have $Y_{s,0} = 0$ for any $s \in \mathcal{S}_{2\theta/\pi}$.

Using these, together with the previous two lemmas bounding the γ_1 and γ_2 functionals, we may apply Lemma 5.3 to see that

$$\sup_{s \in \mathcal{S}_{2\theta/\pi}, v \in B_2^n} Y_{s,v} \leq C(((2\theta/\pi) \log(\pi/2\theta)m + n) + \sqrt{2\theta/\pi}(m + \sqrt{mn}) + u + \sqrt{u}\sqrt{2m\theta/\pi}).$$

with probability at least $1 - e^{-u}$.

Using our assumptions on m , we may simplify this bound to obtain (5.10). \square

Finally, we show that $\frac{1}{m} \sum_{i=1}^m a_i a_i^T$ is well-behaved.

Lemma 5.8. *Let $\delta > 0$. Then if $m \geq C(n + \sqrt{\log(1/\delta)})$, with probability at least $1 - \delta$, we have*

$$\left\| \frac{n}{m} \sum_{i=1}^m a_i a_i^T - I_n \right\| \leq 0.1$$

Proof. Note, as before, that the $\sqrt{n}a_i$'s are isotropic subgaussian random variables with subgaussian norm bounded by an absolute constant. The claim then follows immediately from Theorem 5.39 in [23], which itself is proved using Bernstein and a simple net argument. \square

Theorem 5.9 (Finite measurement sets satisfy ACW condition). *There is some $\theta_0 > 0$ and an absolute constant C such that for all angles $0 < \theta \leq \theta_0$, for all dimensions n , and any $\delta > 0$, if m satisfies*

$$m \geq C(\pi/2\theta)^2(n \log(m/n) + \log(1/\delta)), \quad (5.11)$$

then with probability at least $1 - \delta$, the measurement set A comprising m independent random vectors drawn uniformly from S^{n-1} satisfies the $\text{ACW}(\theta, \alpha)$ condition with $\alpha = 1/2$.

Proof. Fix $n, \delta > 0$. Choose θ_0 such that the constant C in the statement in Theorem 5.7 satisfies $C\sqrt{2\theta_0/\pi} \leq 0.1$. Fix $0 < \theta \leq \theta_0$, and let Ω_1 , Ω_2 , and Ω_3 denote the good events in Lemma 5.1, Theorem 5.7, and Lemma 5.8 with this choice of θ . Whenever m satisfies our assumption (5.11), the intersection of these events occurs with probability at least $1 - 3\delta$ by the union bound.

Let us condition on being in the intersection of these events. For any wedge $W \in \mathcal{W}_\theta$ (i.e. of angle less than θ), Lemma 5.1 tells us that its associated selector vector satisfies $s_{W,A} \in \mathcal{S}_{2\theta/\pi}$ (i.e. that it has at most $2m\theta/\pi$ ones,). By Theorem 5.7 and our assumption on θ_0 , we then have

$$\lambda_{\max} \left(\frac{1}{m} \sum_{i=1}^m a_i a_i^T 1_W(a_i) \right) \leq \frac{1}{nm} \sup_{s \in \mathcal{S}_{2\theta/\pi}, v \in B_2^n} Y_{s,v} \leq \frac{0.1}{n}.$$

On the other hand, Lemma 5.8 guarantees that

$$\lambda_{\min} \left(\frac{1}{m} \sum_{i=1}^m a_i a_i^T \right) \geq \frac{0.9}{n}.$$

Combining these, we get

$$\lambda_{\min} \left(\frac{1}{m} \sum_{i=1}^m a_i a_i^T - \frac{4}{m} \sum_{i=1}^m a_i a_i^T 1_W(a_i) \right) \geq \frac{1}{2n},$$

which was to be shown. \square

6 Proof and discussion of Theorem 1.2

We restate the theorem here for convenience.

Theorem 6.1. *Fix $\epsilon > 0$, $0 < \delta_1 \leq 1/2$, and $0 < \delta, \delta_2 \leq 1$. There are absolute constants $C, c > 0$ such that if*

$$m \geq C(n \log(m/n) + \log(1/\delta)),$$

then with probability at least $1 - \delta$, m sampling vectors selected uniformly and independently from the unit sphere S^{n-1} form a set such that the following holds: Let $x \in \mathbb{R}^n$ be a signal vector and let x_0 be an initial estimate satisfying $\|x_0 - x\|_2 \leq c\sqrt{\delta_1}\|x\|_2$. Then for any $\epsilon > 0$, if

$$K \geq 2(\log(1/\epsilon) + \log(2/\delta_2))n,$$

then the K -th step randomized Kaczmarz estimate x_K satisfies $\|x_K - x\|_2^2 \leq \epsilon \|x_0 - x\|_2^2$ with probability at least $1 - \delta_1 - \delta_2$.

