

Chapter 18

Uncertainty Quantification for Matrix Compressed Sensing and Quantum Tomography Problems



Alexandra Carpentier, Jens Eisert, David Gross, and Richard Nickl

Abstract We construct minimax optimal non-asymptotic confidence sets for low rank matrix recovery algorithms such as the Matrix Lasso or Dantzig selector. These are employed to devise adaptive sequential sampling procedures that guarantee recovery of the true matrix in Frobenius norm after a data-driven stopping time \hat{n} for the number of measurements that have to be taken. With high probability, this stopping time is minimax optimal. We detail applications to quantum tomography problems where measurements arise from Pauli observables. We also give a theoretical construction of a confidence set for the density matrix of a quantum state that has optimal diameter in nuclear norm. The non-asymptotic properties of our confidence sets are further investigated in a simulation study.

Keywords Low rank recovery · Quantum information · Confidence sets · Sequential sampling

A. Carpentier
Institut für Mathematische Stochastik, Otto von Guericke Universität Magdeburg, Magdeburg,
Germany
e-mail: alexandra.carpentier@ovgu.de

J. Eisert
Department of Physics, Dahlem Center for Complex Quantum Systems, Freie Universität zu
Berlin, Berlin, Germany

D. Gross
Institute of Theoretical Physics, Universität Köln, Köln, Germany
e-mail: david.gross@thp.uni-koeln.de

R. Nickl (✉)
Statistical Laboratory, University of Cambridge, Cambridge, UK
e-mail: r.nickl@statslab.cam.ac.uk

18.1 Introduction

18.1.1 *Uncertainty Quantification in Compressed Sensing*

Compressed sensing and related convex relaxation algorithms have had a profound impact on high-dimensional statistical modelling in recent years. They provide efficient recovery of low-dimensional objects that sit within a high-dimensional ill-posed system of linear equations. Prototypical low-dimensional structures are described by *sparsity* and *low rank* hypotheses. The statistical analysis of such algorithms has mostly been concerned with recovery rates, or with the closely related question of how many measurements are sufficient to reach a prescribed recovery level—some key references are Refs. [3, 5–7, 16, 24, 25, 33].

A statistical question of fundamental importance that has escaped a clear answer so far is the question of *uncertainty quantification*: Can we tell from the observations how well the algorithm has worked? In technical terms: can we report confidence sets for the unknown parameter? Or, in the sequential sampling setting, can we give *data-driven* rules that ensure recovery of the true parameter at a given precision? Answers to this question are of great importance in various applications of compressed sensing. For one-dimensional subproblems, such as projection onto a fixed coordinate of the parameter vector, recent advances have provided some useful confidence intervals (see Refs. [8, 9, 22, 36]), but our understanding of valid inference procedures for the *entire* parameter remains limited.

Whereas the ‘estimation theory’ for compressed sensing is quite similar for sparsity and low rank constraints, this is not so for the theory of confidence sets. On the one hand, if one is interested in inference on the full parameter, *sparsity* conditions induce information theoretic barriers, as shown in Ref. [30]: unless one is willing to make additional signal strength assumptions (inspired by the literature on nonparametric confidence sets, such as Refs. [13, 19]), a uniformly valid confidence set for the unknown parameter vector θ cannot have a better performance than $1/\sqrt{n}$ in quadratic loss. This significantly falls short of the optimal recovery rates $(k \log p)/n$ for the most interesting sparsity levels k . On the other hand, and perhaps surprisingly, we will show in this article that the *low-rank constraint* is naturally compatible with certain risk estimation approaches to confidence sets. This will be seen to be true for general sub-Gaussian sensing matrices, but also for sensing matrices arising from Pauli observables, as is specifically relevant in quantum state tomography problems (see the next section). In the latter case it will be helpful to enforce the additional ‘quantum state shape constraint’ on the unknown matrix to obtain optimal results. One can conclude that, in contrast to ‘sparse models’, *no signal strength assumptions* are necessary for the existence of adaptive confidence statements in low rank recovery problems. Our findings are confirmed in a simple simulation study, see Sect. 18.4.

