Subspace estimation under coarse quantization

Sjoerd Dirksen
Mathematical Institute
Utrecht University
Utrecht, Netherlands
s.dirksen@uu.nl

Weilin Li
Department of Mathematics
City College of New York
NY, USA
wli6@ccny.cuny.edu

Johannes Maly
Department of Mathematics, LMU Munich and
Munich Center for Machine Learning (MCML)
Munich, Germany
maly@math.lmu.de

***Abstract*—We study subspace estimation from coarsly quantized data. In particular, we analyze two stochastic quantization schemes which use dithering: a one-bit quantizer combined with rectangular dither and a multi-bit quantizer with triangular dither. For each quantizer, we derive rigorous high probability bounds for the distances between the true and estimated signal subspaces. Using our analysis, we identify scenarios in which subspace estimation via triangular dithering qualitatively outperforms rectangular dithering. We verify in numerical simulations that our estimates are optimal in their dependence on the smallest non-zero eigenvalue of the target matrix.**

***Index Terms*—Subspace estimation, quantization, DOA estimation**

I. Introduction

The question of how to estimate an underlying subspace from noisy data appears in various contexts, e.g., direction-of-arrivals (DOAs) estimation [23]. Two state-of-the-art methods for DOA estimation, MUSIC [24] and ESPRIT [22] require such an estimate as initial step. To be more precise, given unknown signals $\boldsymbol{x}_1, \dots, \boldsymbol{x}_n \in \mathbb{C}^p$ that are concentrated around an s-dimensional subspace $\boldsymbol{U} \subseteq \mathbb{C}^p$, e.g., complex p-variate Gaussians with covariance matrix of rank s and leading eigenspace $\boldsymbol{U}$, the task is to estimate $\boldsymbol{U}$ from noisy observations

$$\boldsymbol{y}_k = \boldsymbol{x}_k + \boldsymbol{e}_k, \qquad \text{for } k \in [n], \tag{I.1}$$

where the vectors $\boldsymbol{e}_k \in \mathbb{C}^p$ model random noise. Here, we focus on the setting where the entries of $\boldsymbol{e}_k$ are i.i.d. samples from a fixed subgaussian distribution with mean zero and variance ν^2, e.g., complex Gaussian with i.i.d. entries.

Defining the observation matrix $\boldsymbol{Y}_n = [\boldsymbol{y}_1, \dots, \boldsymbol{y}_n] \in \mathbb{C}^{p\times n}$, a natural estimator $\widehat{\boldsymbol{U}}$ of $\boldsymbol{U}$ is obtained by computing the left singular vectors of a truncated SVD of $\boldsymbol{Y}_n$, or equivalently, the leading eigenvectors of the sample covariance matrix $\frac{1}{n}\boldsymbol{Y}_n\boldsymbol{Y}_n^*$. The estimator $\widehat{\boldsymbol{U}}$ is consistent because the calculation

$$\mathbb{E}(\boldsymbol{y}_k\boldsymbol{y}_k^*) = \mathbb{E}(\boldsymbol{x}_k\boldsymbol{x}_k^*) + \nu^2\boldsymbol{I}_p,$$

shows that on average, the noise shifts the eigenvalues of $\frac{1}{n}\boldsymbol{Y}_n\boldsymbol{Y}_n^*$, but not their eigenspaces. In the finite sample setting, the performance of $\widehat{\boldsymbol{U}}$ can be measured in terms of a suitable subspace angle distance $\operatorname{dist}(\widehat{\boldsymbol{U}}, \boldsymbol{U})$. Combining well-known results on covariance estimation [17] with the Davis-Kahan theorem [8] then leads to non-asymptotic error bounds for $\operatorname{dist}(\widehat{\boldsymbol{U}}, \boldsymbol{U})$ in terms of the oversampling ratio n/s. Interestingly, if the noise $\boldsymbol{e}_k$ is drawn from certain distributions and has i.i.d. entries, a better error bound for $\operatorname{dist}(\widehat{\boldsymbol{U}}, \boldsymbol{U})$ was derived in [3].

Since applications like DOA estimation normally lie at the interface of analog and digital domain, an additional challenge is to take into account the impact of analog-to-digital conversion, i.e., quantization [13] of observations. In this case one has only access to quantized observations $Q(\boldsymbol{y}_k)$ for $k \in [n]$, instead of the analog $\boldsymbol{y}_k$. Hereby Q is a suitably designed *quantizer*. We restrict ourselves in this work to *memoryless scalar quantization*, i.e., $Q \colon \mathbb{C} \to \mathcal{A}$ is a univariate function from $\mathbb{C}$ to a finite alphabet $\mathcal{A} \subseteq \mathbb{C}$ that is applied entry-wise to vectors in $\mathbb{C}^p$. Especially in view of modern large scale applications, there has been growing interest in coarse quantization schemes for which the number of bits per scalar is small. Several works [1], [30], [16], [19], [26], [31], [25], [18] have been examining DOA estimation from samples that are quantized by the simplest coarse quantizer $Q(\cdot) = \operatorname{sign}(\cdot)$. While this quantizer is cheap to implement, it looses scaling information and allows recovery only under restrictive assumptions. A second, more recent line of works [20], [29] mitigates the shortcomings of this simple quantizing scheme by considering dithered sign-quantizers of the shape $Q(\cdot) = \operatorname{sign}(\cdot + \tau)$ where the dither τ is designed noise. Whereas [20] examines Gaussian noise as dither, the authors of [29] follow the ideas in [12] and use uniform dithering noise together with two-bit observations per sample.