Proof. Let A be our m by n measurement matrix. By Theorem 5.9, there is an angle θ_0 , and a constant C such that for $m \geq C(n \log(m/n) + \log(1/\delta))$, A is ACW($\theta_0, 1/2$) with probability at least $1 - \delta$.

We can then use Corollary 4.4 to guarantee that with probability at least $1 - \delta_1 - \delta_2$, running the randomized Kaczmarz update K times gives an estimate x_K satisfying

$$\|x_K - x\|_2^2 \leq \epsilon \|x_0 - x\|_2^2.$$

This completes the proof of the theorem. \square

Inspecting the statement of the theorem, we see that we can make the failure probability δ as small as possible by making m large enough. Likewise, we can do the same with δ_2 by adjusting K . Proposition B.1 shows that we can also make δ_2 smaller by increasing m . However, while the dependence of m and K on δ and δ_2 respectively is logarithmic, the dependence of m on δ_1 is *polynomial* (we need $m \gtrsim 1/\delta_1^2$). This is rather unsatisfactory, but can be overcome by a simple ensemble method. We encapsulate this idea in the following algorithm.

Algorithm 1 ENSEMBLE RANDOMIZED KACZMARZ

Input: Measurements b_1, \dots, b_m , sampling vectors a_1, \dots, a_m , relative error tolerance ϵ , iteration count K , trial count L .

Output: An estimate \hat{x} for x .

- 1: Obtain an initial estimate x_0 using the truncated spectral initialization method (see Appendix B).
 - 2: **for** $l = 1, \dots, L$, run K randomized Kaczmarz update steps starting from x_0 to obtain an estimate $x_K^{(l)}$.
 - 3: **for** $l = 1, \dots, L$, **do**
 - 4: **if** $|B(x_K^{(l)}, 2\sqrt{\epsilon}) \cap \{x_K^{(1)}, \dots, x_K^{(L)}\}| \geq L/2$
 - 5: **return** $\hat{x} := x_K^{(l)}$.
-

Proposition 6.2 (Guarantee for ensemble method). *Given the assumptions of Theorem 6.1, further assume that $\delta_1 + \delta_2 \leq 1/3$. For any $\delta' > 0$, there is an absolute constant C such that if $L \geq C \log(1/\delta')$, then the estimate \hat{x} given by Algorithm 1 satisfies $\|\hat{x} - x\|_2^2 \leq 9\epsilon \|x_0 - x\|_2^2$ with probability at least $1 - \delta'$.*

Proof. For $1 \leq l \leq L$, let χ_l be the indicator variable for $\|x_K^{(l)} - x\|_2^2 \leq \epsilon \|x_0 - x\|_2^2$. Then χ_1, \dots, χ_L are i.i.d. Bernoulli random variables each with success probability at least $2/3$. Let I be the set of indices l for which $\chi_l = 1$. Using a Chernoff bound [22], we see that with probability at least $1 - e^{-cL}$, $|I| \geq L/2$. Now let I' be the set of indices for which $|B(x_K^{(l)}, 2\epsilon) \cap \{x_K^{(1)}, \dots, x_K^{(L)}\}| \geq L/2$. Observe that for all $l, l' \in I$, we have

$$\|x_K^{(l)} - x_K^{(l')}\|_2 \leq \|x_K^{(l)} - x\|_2 + \|x - x_K^{(l')}\|_2 \leq 2\sqrt{\epsilon}.$$

This implies that $I \subset I'$, so $I' \neq \emptyset$. Furthermore, for all $l' \in I'$, there is $l \in I$ for which $\|x_K^{(l)} - x_K^{(l')}\|_2 \leq 2\sqrt{\epsilon}$. As such, we have

$$\|x_K^{(l')} - x\|_2 \leq \|x_K^{(l')} - x_K^{(l)}\|_2 + \|x_K^{(l)} - x\|_2 \leq 3\sqrt{\epsilon}.$$

Now, observe that the estimate \hat{x} returned by Algorithm 1 is precisely some $x_K^{(l')}$ for which $l' \in I'$. This shows that on the good event, we indeed have $\|\hat{x} - x\|_2^2 \leq 9\epsilon \|x_0 - x\|_2^2$. By our assumption on L , we see that the failure probability is bounded by δ' . \square

In practice however, the ensemble method is not required. Numerical experiments show that the randomized Kaczmarz method always eventually converges from any initial estimate.