The honest non-asymptotic confidence sets we will derive below can be used for the construction of *adaptive sampling procedures*: An experimenter wants to know—at least with a prescribed probability of $1 - \alpha$ —that the matrix recovery

algorithm, the ‘estimator’, has produced an output $\tilde{\theta}$ that is close to the ‘true state’ θ . The sequential protocol advocated here—which is related to ‘active learning algorithms’ from machine learning, e.g., Ref. [29]—should tell the experimenter whether a new batch of measurements has to be taken to decrease the recovery error, or whether the collected observations are already sufficient. The data-driven stopping time for this protocol should not exceed the minimax optimal stopping time, again with high probability. We shall show that for Pauli and sub-Gaussian sensing ensembles, such algorithms exist under mild assumptions on the true matrix θ . These assumptions are in particular always satisfied under the ‘quantum shape constraint’ that naturally arises in quantum tomography problems.

Our results depend on the choice of the Frobenius norm and the Hilbert space geometry induced by it. For other natural matrix norms, such as for instance the trace-(nuclear) norm, the theory is more difficult. We show as a first step that at least theoretically a trace-norm optimal confidence set can be constructed for the unknown quantum state (Theorem 18.4)—this suggests interesting directions for future research.

18.1.2 Application to Quantum State Estimation

This work was partly motivated by a problem arising in present-day physics experiments that aim at estimating quantum states. Conceptually, a quantum mechanical experiment involves two stages (c.f. Fig. 18.1): A *source* (or *preparation procedure*) that emits quantum mechanical systems with unknown properties, and a *measurement device* that interacts with incoming quantum systems and produces real-valued measurement outcomes, e.g. by pointing a dial to a value on a scale. Quantum mechanics stipulates that both stages are completely described by certain matrices.

The properties of the source are represented by a positive semi-definite unit trace matrix θ , the *quantum state*, also referred to as *density matrix*. In turn, the measurement device is modelled by a Hermitian matrix X , which is referred to as an *observable* in physics jargon. A key axiom of the quantum mechanical formalism



Fig. 18.1 Caricature of a quantum mechanical experiment. With every source of quantum systems, one associates a *density matrix* θ . Observations systems are performed by measurement devices, which interact with incoming systems and produce real-valued outcomes. Each such devices is modelled mathematically by a Hermitian matrix X , referred to as an *observable*

states that if the measurement X is repeatedly performed on systems emitted by the source that is preparing θ , then the real-valued measurement outcomes will fluctuate randomly with expected value

$$\langle X, \theta \rangle_F = \text{tr}(X\theta) \quad (18.1)$$

(referred to as *expectation value* in the quantum physics literature). The precise way in which physical properties are represented by these matrices is immaterial to our discussion (cf. any textbook, e.g. Ref. [32]). We merely note that, while in principle *any* Hermitian X can be measured by some physical apparatus, the required experimental procedures are prohibitively complicated for all but a few highly structured matrices. This motivates the introduction of *Pauli designs* below, which correspond to fairly tractable ‘spin parity measurements’.

The *quantum state estimation* or *quantum state tomography*¹ problem is to estimate an unknown density matrix θ from the measurement of a collection of observables X^1, \dots, X^n . This task is of particular importance to the young field of quantum information science [31]. There, the sources might be a carefully engineered component used for technological applications such as quantum key distribution or quantum computing. In this context, quantum state estimation is the process of characterising the components one has built—clearly an important capability for any technology.