Dithering has a long history in signal processing [21]. In particular, the last decade showed substantial progress in deriving rigorous non-asymptotic performance guarantees for signal reconstruction from one-bit measurements [2], [9]. Only recently the first non-asymptotic (and near-minimax optimal) guarantees for estimating covariance matrices from one-bit samples with uniform dither have been derived [12]. In [28], the results of [12] have been generalized to the complex domain and applied to massive MIMO; in [11] a data-adaptive variant of the estimator in [12] has been developed. A recent work [6] modified the strategy of [12] to cover heavy-tailed distributions (by using truncation before quantizing). The work [4] introduced the idea of using triangular (rather than uniform) dithering.

This is an excerpt of the recent journal article [10]. All proof details, applications of the resulting bounds to DOA estimation, and further numerical studies will be contained therein.

A. Contribution

In this work, we analyze subspace identification from samples that are collected via a dithered quantizer. In particular,

(i) we provide non-asymptotic bounds for $\operatorname{dist}(\widehat{\boldsymbol{U}}, \boldsymbol{U})$ when $\widehat{\boldsymbol{U}}$ is obtained from quantized observations following the quantization model with uniformly distributed dither as proposed in [12], [29], see Theorem II.2. Due to the characteristic shape of the dither density, we refer to this setting as *rectangular dithering*, cf. Figure 1.

(ii) we show that, for low noise and higher bit-rates, *triangular dithering* [14] yields superior error bounds on $\operatorname{dist}(\widehat{\boldsymbol{U}}, \boldsymbol{U})$, see Theorem II.6. Here, the dither τ is a sum of two independent copies of a uniform random variable, and the name triangular dithering again stems from the characteristic shape of the dither density, cf. Figure 1. *The superior performance of triangular dithering in subspace estimation is particularly surprising since it stands in stark contrast to recently derived covariance estimation bounds.* Indeed, for covariance estimation the operator norm error bounds under rectangular dithering [12] and triangular dithering [5] are comparable.

(iii) we provide numerical simulations to verify that the theoretical error bounds are sharp, see Section III.

We are not aware of any comparable work that provides non-asymptotic error bounds for subspace estimation under coarse quantization. Indeed, almost all previously mentioned studies on DOA estimation from coarsely quantized samples are of empirical nature, proposing algorithmic approaches to the problem and evaluating their performance in simulations. The only exceptions are [25], [29], and [15]: The authors of [25] focus on one-bit DOA estimation via Sparse Linear Arrays. They provide conditions under which the identifiability of source DOAs from unquantized data is equivalent to the one of one-bit data. Furthermore, they provide a Cramér Rao bound analysis and a MUSIC-based reconstruction approach with asymptotic error guarantees. The authors of [29] build upon the idea of [27] in which DOA estimation from undithered one-bit samples is performed by Learned Iterative Soft-Thresholding (LISTA), i.e., unfolding and training ISTA as a network. By adding a uniform dither to the quantization model and relying on the theoretical analysis of the corresponding quantized covariance estimator in [28], they can derive performance guarantees for the one-bit LISTA approach from the results in [7]. The authors of [15] used ESPRIT for the quantized single-snapshot DOA problem, where information is only collected at one time instance. This setting cannot be treated using statistical methods. Instead, they exploited analytic properties of the Fourier transform by using a two-bit beta-quantization method, which is very different from dithered quantization.

In the extended version of this paper [10], we furthermore generalize the theory developed in [3] and apply the presented results to derive novel performance bounds for ESPRIT for DOA estimation.

B. Notation

We write $[n] = \{1, ..., n\}$ for $n \in \mathbb{N}$ and let $\mathbb{F} \in \{\mathbb{R}, \mathbb{C}\}$ be a placeholder for any of the two fields. We use the notation $a \lesssim_\alpha b$ (resp. $\gtrsim_\alpha$) to abbreviate $a \leq C_\alpha b$ (resp. $\geq$), for a constant $C_\alpha > 0$ depending only on α. Similarly, we write $a \lesssim b$ if $a \leq Cb$ for an absolute constant $C > 0$. We write $a \simeq b$ if both $a \lesssim b$ and $b \lesssim a$ hold (with possibly different implicit constants). Whenever we use absolute constants $c, C > 0$, their values may vary from line to line.

We let scalar-valued functions act entry-wise on vectors and matrices. In particular, the real-valued sign function is given by

$$[\operatorname{sign}_{\mathbb{R}}(\boldsymbol{x})]_i = \begin{cases} 1 & \text{if } x_i \geq 0 \\ -1 & \text{if } x_i < 0, \end{cases}$$

for all $\boldsymbol{x} \in \mathbb{R}^p$ and $i \in [p]$. We furthermore define the complex-valued sign-function by

$$\operatorname{sign}_{\mathbb{C}}(\boldsymbol{z}) = \operatorname{sign}_{\mathbb{R}}(\operatorname{Re}(\boldsymbol{z})) + \mathrm{i} \operatorname{sign}_{\mathbb{R}}(\operatorname{Im}(\boldsymbol{z})) \in \{\pm 1 \pm \mathrm{i}\}^p,$$