7 Extensions

7.1 Arbitrary initialization

In order to obtain a convergence guarantee, we used a truncated spectral initialization to obtain an initial estimate before running randomized Kaczmarz updates. Since the number of steps that we require is only linear in the dimension, and each step requires only linear time, the iteration phase of the algorithm only requires $O(n^2)$ time, and furthermore does not need to see all the data in order to start running.

The spectral initialization on the other hand requires one to see all the data. Forming the matrix from which we obtain the estimate involves adding m rank 1 matrices, and hence naively requires $O(mn^2)$ time. There is hence an incentive to do away with this step altogether, and ask whether the randomized Kaczmarz algorithm works well even if we start from an arbitrary initialization.

We have some numerical evidence that this is indeed true, at least for real Gaussian measurements. Unfortunately, we do not have any theoretical justification for this phenomenon, and it will be interesting to see if any results can be obtained in this direction.

7.2 Complex Gaussian measurements

We have proved our main results for measurement systems comprising random vectors drawn independently and uniformly from the sphere, or equivalently, for real Gaussian measurements. These are not the measurement sets that are used in practical applications, which often deal with imaging and hence make use of complex measurements.

While most theoretical guarantees for phase retrieval algorithms are in terms of real Gaussian measurements, some also hold for complex Gaussian measurements, even with identical proofs. This is the case for *PhaseMax* [5] and for Wirtnger flow [4]. We believe that a similar situation should hold for the randomized Kaczmarz method, but are not yet able to recalibrate our tools to handle the complex setting.

It is easy to adapt the randomized Kaczmarz update formula (1.3) itself: we simply replace the sign of $\langle a_{r(k)}, x_{k-1} \rangle$ with its phase (i.e. $\frac{\langle a_{r(k)}, x_{k-1} \rangle}{|\langle a_{r(k)}, x_{k-1} \rangle|}$). Numerical experiments also show that convergence does occur for complex Gaussian measurements (and even CDP measurements) [25]. Nonetheless, in trying to adapt the proof to this situation, we meet an obstacle at the first step: when computing the error term, we can no longer simply sum up the influence of “bad measurements” as we did in Lemma 2.1. Instead, *every* term contributes an error that scales with the phase difference

$$\frac{\langle a_i, z \rangle}{|\langle a_i, z \rangle|} - \frac{\langle a_i, x \rangle}{|\langle a_i, x \rangle|}.$$

We leave it to future work to prove convergence in this setting, whether by adapting our methods, or by proposing completely new ones.

7.3 Deterministic constructions of measurement sets

The theory that we have developed in this paper does not apply solely to Gaussian measurements, and generalizes to any measurement sets that satisfy the ACW condition that we introduced in Section 5. It will be interesting to investigate what natural classes of measurement sets satisfy this condition.

Acknowledgments

Y.T. was partially supported by the Juha Heinonen Memorial Graduate Fellowship at the University of Michigan. R.V. was partially supported by NSF Grant DMS 1265782 and U.S. Air Force Grant FA9550-18-1-0031. Y.T. would like to thank Halyun Jeong for insightful discussions on this topic. We would also like to thank the anonymous reviewers for their many helpful comments.