A major challenge lies in the fact that relevant instances are described by $d \times d$ -matrices for fairly large dimensions d ranging from 100 to 10,000 in presently performed experiments [18]. Such high-dimensional estimation problems can benefit substantially from structural properties of the objects to be recovered. Fortunately, the density matrices occurring in quantum information experiments are typically well-approximated by matrices of *low rank* $r \ll d$. In fact, in the practically most important applications, one usually even aims at preparing a state of rank one—a so-called *pure quantum state*. While environmental noise will drive the actual state away from the perfect rank-one case, the error will usually be small.

As a result, quantum physicists have early on shown an interest in low-rank matrix recovery methods [12, 15–17, 28]. Initial works [15, 16] focused on the minimal number n of observables X^1, \dots, X^n required for reconstructing a rank- r density matrix θ in the noiseless case, i.e. under the idealised assumption that the expectation values $\text{tr}(\theta X^i)$ are known exactly. The practically highly relevant problem of quantifying the uncertainty of an estimate $\hat{\theta}$ arising from noisy observations on low-rank states was addressed only later [12] and remains less well understood.

¹ The term ‘tomography’ goes back to the use of *Radon transforms* in early schemes for estimating quantum states of electromagnetic fields [1, 27]. It has become synonymous with ‘quantum density matrix estimation’, even though current methods applied to quantum systems with a finite dimension d have no technical connection to classical tomographic reconstruction algorithms.

More concretely, the basic approach taken in Ref. [12] for uncertainty quantification is similar to the one pursued in the present paper. In a first step, one uses a Matrix Lasso or Dantzig Selector to construct an estimate. Then, a confidence region is obtained by comparing predictions derived from the initial estimate to new samples. However, Ref. [12] suffers from two demerits. First, and most importantly, the performance analysis of the scheme relies on a bound on the rank r of the unknown true θ . Such a bound is not available in practice. Second, the dependence of the rate on r is not tight. Both of these demerits will be addressed here.

We close this section pointing to more broadly related works. Uncertainty quantification in quantum state tomography in general has been treated by numerous authors—a highly incomplete list is Refs. [2, 4, 10, 34, 35]. However, the concept of dimension reduction for low-rank states does not feature explicitly in these papers. This contrasts with Ref. [17], where the authors propose model selection techniques based on information criteria to arrive at low-rank estimates. The use of general-purpose methods—like maximum likelihood estimation and the Akaike Information Criterion—in Ref. [17] means that it is applicable to very general experimental designs. In contrast to this, the present paper relies on compressed sensing ideas to arrive at rigorous a priori guarantees on statistical and computational performance. Also, it remains non-obvious how such model selection steps can be transformed into ‘post-model selection’ confidence sets—typically such constructions result in sub-optimal signal strength conditions that ensure model selection consistency (see Ref. [26] and also the discussion after Theorem 2 in Ref. [30]). Our confidence procedures never estimate the unknown rank of the quantum state—not even implicitly. Rather, they estimate the performance of a dimension-reduction technique directly based on sample splitting.

18.2 Matrix Compressed Sensing

We consider inference on a $d \times d$ matrix θ that is symmetric, or, if it consists of possibly complex entries, assumed to be Hermitian (that is $\theta = \theta^*$ where θ^* is the conjugate transpose of θ). Denote by $\mathbb{M}_d(\mathbb{K})$ the space of $d \times d$ matrices with entries in $\mathbb{K} = \mathbb{C}$ or $\mathbb{K} = \mathbb{R}$. We write $\|\cdot\|_F$ for the usual Frobenius norm on $\mathbb{M}_d(\mathbb{K})$ arising from the inner product $\text{tr}(AB) = \langle A, B \rangle_F$. Moreover let $\mathbb{H}_d(\mathbb{C})$ be the set of all Hermitian matrices, and $\mathbb{H}_d(\mathbb{R})$ for the set of all symmetric $d \times d$ matrices with real entries. The norm symbol $\|\cdot\|$ without subindex denotes the standard Euclidean norm on \mathbb{R}^n or on \mathbb{C}^n arising from the Euclidean inner product $\langle \cdot, \cdot \rangle$.