for any $\boldsymbol{z} \in \mathbb{C}^p$. Note that $\operatorname{sign}(\boldsymbol{x})_{\mathbb{R}} \neq \operatorname{sign}_{\mathbb{C}}(\boldsymbol{x})$, for $\boldsymbol{x} \in \mathbb{R}^p$ (since for instance $\operatorname{sign}_{\mathbb{R}}(0) \neq \operatorname{sign}_{\mathbb{C}}(0)$). Whenever we make a statement regarding the space $\mathbb{F}^p$ where $\mathbb{F} \in \{\mathbb{R}, \mathbb{C}\}$, we use $\operatorname{sign}_{\mathbb{F}}$ to refer to the respective sign-function. We abbreviate the squared modulus of the sign function by $c_{\mathbb{F}} := |\operatorname{sign}_{\mathbb{F}}(\cdot)|^2$, i.e., $c_{\mathbb{R}} = 1$ and $c_{\mathbb{C}} = 2$.

For $\boldsymbol{Z} \in \mathbb{F}^{p \times p}$, we denote the operator norm by $\|\boldsymbol{Z}\| = \sup_{\boldsymbol{u} \in \mathbb{S}^{p-1}} \|\boldsymbol{Z}\boldsymbol{u}\|_2$ and the entry-wise max-norm by $\|\boldsymbol{Z}\|_\infty = \max_{i,j} |Z_{i,j}|$. For $m \geq n$, we let $\mathbb{O}^{m \times n}$ be the set of $m \times n$ matrices over $\mathbb{F}$ whose columns are orthonormal. For a Hermitian matrix $\boldsymbol{A} \in \mathbb{C}^{m \times m}$, we let $\lambda_1(\boldsymbol{A}) \geq \lambda_2(\boldsymbol{A}) \geq \cdots \geq \lambda_m(\boldsymbol{A})$ denote its eigenvalues in non-increasing order. For a general matrix $\boldsymbol{B} \in \mathbb{C}^{m \times n}$, we let $\sigma_k(\boldsymbol{B}) := \sqrt{\lambda_k(\boldsymbol{B}^*\boldsymbol{B})}$ be its k-th largest singular value. Its r-th condition number is denoted by $\kappa_r(\boldsymbol{B}) := \sigma_1(\boldsymbol{B})/\sigma_r(\boldsymbol{B}) \in [1, \infty]$. For Hermitian $\boldsymbol{W}, \boldsymbol{Z} \in \mathbb{C}^{p \times p}$ we write $\boldsymbol{W} \preceq \boldsymbol{Z}$ if $\boldsymbol{Z} - \boldsymbol{W}$ is positive semidefinite.

We denote a real Gaussian random vector with mean $\boldsymbol{\mu} \in \mathbb{R}^p$ and covariance matrix $\boldsymbol{\Sigma} \in \mathbb{R}^{p \times p}$ by $\boldsymbol{x} \sim \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ and a circularly symmetric complex Gaussian random vector with covariance matrix $\boldsymbol{\Sigma} \in \mathbb{C}^{p \times p}$ by $\boldsymbol{z} \sim \mathcal{CN}(\boldsymbol{0}, \boldsymbol{\Sigma})$. If $\boldsymbol{\mu} = \boldsymbol{0}$ and $\boldsymbol{\Sigma} = \boldsymbol{I}_p$, we call $\boldsymbol{x}$ resp. $\boldsymbol{z}$ (complex) standard normal. The subgaussian (ψ_2-)norm of a random variable X is defined by

$$\|X\|_{\psi_2} = \inf\left\{t > 0 \colon \mathbb{E}\, e^{|X|^2/t^2} \leq 2\right\}.$$

A mean-zero random vector $\boldsymbol{y}$ in $\mathbb{R}^p$ is called K-subgaussian if

$$\|\langle \boldsymbol{y}, \boldsymbol{x}\rangle\|_{\psi_2} \leq K \|\langle \boldsymbol{y}, \boldsymbol{x}\rangle\|_{L_2} \quad \text{for all } \boldsymbol{x} \in \mathbb{R}^p.$$

We call a random vector $\boldsymbol{y} \in \mathbb{C}^p$ K-subgaussian if both $\operatorname{Re}(\boldsymbol{y})$ and $\operatorname{Im}(\boldsymbol{y})$ are subgaussian. For $p \geq 1$, we denote by $\mathcal{U}_{\mathbb{R}}[a, b]^p$ the uniform distribution on the rectangle $[a, b]^p$ and by $\mathcal{U}_{\mathbb{C}}[a, b]^p$ the uniform distribution on the set $[a, b]^p + \mathrm{i}[a, b]^p$, i.e., $X \sim \mathcal{U}_{\mathbb{C}}[a, b]^p$ if and only if $\operatorname{Re}(X) \sim \mathcal{U}_{\mathbb{R}}[a, b]^p$ and $\operatorname{Im}(X) \sim \mathcal{U}_{\mathbb{R}}[a, b]^p$ are independent.

II. Subspace identification from quantized observations

Given the observation model in (I.1), we will work with the following set of assumptions which is used in various applications.

Assumption II.1. Fix $n, p \in \mathbb{N}$.