References

- [1] ARTSTEIN-AVIDAN, S., GIANNOPOULOS, A., AND MILMAN, V. D. *Asymptotic Geometric Analysis, Part I*. AMS, 2015.
- [2] BAHMANI, S., AND ROMBERG, J. Phase Retrieval Meets Statistical Learning Theory: A Flexible Convex Relaxation. Preprint, available at <http://arxiv.org/abs/1610.04210>.
- [3] BLUMER, A., EHRENFEUCHT, A., HAUSSLER, D., AND WARMUTH, M. K. Learnability and the Vapnik-Chervonenkis dimension. *Journal of the ACM* 36, 4 (1989), 929–965.
- [4] CANDES, E. J., LI, X., AND SOLTANOLKOTABI, M. Phase Retrieval via Wirtinger Flow: Theory and Algorithms. *IEEE Transactions on Information Theory* 61, 4 (2015), 1985–2007.
- [5] CANDES, E. J., STROHMER, T., AND VORONINSKI, V. PhaseLift: Exact and stable signal recovery from magnitude measurements via convex programming. *Communications on Pure and Applied Mathematics* 66, 8 (2013), 1241–1274.
- [6] CHEN, Y., AND CANDES, E. Solving random quadratic systems of equations is nearly as easy as solving linear systems. In *Advances in Neural Information Processing Systems* 28. 2015, pp. 739–747.
- [7] CHI, Y., AND LU, Y. M. Kaczmarz Method for Solving Quadratic Equations. *IEEE Signal Processing Letters* 23, 9 (2016), 1183–1187.
- [8] DAVIS, C., AND KAHAN, W. M. The rotation of eigenvectors by a perturbation. III. *SIAM Journal on Numerical Analysis* 7, 1 (1970), 1–46.
- [9] DAVIS, D., DRUSVYATSKIY, D., AND PAQUETTE, C. The nonsmooth landscape of phase retrieval.
- [10] DIRKSEN, S. Tail bounds via generic chaining. *Electronic Journal of Probability* 20 (2015), 1–29.
- [11] DUCHI, J. C., AND RUAN, F. Solving (most) of a set of quadratic equalities: Composite optimization for robust phase retrieval. Preprint, available at <http://arxiv.org/abs/1705.02356>.
- [12] ELDAR, Y. C., AND MENDELSON, S. Phase retrieval: Stability and recovery guarantees. *Applied and Computational Harmonic Analysis* 36, 3 (2014), 473–494.
- [13] FIENUP, J. R. Phase retrieval algorithms: a comparison. *Appl. Opt.* 21, 15 (1982), 2758.
- [14] GOLDSTEIN, T., AND STUDER, C. PhaseMax: Convex Phase Retrieval via Basis Pursuit. Preprint, available at <http://arxiv.org/abs/1610.07531>.

- [15] HAND, P., AND VORONINSKI, V. An Elementary Proof of Convex Phase Retrieval in the Natural Parameter Space via the Linear Program PhaseMax. Preprint, available at <http://arxiv.org/abs/1611.03935>.
- [16] JEONG, H., AND GÜNTÜRK, C. S. Convergence of the randomized Kaczmarz method for phase retrieval. Preprint, available at <http://arxiv.org/abs/1706.10291>.
- [17] MOHRI, M., ROSTAMIZADEH, A., AND TALWALKAR, A. *Foundations of Machine Learning*. MIT press, 2012.
- [18] STROHMER, T., AND VERSHYNIN, R. A randomized kaczmarz algorithm with exponential convergence. *Journal of Fourier Analysis and Applications* 15, 2 (2009), 262–278.
- [19] SUN, J., QU, Q., AND WRIGHT, J. A Geometric Analysis of Phase Retrieval. *Foundations of Computational Mathematics* (2017).
- [20] TALAGRAND, M. *The Generic Chaining: Upper and Lower Bounds of Stochastic Processes*. Springer Monographs in Mathematics. Springer Berlin Heidelberg, 2005.
- [21] VAPNIK, V. N., AND CHERVONENKIS, A. Y. On the Uniform Convergence of Relative Frequencies of Events to Their Probabilities. *Theory of Probability & Its Applications* 16, 2 (Jan 1971), 264–280.
- [22] VERSHYNIN, R. *High-Dimensional Probability*. Book draft, available at <https://www.math.uci.edu/~rvershyn/papers/HDP-book/HDP-book.html>.
- [23] VERSHYNIN, R. Introduction to the non-asymptotic analysis of random matrices. In *Compressed Sensing*, Y. C. Eldar and G. Kutyniok, Eds. Cambridge University Press, Cambridge, 2011, pp. 210–268.
- [24] WANG, G., GIANNAKIS, G. B., AND ELDAR, Y. C. Solving Systems of Random Quadratic Equations via Truncated Amplitude Flow. *IEEE Transactions on Signal Processing* 65, 8 (May 2016), 1961–1974.
- [25] WEI, K. Solving systems of phaseless equations via Kaczmarz methods: a proof of concept study. *Inverse Problems* 31, 12 (Dec 2015), 125008.
- [26] ZHANG, H., AND LIANG, Y. Reshaped wirtinger flow for solving quadratic system of equations. In *Advances in Neural Information Processing Systems* 29. 2016, pp. 2622–2630.