We denote the usual operator norm on $\mathbb{M}_d(\mathbb{C})$ by $\|\cdot\|_{op}$. For $M \in \mathbb{M}_d(\mathbb{C})$ let $(\lambda_k^2 : k = 1, \dots, d)$ be the eigenvalues of $M^T M$ (which are all real-valued and positive). The l_1 -Schatten, trace, or nuclear norm of M is defined as

$$\|M\|_{S_1} = \sum_{j \leq d} |\lambda_j|.$$

Note that for any matrix M of rank $1 \leq r \leq d$ the following inequalities are easily shown,

$$\|M\|_F \leq \|M\|_{S_1} \leq \sqrt{r} \|M\|_F. \quad (18.2)$$

We will consider parameter subspaces of $\mathbb{H}_d(\mathbb{C})$ described by low rank constraints on θ , and denote by $R(k)$ the space of all Hermitian $d \times d$ matrices that have rank at most k , $k \leq d$. In quantum tomography applications, we may assume an additional ‘shape constraint’, namely that θ is a density matrix of a quantum state, and hence contained in *state space*

$$\Theta_+ = \{\theta \in \mathbb{H}_d(\mathbb{C}) : \text{tr}(\theta) = 1, \theta \succeq 0\},$$

where $\theta \succeq 0$ means that θ is positive semi-definite. In fact, in most situations, we will only require the bound $\|\theta\|_{S_1} \leq 1$ which trivially holds for any θ in Θ_+ .

We have at hand measurements arising from inner products $\langle X^i, \theta \rangle_F = \text{tr}(X^i \theta)$, $i = 1, \dots, n$, of θ with $d \times d$ (random) matrices X^i . This measurement process is further subject to independent additive noise ε . Formally, the measurement model is

$$Y_i = \text{tr}(X^i \theta) + \varepsilon_i, \quad i = 1, \dots, n, \quad (18.3)$$

where the ε_i ’s and X^i ’s are independent of each other. We write $Y = (Y_1, \dots, Y_n)^T$, and for probability statements under the law of Y, X, ε given fixed θ we will use the symbol \mathbb{P}_θ . Unless mentioned otherwise we will make the basic assumption of Gaussian noise

$$\varepsilon = (\varepsilon_1, \dots, \varepsilon_n)^T \sim N(0, \sigma^2 I_n),$$

where $\sigma > 0$ is known. See Remark 18.6 for some discussion of the unknown variance case. In the context of quantum mechanics, the inner product $\text{tr}(X^i \theta)$ gives the expected value of the observable X^i when measured on a system in state θ (cf. Sect. 18.1.2). A class of physically realistic measurements (correlations among spin-1/2 particles) is described by X^i ’s drawn from the *Pauli basis*. Our main results also hold for measurement processes of this type. Before we describe this in Sect. 18.2.2, let us first discuss our assumptions on the matrices X^i .

18.2.1 Sensing Matrices and the RIP

When $\theta \in \mathbb{M}_d(\mathbb{R})$, we shall restrict to design matrices X^i that have real-valued entries, too, and when $\theta \in \mathbb{H}_d(\mathbb{C})$ we shall consider designs where $X^i \in \mathbb{H}_d(\mathbb{C})$. This way, in either case, the measurements $\text{tr}(X_i \theta)$ ’s and hence the Y_i ’s are all real-valued. More concretely, the sensing matrices X^i that we shall consider are

described in the following assumption, which encompasses both a prototypical compressed sensing setting—where we can think of the matrices X^i as i.i.d. draws from a Gaussian ensemble $(X_{m,k}) \sim^{iid} N(0, 1)$ —as well as the ‘random sampling from a basis of $\mathbb{M}_d(\mathbb{C})$ ’ scenario. The systematic study of the latter has been initiated by quantum physicists [15, 28], as it contains, in particular, the case of Pauli basis measurements [12, 16] frequently employed in quantum tomography problems. Note that in Part (a) the design matrices are not Hermitian but our results can easily be generalised to symmetrised sub-Gaussian ensembles (as those considered in Ref. [24]).