(i) Let $\boldsymbol{x}_1, \dots, \boldsymbol{x}_n \in \mathbb{F}^p$ be deterministic (d) or stochastic (s) vectors. We then set $\boldsymbol{X}_n = \frac{1}{\sqrt{n}}(\boldsymbol{x}_1, \dots, \boldsymbol{x}_n)$ and either
  (d) set $\boldsymbol{\Sigma}_{\boldsymbol{x}} := \boldsymbol{X}_n \boldsymbol{X}_n^* = \frac{1}{n}\sum_{k=1}^n \boldsymbol{x}_k \boldsymbol{x}_k^*$ or
  (s) assume $\boldsymbol{x}_1, \dots, \boldsymbol{x}_n \overset{\text{i.i.d.}}{\sim} \boldsymbol{x}$, where $\boldsymbol{x} \in \mathbb{F}^p$ is K-subgaussian, and set $\boldsymbol{\Sigma}_{\boldsymbol{x}} := \mathbb{E}(\boldsymbol{x}\boldsymbol{x}^*)$.

(ii) Let $\boldsymbol{e}_1, \dots, \boldsymbol{e}_n \overset{\text{i.i.d.}}{\sim} \boldsymbol{e}$, where $\boldsymbol{e} \in \mathbb{F}^p$ is K-subgaussian with uncorrelated entries that are mean-zero and have variance ν^2. If $\mathbb{F} = \mathbb{C}$, we assume that the real and imaginary parts of the entries of $\boldsymbol{e}$ are independent.

(iii) Let $\boldsymbol{x}_1, \dots, \boldsymbol{x}_n, \boldsymbol{e}_1, \dots, \boldsymbol{e}_n$ be independent.

We define $\boldsymbol{\Sigma}_{\boldsymbol{y}} := \boldsymbol{\Sigma}_{\boldsymbol{x}} + \boldsymbol{\Sigma}_{\boldsymbol{e}}$ with $\boldsymbol{\Sigma}_{\boldsymbol{e}} := \mathbb{E}(\boldsymbol{e}\boldsymbol{e}^*)$. To measure the distance between subspaces, we will rely on the sine-theta distance [8] which is defined for two s-dimensional subspaces of $\mathbb{F}^p$ with bases $\boldsymbol{U}, \boldsymbol{V} \in \mathbb{O}^{p\times s}$ as

$$\operatorname{dist}(\boldsymbol{U}, \boldsymbol{V}) = \|\boldsymbol{U}\boldsymbol{U}^* - \boldsymbol{V}\boldsymbol{V}^*\|. \tag{II.1}$$

A. Rectangular dithering

We begin our study with the quantization model considered in [12], [29], i.e., a sign-quantizer with uniform dithers. Given observations $\boldsymbol{y}_k \in \mathbb{F}^p$ as in (I.1), we thus collect quantized samples $\boldsymbol{q}_k^{\square}$ and $\dot{\boldsymbol{q}}_k^{\square}$ where

$$\left\{\boldsymbol{q}_k^{\square}, \dot{\boldsymbol{q}}_k^{\square}\right\} := \left\{\operatorname{sign}_{\mathbb{F}}\big(\boldsymbol{y}_k + \boldsymbol{\tau}_k^{\square}\big), \operatorname{sign}_{\mathbb{F}}\big(\boldsymbol{y}_k + \dot{\boldsymbol{\tau}}_k^{\square}\big)\right\} \tag{II.2}$$

and the *dithering vectors* $\boldsymbol{\tau}_k^{\square}$ and $\dot{\boldsymbol{\tau}}_k^{\square}$ are independently drawn from $\mathcal{U}_{\mathbb{F}}[-\lambda, \lambda]^p$, for $\lambda > 0$ to be determined later. In the case of real measurements, each entry of $\boldsymbol{\tau}_k^{\square}, \dot{\boldsymbol{\tau}}_k^{\square}$ is uniformly distributed in $[-\lambda, \lambda]$ and $\boldsymbol{q}_k^{\square}, \dot{\boldsymbol{q}}_k^{\square} \in \{\pm 1\}^p$, i.e., $\mathcal{A}_{\mathbb{R}} = \{\pm 1\}$ is a one bit alphabet; in the complex case, the real and complex part of each entry of $\boldsymbol{\tau}_k^{\square}, \dot{\boldsymbol{\tau}}_k^{\square}$ are independently drawn from $\mathcal{U}[-\lambda, \lambda]$ and $\boldsymbol{q}_k^{\square}, \dot{\boldsymbol{q}}_k^{\square} \in \{\pm 1 \pm \mathrm{i}\}^p$, i.e., $\mathcal{A}_{\mathbb{C}} = \{\pm 1 \pm \mathrm{i}\}$ is a two bit alphabet.

According to [12] a natural estimator for $\boldsymbol{\Sigma}_{\boldsymbol{y}}$ from quantized samples as in (II.2) is given by $\widehat{\boldsymbol{\Sigma}}_n^{\square}$ where

$$\widehat{\boldsymbol{\Sigma}}_n^{\square} = \frac{1}{2}\widehat{\boldsymbol{\Sigma}}_n' + \frac{1}{2}(\widehat{\boldsymbol{\Sigma}}_n')^* \tag{II.3}$$

and

$$\widehat{\boldsymbol{\Sigma}}_n' = \frac{\lambda^2}{n}\sum_{k=1}^n \boldsymbol{q}_k^{\square}(\dot{\boldsymbol{q}}_k^{\square})^*. \tag{II.4}$$