A Growth functions and VC dimension

In this section, we define growth functions and VC dimension. We also state some standard results on these topics that we require for our proofs in Section 5. We refer the interested reader to [17] for a more in-depth exposition on these topics.

Let \mathcal{X} be a set and \mathcal{C} be a family of subsets of \mathcal{X} . For a given set $C \in \mathcal{C}$, we slightly abuse notation and identify it with its indicator function $1_C: \mathcal{X} \rightarrow \{0, 1\}$.

The *growth function* $\Pi_{\mathcal{C}}: \mathbb{N} \rightarrow \mathbb{R}$ of \mathcal{C} is defined via

$$\Pi_{\mathcal{C}}(m) := \max_{x_1, \dots, x_m \in \mathcal{X}} |\{(C(x_1), C(x_2), \dots, C(x_m)) : C \in \mathcal{C}\}|.$$

Meanwhile, the *VC dimension* of \mathcal{C} is defined to be the largest integer m for which $\Pi_{\mathcal{C}}(m) = 2^m$. These two concepts are fundamental to statistical learning theory. The key connection between them is given by the Sauer-Shelah lemma.

Lemma A.1 (Sauer-Shelah, Corollary 3.3 in [17]). *Let \mathcal{C} be a collection of subsets of VC dimension d . Then for all $m \geq d$, have*

$$\Pi_{\mathcal{C}}(m) \leq \left(\frac{em}{d}\right)^d.$$

The reason why we are interested in the growth function of a family of subsets \mathcal{C} is because we have the following guarantee for the uniform convergence for the empirical measures of sets belonging to \mathcal{C} .

Proposition A.2 (Uniform deviation, Theorem 2 in [21]). *Let \mathcal{C} be a family of subsets of a set \mathcal{X} . Let μ be a probability measure on \mathcal{X} , and let $\hat{\mu}_m := \frac{1}{m} \sum_{i=1}^m \delta_{X_i}$ be the empirical measure obtained from m independent copies of a random variable X with distribution μ . For every u such that $m \geq 2/u^2$, the following deviation inequality holds:*

$$\mathbb{P}(\sup_{C \in \mathcal{C}} |\hat{\mu}_m(C) - \sigma(C)| \geq u) \leq 4\Pi_{\mathcal{C}}(2m) \exp(-mu^2/16). \quad (\text{A.1})$$

We now state and prove two simple claims.

Claim A.3. *Let \mathcal{C} be the collection of all hemispheres in S^{n-1} . Then the VC dimension of \mathcal{C} is bounded from above by $n + 1$.*

Proof. It is a standard fact from statistical learning theory [17] that the VC dimension of half-spaces in \mathbb{R}^n is $n + 1$. Since S^{n-1} is a subset of \mathbb{R}^n , the claim follows by the definition of VC dimension. \square

Claim A.4. *Let \mathcal{C} and \mathcal{D} be two collections of functions from a set \mathcal{X} to $\{0, 1\}$. Using Δ to denote symmetric difference, we define*

$$\mathcal{C} \Delta \mathcal{D} := \{C \Delta D \mid C \in \mathcal{C}, D \in \mathcal{D}\}. \quad (\text{A.2})$$

Then the growth function $\Pi_{\mathcal{C} \Delta \mathcal{D}}$ of $\mathcal{C} \Delta \mathcal{D}$ satisfies $\Pi_{\mathcal{C} \Delta \mathcal{D}}(m) \leq \Pi_{\mathcal{C}}(m) \cdot \Pi_{\mathcal{D}}(m)$ for all $m \in \mathbb{Z}_+$.

Proof. Fix m , and points $x_1, \dots, x_m \in \mathcal{X}$. Then every possible configuration $(f(x_1), f(x_2), \dots, f(x_m))$ arising from some $f \in \mathcal{C} \Delta \mathcal{D}$ is the point-wise symmetric difference

$$(f(x_1), f(x_2), \dots, f(x_m)) = (C(x_1), C(x_2), \dots, C(x_m)) \Delta (D(x_1), D(x_2), \dots, D(x_m))$$

of configurations arising from some $C \in \mathcal{C}$ and $D \in \mathcal{D}$. By the definition of growth functions, there are at most $\Pi_{\mathcal{C}}(m) \cdot \Pi_{\mathcal{D}}(m)$ pairs of these configurations, from which the bound follows. \square

Remark A.5. There is an extensive literature on how to bound the VC dimension of concept classes that arise from finite intersections or unions of those from a known collection of concept classes, each of which has bounded VC dimension. We won't require this much sophistication here, and refer the reader to [3] for more details.