Condition 18.1

(a) $\theta \in \mathbb{H}_d(\mathbb{R})$, **‘isotropic’ sub-Gaussian design**: The random variables $(X_{m,k}^i)$, $1 \leq m, k \leq d, i = 1, \dots, n$, generating the entries of the random matrix X^i are i.i.d. distributed across all indices i, m, k with mean zero and unit variance. Moreover, for every $\theta \in \mathbb{M}_d(\mathbb{R})$ such that $\|\theta\|_F \leq 1$, the real random variables $Z_i = \text{tr}(X^i \theta)$ are sub-Gaussian: for some fixed constants $\tau_i > 0$ independent of θ ,

$$\mathbb{E}e^{\lambda Z_i} \leq \tau_i e^{\lambda^2 \tau_i^2} \quad \forall \lambda \in \mathbb{R}.$$

(b) $\theta \in \mathbb{H}_d(\mathbb{C})$, **random sampling from a basis (‘Pauli design’)**: Let $\{E_1, \dots, E_{d^2}\} \subset \mathbb{H}_d(\mathbb{C})$ be a basis of $\mathbb{M}_d(\mathbb{C})$ that is orthonormal for the scalar product $\langle \cdot, \cdot \rangle_F$ and such that the operator norms satisfy, for all $i = 1, \dots, d^2$,

$$\|E_i\|_{op} \leq \frac{K}{\sqrt{d}},$$

for some universal ‘coherence’ constant K . [In the Pauli basis case we have $K = 1$.] Assume the $X^i, i = 1, \dots, n$, are draws from the finite family $\mathcal{E} = \{dE_i : i = 1, \dots, d^2\}$ sampled uniformly at random.

The above examples all obey the *matrix restricted isometry property*, that we describe now. Note first that if $\mathcal{X} : \mathbb{R}^{d \times d} \rightarrow \mathbb{R}^n$ is the linear ‘sampling’ operator

$$\mathcal{X} : \theta \mapsto \mathcal{X}\theta = (\text{tr}(X^1 \theta), \dots, \text{tr}(X^n \theta))^T, \quad (18.4)$$

so that we can write the model equation (18.3) as $Y = \mathcal{X}\theta + \varepsilon$, then in the above examples we have the ‘expected isometry’

$$\mathbb{E} \frac{1}{n} \|\mathcal{X}\theta\|^2 = \|\theta\|_F^2.$$

Indeed, in the isotropic design case we have

$$\frac{1}{n} \mathbb{E} \|\mathcal{X}\theta\|^2 = \frac{1}{n} \sum_{i=1}^n \mathbb{E} \left(\sum_m \sum_k X_{m,k}^i \theta_{m,k} \right)^2 = \sum_m \sum_k \mathbb{E} X_{m,k}^2 \theta_{m,k}^2 = \|\theta\|_F^2, \tag{18.5}$$

and in the ‘basis case’ we have, from Parseval’s identity and since the X^i ’s are sampled uniformly at random from the basis,

$$\frac{1}{n} \mathbb{E} \|\mathcal{X}\theta\|^2 = \frac{d^2}{n} \sum_{i=1}^n \sum_{j=1}^d \Pr(X^i = E_j) |\langle E_j, \theta \rangle_F|^2 = \|\theta\|_F^2. \tag{18.6}$$