By combining the ideas in [12], [28] with the Davis-Kahan theorem [8], we derive the following result which depends on the quantities

$$C_\infty = \max_{k\in[n]} \|\mathbb{E}(\boldsymbol{y}_k\boldsymbol{y}_k^*)\|_\infty \quad \text{and} \quad C_{\mathrm{op}} = \max_{k\in[n]} \|\mathbb{E}(\boldsymbol{y}_k\boldsymbol{y}_k^*)\|. \tag{II.5}$$

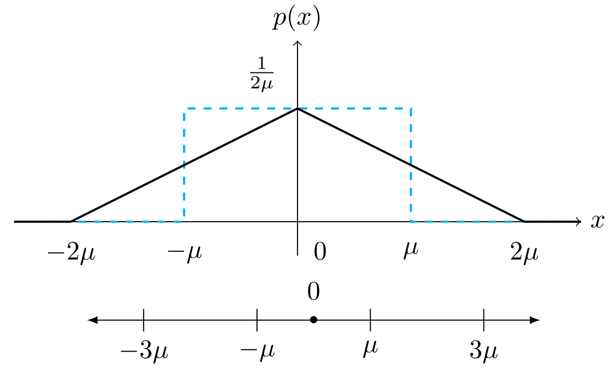

Fig. 1. Triangular vs Uniform distribution and $\mathcal{A}_\mu$.

Theorem II.2. *Suppose Assumption II.1 holds and we have an eigengap $\lambda_s(\boldsymbol{\Sigma}_{\boldsymbol{x}}) > \lambda_{s+1}(\boldsymbol{\Sigma}_{\boldsymbol{x}})$. Then, there exist constants $B_K, C_K > 0$ that depend only on K such that the following hold. For any $\lambda^2 \geq B_K \log(n) C_\infty$, let $\boldsymbol{U}$, $\boldsymbol{U}_n^{\square} \in \mathbb{O}^{p\times s}$ denote orthonormal bases of the leading eigenspaces of $\boldsymbol{\Sigma}_{\boldsymbol{x}}$ and $\widehat{\boldsymbol{\Sigma}}_n^{\square}$. We have with probability at least $1 - e^{-t}$ that*

$$\begin{aligned} &\operatorname{dist}(\boldsymbol{U}_n^{\square}, \boldsymbol{U}) \\ &\leq \min\left\{1, \frac{C_K\,(C_{\mathrm{op}}^{1/2} + \lambda)}{\lambda_s(\boldsymbol{\Sigma}_{\boldsymbol{x}}) - \lambda_{s+1}(\boldsymbol{\Sigma}_{\boldsymbol{x}})}\sqrt{\frac{p(\log(p)+t)}{n}}\right. \\ &\qquad\left. + \frac{C_K\lambda^2}{\lambda_s(\boldsymbol{\Sigma}_{\boldsymbol{x}}) - \lambda_{s+1}(\boldsymbol{\Sigma}_{\boldsymbol{x}})}\frac{p(\log(p)+t)}{n}\right\}. \end{aligned} \tag{II.6}$$

In particular, if $\boldsymbol{x}$ and $\mathbf{e}$ are bounded and $\lambda^2 \geq C_\infty$ a.s., we have with probability at least $1 - e^{-t}$ that (II.6) *holds.*

According to Theorem II.2, the estimation error decays like $\mathcal{O}(1/\lambda_s(\boldsymbol{\Sigma}_{\boldsymbol{x}})\sqrt{n})$ if $\operatorname{rank}(\boldsymbol{\Sigma}_{\boldsymbol{x}}) = s$. While this bound is optimal in n due to minimax optimality of $\widehat{\boldsymbol{\Sigma}}_n^{\square}$ [12], and we confirm in Section III that it is tight in $\lambda_s(\boldsymbol{\Sigma}_{\boldsymbol{x}})$, it does not match the best known guarantees in the unquantized case [3] which are of order $\mathcal{O}(1/\sqrt{\lambda_s(\boldsymbol{\Sigma}_{\boldsymbol{x}})n})$. As it turns out, this discrepancy can be explained by the statistical properties of the dither.

B. Triangular dithering

In recent works on covariance estimation from quantized samples [5], [4] the authors suggest to combine a uniform infinite-range quantizer $Q_\mu \colon \mathbb{R} \to \mathcal{A}_\mu$, $Q_\mu(x) := 2\mu\left(\left\lfloor \frac{x}{2\mu}\right\rfloor + \frac{1}{2}\right)$ of resolution $\mu > 0$ with triangular dithering which is the convolution of two uniform distributions, see Figure 1. Here, $\mathcal{A}_\mu := 2\mu\mathbb{Z} + \mu$.

Remark II.3. For bounded or strongly concentrating data and noise distributions as in Assumption II.1, Q_μ effectively behaves like the sign-quantizer $Q(\cdot) = \operatorname{sign}(\cdot)$ if μ is chosen sufficiently large. The infinite range is only assumed for simplifying the analysis.

The statistical properties of such dithered quantizers were carefully studied [14]. In particular, the *quantization noise*

$$\boldsymbol{\xi}_k := \boldsymbol{q}_k^{\triangle} - \boldsymbol{x}_k = Q_\mu(\boldsymbol{x}_k + \boldsymbol{e}_k + \boldsymbol{\tau}_k^{\triangle}) - \boldsymbol{x}_k. \tag{II.7}$$

has several favorable properties under a mild regularity assumption for $\boldsymbol{e}$.

Assumption II.4. Assume that the coordinates of the distribution of $\boldsymbol{e}$ are uncorrelated, absolutely continuous with respect to the Lebesgue measure on $\mathbb{F}$, and that the characteristic function of each coordinate of $\boldsymbol{e}$ is twice differentiable.