B Initialization

Several different schemes have been proposed for obtaining initial estimates for *PhaseMax* and gradient descent methods for phase retrieval. Surprisingly, these are all spectral in nature: the initial estimate x_0 is obtained as the leading eigenvector to a matrix that is constructed out of the sampling vectors a_1, \dots, a_m and their associated measurements b_1, \dots, b_m [4, 6, 26, 24].

There seems to be empirical evidence, at least for Gaussian measurements, that the best performing method is the *orthogonality-promoting method* of [24]. Nonetheless, for any given relative error tolerance, all the methods seem to require sample complexity of the same order. Hence, we focus on the *truncated spectral method* of [6] for expositional clarity, and refer the reader to the respective papers on the other methods for more details.

The truncated spectral method initializes $x_0 := \lambda_0 \tilde{x}_0$, where $\lambda_0 = \sqrt{\frac{1}{m} \sum_{i=1}^m b_i^2}$, and \tilde{x}_0 is the leading eigenvector of

$$Y = \frac{1}{m} \sum_{i=1}^m b_i^2 a_i a_i^T 1(b_i \leq 3\lambda_0).$$

Note that when constructing Y , we sum up only those sampling vectors whose corresponding measurements satisfy $b_i \leq 3\lambda_0$. The point of this is to remove the influence of unduly large measurements, and allow for good concentration estimates, as we shall soon demonstrate.

Suppose from now on that the a_i 's are independent standard Gaussian vectors. In [6], the authors prove that with probability at least $1 - \exp(-\Omega(m))$, we have $\|\tilde{x}_0 - x\|_2 \leq \epsilon \|x\|_2$ for any fixed relative error tolerance ϵ (see their Proposition 3). They do not, however, examine the dependence of the probability bound on ϵ . Nonetheless, by examining the proof more carefully, we can make this dependence explicit. In doing so, we obtain the following proposition.

Proposition B.1 (Relative error guarantee for initialization). *Let $a_1, \dots, a_m, b_1, \dots, b_m, Y$ and x_0 be defined as in the preceding discussion. Fix $\epsilon > 0$ and $0 < \delta < 1$. Then with probability at least $1 - \delta$, we have $\|x_0 - x\|_2 \leq \epsilon \|x\|_2$ so long as $m \geq C(\log(1/\delta) + n)/\epsilon^2$.*

Proof. We simply make the following observations while following the proof of Proposition 3 in [6]. First, since all quantities are 2-homogeneous in $\|x\|_2$, we may assume without loss of generality that $\|x\|_2 = 1$. Next, there is some absolute constant c such that if we define Y_1 and Y_2 by choosing $\gamma_1 = 3 + c\epsilon$, $\gamma_2 = 3 - c\epsilon$, we have the bound $\|\mathbb{E}Y_1 - \mathbb{E}Y_2\| \leq C\epsilon$. Note also that the deviation estimates $\|Y_1 - \mathbb{E}Y_1\|$, $\|Y_2 - \mathbb{E}Y_2\|$ are bounded by $C\epsilon$ given our assumptions on m . This implies that with high probability,

$$\|Y - \beta_1 x x^T - \beta_2 I_n\| \leq C\epsilon.$$

Adjust our constants so that C in the last equation is bounded by $\beta_1 - \beta_2$. We may then apply Davis-Kahan [8] to get

$$\|\tilde{x}_0 - x\|_2 \leq \frac{\|Y - \beta_1 x x^T - \beta_2 I_n\|}{\beta_1 - \beta_2} \leq \epsilon$$

as we wanted. \square

By examining the proof carefully, the astute reader will observe that the crucial properties that we used were the rotational invariance of the a_i 's (to compute the formulas for $\mathbb{E}Y_1$ and $\mathbb{E}Y_2$) and their subgaussian tails (to derive the deviation estimates). These properties also hold for sampling vectors that are uniformly distributed on the sphere. As such, a more lengthy and

tedious calculation can be done to show that the guarantee also holds for such sampling vectors. If the reader has any residual doubt, perhaps this can be assuaged by noting that a uniform sampling vector and its associated measurement (a_i, b_i) can be turned into an honest real Gaussian vector by multiplying both quantities by an independent χ^2 random variable with n degrees of freedom.