The restricted isometry property (RIP) then requires that this ‘expected isometry’ actually holds, up to constants and with probability $\geq 1 - \delta$, for a given realisation of the sampling operator, and for all $d \times d$ matrices θ of rank at most k :

$$\sup_{\theta \in R(k)} \left| \frac{\frac{1}{n} \|\mathcal{X}\theta\|^2 - \|\theta\|_F^2}{\|\theta\|_F^2} \right| \leq \tau_n(k), \tag{18.7}$$

where $\tau_n(k)$ are some constants that may depend, among other things, on the rank k and the ‘exceptional probability’ δ . For the above examples of isotropic and Pauli basis design inequality (18.7) can be shown to hold with

$$\tau_n^2(k) = c^2 \frac{kd \cdot \overline{\log d}}{n}, \tag{18.8}$$

where

$$\overline{\log x} := (\log x)^\eta,$$

for some $\eta > 0$ denotes a ‘polylog function’, and where $c = c(\delta)$ is a constant. See Refs. [6, 28] for these results, where it is also shown that $c(\delta)$ can be taken to be at least $O(1/\delta^2)$ as $\delta \rightarrow 0$ (sufficient for our purposes below).

18.2.2 Quantum Measurements

Here, we introduce a paradigmatic set of quantum measurements that is frequently used in both theoretical and practical treatments of quantum state estimation (e.g. [16, 18]). For a more general account, we refer to standard textbooks [20, 31]. The purpose of this section is to motivate the ‘Pauli design’ case (Condition 18.1(b) of

the main theorem, as well as the approximate Gaussian noise model. Beyond this, the technical details presented here will not be used.

18.2.2.1 Pauli Spin Measurements on Multiple Particles

We start by describing ‘spin measurements’ on a single ‘spin-1/2 particle’. Such a measurement corresponds to the situation of having $d = 2$. Without worrying about the physical significance, we accept as fact that on such particles, one may measure one of three properties, referred to as the ‘spin along the x , y , or z -axis’ of \mathbb{R}^3 . Each of these measurements may yield one of two outcomes, denoted by $+1$ and -1 respectively.

The mathematical description of these measurements is derived from the *Pauli matrices*

$$\sigma^1 = \begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}, \sigma^2 = \begin{bmatrix} 0 & -i \\ i & 0 \end{bmatrix}, \sigma^3 = \begin{bmatrix} 1 & 0 \\ 0 & -1 \end{bmatrix} \quad (18.9)$$

in the following way. Recall that the Pauli matrices have eigenvalues ± 1 . For $x \in \{1, 2, 3\}$ and $j \in \{+1, -1\}$, we write ψ_j^x for the normalised eigenvector of σ^x with eigenvalue j . The spectral decomposition of each Pauli spin matrix can hence be expressed as

$$\sigma^x = \pi_+^x - \pi_-^x, \quad (18.10)$$

with

$$\pi_\pm^x = \psi_\pm^x (\psi_\pm^x)^* \quad (18.11)$$

denoting the projectors onto the eigenspaces. Now, a physical measurement of the ‘spin along direction x ’ on a system in state θ will give rise to a $\{-1, 1\}$ -valued random variable C^x with

$$\mathbb{P}(C^x = j) = \text{tr}(\pi_j^x \theta), \quad (18.12)$$

where $\theta \in \mathbb{H}_2(\mathbb{C})$. Using Eq. (18.10), this is equivalent to stating that the expected value of C^x is given by

$$\mathbb{E}(C^x) = \text{tr}(\sigma^x \theta). \quad (18.13)$$

Next, we consider the case of joint spin measurements on a collection of N particles. For each, one has to decide on an axis for the spin measurement. Thus, the joint *measurement setting* is now described by a word $x = (x_1, \dots, x_N) \in \{1, 2, 3\}^N$. The axioms of quantum mechanics posit that the joint state θ of the N particles acts on the tensor product space $(\mathbb{C}^2)^{\otimes N}$, so that $\theta \in \mathbb{H}_{2^N}(\mathbb{C})$.