Lemma II.5. *There is an absolute constant $C > 0$ such that the following holds. Let $\boldsymbol{x}_1, \ldots, \boldsymbol{x} \in \mathbb{C}^p$. Let $\boldsymbol{e}_1, \ldots, \boldsymbol{e}_n \overset{\text{i.i.d.}}{\sim} \boldsymbol{e}$. Suppose that $\mathbf{e}$ is K-subgaussian for some $K \geq 1$ and satisfies Assumption II.4 with mean zero and entry-wise variance ν^2. Then $\boldsymbol{\xi}_1, \ldots, \boldsymbol{\xi}_n$ are independent, CK-subgaussian with mean zero and covariance $2(\mu^2 + \nu^2)\boldsymbol{I}_p$.*

Given samples $\boldsymbol{y}_1, \ldots, \boldsymbol{y}_n \in \mathbb{R}^p$ we thus consider

$$\boldsymbol{q}_k^{\triangle} = Q_\mu(\boldsymbol{y}_k + \boldsymbol{\tau}_k^{\triangle}) \in \mathcal{A}_\mu^p,$$

where $\boldsymbol{\tau}_k^{\triangle} \overset{i.i.d.}{\sim} \mathcal{U}_{\mathbb{R}}(-\mu, \mu)^p * \mathcal{U}_{\mathbb{R}}(-\mu, \mu)^p$. Since $\mathbb{E}(\boldsymbol{q}_k^{\triangle}(\boldsymbol{q}_k^{\triangle})^*) = \boldsymbol{\Sigma}_{\boldsymbol{x}} + \sigma^2 \boldsymbol{I}_p$ by Lemma II.5, for $\sigma^2 := c_{\mathbb{F}}^2(\nu^2 + \mu^2)$ where $c_{\mathbb{R}} = 1$ and $c_{\mathbb{C}} = 2$, we deduce the following result for the estimator

$$\widehat{\boldsymbol{\Sigma}}_n^{\triangle} := \frac{1}{n}\sum_{k=1}^{n} \boldsymbol{q}_k^{\triangle}(\boldsymbol{q}_k^{\triangle})^*. \tag{II.8}$$

Theorem II.6. *Suppose Assumptions II.1 and II.4 hold and that $\operatorname{rank}(\boldsymbol{\Sigma}_{\boldsymbol{x}}) = s$. There are absolute constants $C_1, C_2 > 0$, and constants $c, \alpha_0 > 0$ depending only on K such that the following holds. Let $\boldsymbol{U}, \boldsymbol{U}_n^{\triangle} \in \mathbb{O}^{p \times s}$ denote orthonormal bases for the leading left singular spaces of $\boldsymbol{\Sigma}_{\boldsymbol{x}}$ and $\widehat{\boldsymbol{\Sigma}}_n^{\triangle}$.*

(i) For deterministic $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n \subseteq \mathbb{F}^p$ and any $\alpha \geq \alpha_0$, we have with probability at least $1 - e^{-c\alpha p}$

$$\operatorname{dist}(\boldsymbol{U}_n^{\triangle}, \boldsymbol{U}) \leq \tag{II.9}$$
$$\min\left\{1,\, C_1 \sqrt{\frac{\nu^2 + \mu^2}{\lambda_s(\boldsymbol{\Sigma}_{\boldsymbol{x}})}} \left(1 + \sqrt{\frac{\nu^2 + \mu^2}{\lambda_s(\boldsymbol{\Sigma}_{\boldsymbol{x}})}}\right) \sqrt{\frac{\alpha p}{n}}\right\}.$$

(ii) For $\boldsymbol{x}_1, \ldots, \boldsymbol{x}_n \overset{i.i.d.}{\sim} \boldsymbol{x}$ with $\boldsymbol{x} \sim \mathcal{CN}(\boldsymbol{0}, \boldsymbol{\Sigma}_{\boldsymbol{x}})$ and any $\alpha \geq \alpha_0$, we have that if $n \geq C_2 \frac{\sum_{j=1}^{s} \lambda_j(\boldsymbol{\Sigma}_{\boldsymbol{x}})}{\lambda_s(\boldsymbol{\Sigma}_{\boldsymbol{x}})}$ then (II.9) holds with probability at least $1 - 3e^{-c\min\{\alpha p, n\}}$.

Theorem II.6 shows that as soon as $\nu^2 + \mu^2 \lesssim \lambda_s(\boldsymbol{\Sigma}_{\boldsymbol{x}})$, i.e., the noise is sufficiently small and the quantization resolution is sufficiently high, the error bound improves to $\mathcal{O}(1/\sqrt{\lambda_s(\boldsymbol{\Sigma}_{\boldsymbol{x}})n})$. The key for deriving this result are the properties of the quantization noise in Lemma II.5, cf. [10]. A generalization of Theorem II.6(ii) to heavy-tailed distributions, an in-depth comparison with Theorem II.2, and performance guarantees for ESPRIT in spectral estimation that build on those results can be found in [10].