Likewise, the *measurement outcome* is a word $j = (j_1, \dots, j_N) \in \{1, -1\}^N$, with j_i the value of the spin along axis x_i of particle $i = 1, \dots, N$. As above, this prescription gives rise to a $\{1, -1\}^N$ -valued random variable C^x . Again, the axioms of quantum mechanics imply that the distribution of C^x is given by

$$\mathbb{P}(C^x = j) = \text{tr} \left((\pi_{j_1}^{x_1} \otimes \dots \otimes \pi_{j_N}^{x_N}) \theta \right). \tag{18.14}$$

Note that the components of the random vector C^x are not necessarily independent, as θ will generally not factorise

It is often convenient to express the information in Eq. (18.14) in a way that involves tensor products of Pauli matrices, rather than their spectral projections. In other words, we seek a generalisation of Eq. (18.13) to N particles. As a first step toward this goal, let

$$\chi(j) = \begin{cases} -1 & \text{number of } -1 \text{ elements in } j \text{ is odd} \\ 1 & \text{number of } -1 \text{ elements in } j \text{ is even} \end{cases} \tag{18.15}$$

be the *parity function*. Then one easily verifies

$$\text{tr}((\sigma^{x_1} \otimes \dots \otimes \sigma^{x_N})\theta) = \sum_{j \in \{1, -1\}^N} \chi(j) \text{tr} \left(\theta(\pi_{j_1}^{x_1} \otimes \dots \otimes \pi_{j_N}^{x_N}) \right) = \mathbb{E}(\chi(C^x)). \tag{18.16}$$

In this sense, the tensor product $\sigma^{x_1} \otimes \dots \otimes \sigma^{x_N}$ describes a measurement of the parity of the spins along the respective directions given by x .

In fact, the entire distribution of C^x can be expressed in terms of tensor products of Pauli matrices and suitable parity functions. To this end, we extend the definitions above. Write

$$\sigma^0 = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \tag{18.17}$$

for the identity matrix in $\mathbb{M}_2(\mathbb{C})$. For every subset S of $\{1, \dots, N\}$, define the ‘parity function restricted to S ’ via

$$\chi_S(j) = \begin{cases} -1 & \text{number of } -1 \text{ elements } j_i \text{ for } i \in S \text{ is odd} \\ 1 & \text{number of } -1 \text{ elements } j_i \text{ for } i \in S \text{ is even.} \end{cases} \tag{18.18}$$

Lastly, for $S \subset \{1, \dots, N\}$ and $x \in \{1, 2, 3\}^N$, the *restriction of x to S* is

$$x_i^S = \begin{cases} x_i & i \in S \\ 0 & i \notin S. \end{cases} \tag{18.19}$$

Then for every such x, S one verifies the identity

$$\text{tr}((\sigma^{x_1^S} \otimes \dots \otimes \sigma^{x_N^S})\theta) = \mathbb{E}(\chi_S(C^x)). \tag{18.20}$$

In other words, the distribution of C^x contains enough information to compute the expectation value of all observables $(\sigma^{x_1^S} \otimes \dots \otimes \sigma^{x_N^S})$ that can be obtained by replacing the Pauli matrices on an arbitrary subset S of particles by the identity σ^0 . The converse is also true: the set of all such expectation values allows one to recover the distribution of C^x . The explicit formula reads

$$\mathbb{P}(C^x = j) = \frac{1}{2^N} \sum_{S \subset \{1, \dots, N\}} \chi_S(j) \mathbb{E}(\chi_S(C^x)) = \frac{1}{2^N} \sum_{S \in \{1, \dots, N\}} \chi_S(j) \text{tr}(\theta(\sigma^{x_1^S} \otimes \dots \otimes \sigma^{x_N^S})) \tag{18.21}$$

and can be verified by direct computation.²

In this sense, the information obtainable from joint spin measurements on N particles can be encoded in the 4^N real numbers

$$2^{-N/2} \text{tr}((\sigma^{y_1} \otimes \dots \otimes \sigma^{y_N})\theta), \quad y \in \{0, 1, 2, 3\}^N. \tag{18.22}$$