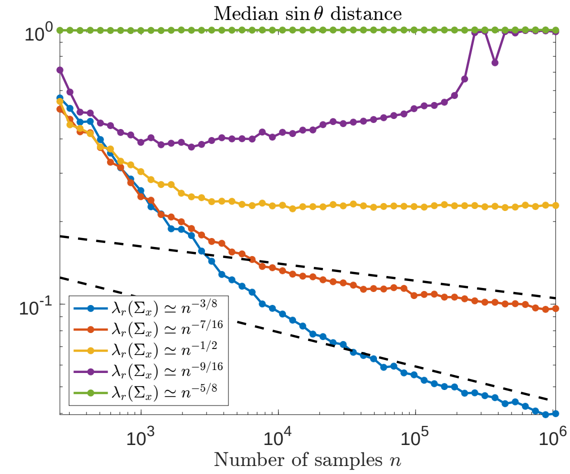

Fig. 2. Dependence on smallest singular value for rectangular dithering. The black dashed lines are $1/(4n^{1/16})$ and $1/(4n^{1/8})$.

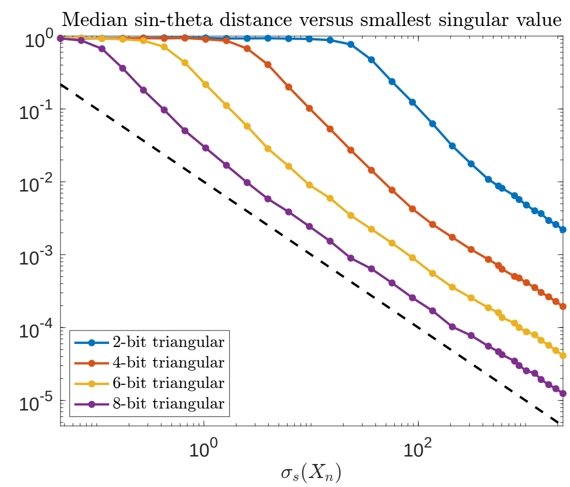

Fig. 3. Dependence on smallest singular value for triangular dither, and the black dashed line is $C/\sigma_s(\boldsymbol{X}_n)$.

III. Numerical Experiments

In the final section, we will validate our theoretical results. Let us first confirm that the worse dependence on $\lambda_s(\boldsymbol{\Sigma}_{\boldsymbol{x}})$ in Theorem II.2 is not an artifact of our proofs. To this end, we consider $\boldsymbol{\Sigma}_{\boldsymbol{x}} = \operatorname{diag}(1, \zeta, \ldots, \zeta, 0, \ldots, 0)$, for $\zeta \in (0, 1)$ and $\operatorname{rank}(\boldsymbol{\Sigma}_{\boldsymbol{x}}) = r$. We set $\nu = 0$, $r = 15$ and $p = 20$, pick an arbitrary rotation matrix $\boldsymbol{R}$, and draw $\boldsymbol{x} = \boldsymbol{R}\boldsymbol{g}$ where $\boldsymbol{g} \sim N(\boldsymbol{0}, \boldsymbol{\Sigma}_{\boldsymbol{x}})$. Theorem II.2 predicts that $\operatorname{dist}(\boldsymbol{U}_n^{\square}, \boldsymbol{U}) \lesssim_{p,\lambda,K} \frac{1}{\lambda_r(\boldsymbol{\Sigma}_{\boldsymbol{x}})\sqrt{n}}$. If we vary n and enforce $\lambda_r(\boldsymbol{\Sigma}_{\boldsymbol{x}}) = \zeta \simeq n^{-\beta}$ for $\beta > 0$, then the upper bound

(i) remains constant at the critical exponent $\beta = 1/2$,
(ii) decreases in n when $\beta < 1/2$,
(iii) achieves the trivial upper bound of 1 when $\beta > 1/2$.

This is confirmed in Figure 2.

Likewise, Theorem II.6 predicts for $\nu^2 + \mu^2 \lesssim \lambda_r(\boldsymbol{\Sigma}_{\boldsymbol{x}})$ that $\operatorname{dist}(\boldsymbol{U}_n^{\triangle}, \boldsymbol{U}) \lesssim_{\nu,\mu,\alpha,p} 1/\sqrt{\lambda_s(\boldsymbol{\Sigma}_{\boldsymbol{x}})n} = 1/\sigma_s(\boldsymbol{X}_n)$. For a different construction of $\boldsymbol{\Sigma}_{\boldsymbol{x}}$ detailed in [10], we get perfect agreement with the predicted bound as depicted in Figure 3.

References

[1] Ofer Bar-Shalom and Anthony J Weiss. DOA estimation using one-bit quantized measurements. *IEEE Transactions on Aerospace and Electronic Systems*, 38(3):868–884, 2002.

[2] Petros T Boufounos, Laurent Jacques, Felix Krahmer, and Rayan Saab. Quantization and compressive sensing. In *Compressed Sensing and its Applications: MATHEON Workshop 2013*, pages 193–237. Springer, 2015.

[3] T. Tony Cai and Anru Zhang. Rate-optimal perturbation bounds for singular subspaces with applications to high-dimensional statistics. *The Annals of Statistics*, 46(1):60–89, 2018.

[4] Junren Chen and Michael K Ng. A parameter-free two-bit covariance estimator with improved operator norm error rate. *arXiv preprint arXiv:2308.16059*, 2023.

[5] Junren Chen, Michael K Ng, and Di Wang. Quantizing heavy-tailed data in statistical estimation:(near) minimax rates, covariate quantization, and uniform recovery. *IEEE Transactions on Information Theory*, 2023.

[6] Junren Chen, Cheng-Long Wang, Michael K Ng, and Di Wang. High dimensional statistical estimation under uniformly dithered one-bit quantization. *IEEE Transactions on Information Theory*, 2023.

[7] Xiaohan Chen, Jialin Liu, Zhangyang Wang, and Wotao Yin. Theoretical linear convergence of unfolded ista and its practical weights and thresholds. *Advances in Neural Information Processing Systems*, 31, 2018.

[8] Chandler Davis and William Morton Kahan. The rotation of eigenvectors by a perturbation. iii. *SIAM Journal on Numerical Analysis*, 7(1):1–46, 1970.

[9] Sjoerd Dirksen. Quantized compressed sensing: A survey. In *Compressed Sensing and Its Applications: Third International MATHEON Conference 2017*, pages 67–95. Applied and Numerical Harmonic Analysis. Birkhäuser, Cham, 2019.

[10] Sjoerd Dirksen, Weilin Li, and Johannes Maly. Subspace and DOA estimation under coarse quantization. *arXiv preprint arXiv:2502.17037*, 2025.

[11] Sjoerd Dirksen and Johannes Maly. Tuning-free one-bit covariance estimation using data-driven dithering. *IEEE Transactions on Information Theory*, 2024.

[12] Sjoerd Dirksen, Johannes Maly, and Holger Rauhut. Covariance estimation under one-bit quantization. *The Annals of Statistics*, 50(6):3538–3562, 2022.

[13] Robert M. Gray and David L. Neuhoff. Quantization. *IEEE Transactions on Information Theory*, 44(6):2325–2383, 1998.

[14] Robert M Gray and Thomas G Stockham. Dithered quantizers. *IEEE Transactions on Information Theory*, 39(3):805–812, 1993.

[15] C Sinan Güntürk and Weilin Li. Quantization for spectral super-resolution. *Constructive Approximation*, 56(3):619–648, 2022.

[16] Xiaodong Huang and Bin Liao. One-bit music. *IEEE Signal Processing Letters*, 26(7):961–965, 2019.

[17] Vladimir Koltchinskii and Karim Lounici. Concentration inequalities and moment bounds for sample covariance operators. *Bernoulli*, pages 110–133, 2017.

[18] Xiao-Peng Li, Zhang-Lei Shi, Lei Huang, Anthony Man-Cho So, and Hing Cheung So. Rocs: Robust one-bit compressed sensing with application to direction of arrival. *IEEE Transactions on Signal Processing*, 2024.

[19] Zeyang Li, Junpeng Shi, Xinhai Wang, and Fangqing Wen. Joint angle and frequency estimation using one-bit measurements. *Sensors*, 19(24):5422, 2019.

[20] Yan Pan, Li Zhang, Liyan Xu, and Fabing Duan. DOA Estimation on One-Bit Quantization Observations through Noise-Boosted Multiple Signal Classification. *Sensors*, 24(14):4719, 2024.

[21] Lawrence Roberts. Picture coding using pseudo-random noise. *IRE Transactions on Information Theory*, 8(2):145–154, 1962.

[22] Richard Roy and Thomas Kailath. ESPRIT-estimation of signal parameters via rotational invariance techniques. *IEEE Transactions on Acoustics, Speech, and Signal Processing*, 37(7):984–995, 1989.

[23] Ningjun Ruan, Han Wang, Fangqing Wen, and Junpeng Shi. DOA estimation in B5G/6G: Trends and challenges. *Sensors*, 22(14):5125, 2022.

[24] Ralph Schmidt. Multiple emitter location and signal parameter estimation. *IEEE Transactions on Antennas and Propagation*, 34(3):276–280, 1986.

[25] Saeid Sedighi, MR Bhavani Shankar, Mojtaba Soltanalian, and Björn Ottersten. On the performance of one-bit DoA estimation via sparse linear arrays. *IEEE Transactions on Signal Processing*, 69:6165–6182, 2021.

[26] Zhenyu Wei, Wei Wang, Fuwang Dong, and Qi Liu. Gridless one-bit direction-of-arrival estimation via atomic norm denoising. *IEEE communications letters*, 24(10):2177–2181, 2020.

[27] Peng Xiao, Bin Liao, and Nikos Deligiannis. DeepFPC: A deep unfolded network for sparse signal recovery from 1-bit measurements with application to DoA estimation. *Signal Processing*, 176:107699, 2020.

[28] Tianyu Yang, Johannes Maly, Sjoerd Dirksen, and Giuseppe Caire. Plug-in channel estimation with dithered quantized signals in spatially non-stationary massive MIMO systems. *IEEE Transactions on Communications*, 2023.

[29] Farhang Yeganegi, Arian Eamaz, Tara Esmaeilbeig, and Mojtaba Soltanalian. Deep Learning-Enabled One-Bit DoA Estimation. *ArXiv:2405.09712*, 2024.

[30] Kai Yu, Yimin D Zhang, Ming Bao, Yu-Hen Hu, and Zhi Wang. DOA estimation from one-bit compressed array data via joint sparse representation. *IEEE Signal Processing Letters*, 23(9):1279–1283, 2016.

[31] Chengwei Zhou, Yujie Gu, Zhiguo Shi, and Martin Haardt. Direction-of-arrival estimation for coprime arrays via coarray correlation reconstruction: A one-bit perspective. In *2020 IEEE 11th Sensor Array and Multichannel Signal Processing Workshop (SAM)*, pages 1–4. IEEE, 2020.