Indeed, every such y arises as $y = x^S$ for some (generally non-unique) combination of x and S . This representation is particularly convenient from a mathematical point of view, as the collection of matrices

$$E^y := 2^{-N/2} \sigma^{y_1} \otimes \dots \otimes \sigma^{y_N}, \quad y \in \{0, 1, 2, 3\}^N \tag{18.23}$$

forms an ortho-normal basis with respect to the $\langle \cdot, \cdot \rangle_F$ inner product. Thus the terms in Eq. (18.22) are just the coefficients of a basis expansion of the density matrix θ .³

From now on, we will use Eq. (18.22) as our model for quantum tomographic measurements. Note that the E^y satisfy Condition 18.1(b) with coherence constant $K = 1$ and $d = 2^N$.

² A more insightful way of proving the first identity is to realise that $\mathbb{E}(\chi_S(C^x))$ is effectively a Fourier coefficient (over the group \mathbb{Z}_2^N) of the distribution function of the $\{-1, 1\}^N$ -valued random variable C^x (e.g., [11]). Equation (18.21) is then nothing but an inverse Fourier transform.

³We note that quantum mechanics allows to design measurement devices that directly probe the observable of $\sigma^{y_1} \otimes \dots \otimes \sigma^{y_N}$, without first measuring the spin of every particle and then computing a parity function. In fact, the ability to perform such correlation measurements is crucial for quantum error correction protocols [31]. For practical reasons these setups are used less commonly in tomography experiments, though.

18.2.2.2 Bernoulli Errors and Pauli Observables

In the model (18.3) under Condition 18.1(b) we wish to approximate $d \cdot \text{tr}(E^y \theta)$ for a fixed observable E^y (we fix the random values of the X^i 's here) and for $d = 2^N$. If $y = x^S$ for some setting x and subset S , then the parity function $B^y := \chi_S(C^x)$ has expected value $2^{N/2} \cdot \text{tr}(E^y \theta) = \sqrt{d} \cdot \text{tr}(E^y \theta)$ (see Eqs. (18.20) and (18.23)), and itself is a Bernoulli variable taking values $\{1, -1\}$ with

$$p = \mathbb{P}(B^y = 1) = \frac{1 + \sqrt{d} \text{tr}(E^y \theta)}{2}.$$

Note that

$$\sqrt{d} |\text{tr}(E^y \theta)| \leq \sqrt{d} \|E^y\|_{op} \|\theta\|_{S_1} \leq 1,$$

so indeed $p \in [0, 1]$ and the variance satisfies

$$\text{Var} B^y = 1 - d \cdot \text{tr}(E^y \theta)^2 \leq 1.$$

This is the error model considered in Ref. [12].

In order to estimate all $Y_i, i = 1, \dots, n$, for given $E_i := E^y$, a total number nT of identical preparations of the quantum state θ are being performed, divided into batches of T Bernoulli variables $B_{i,j} := B_j^y, j = 1, \dots, T$. The measurements of the sampling model Eq. (18.3) are thus

$$Y_i = \frac{\sqrt{d}}{T} \sum_{j=1}^T B_{i,j} = d \cdot \text{tr}(E_i \theta) + \varepsilon_i$$

where

$$\varepsilon_i = \frac{\sqrt{d}}{T} \sum_{j=1}^T (B_{i,j} - \mathbb{E} B_{i,j})$$

is the effective error arising from the measurement procedure making use of T preparations to estimate each quantum mechanical expectation value. Now note that

$$|\varepsilon_i| \leq 2\sqrt{d}, \quad \mathbb{E} \varepsilon_i^2 \leq \frac{d}{T} \text{Var}(B_{i,1}) \leq \frac{d}{T}. \quad (18.24)$$

We see that since the ε_i 's are themselves sums of independent random variables, an approximate Gaussian error model with variance σ^2 will be roughly appropriate. If $T \geq n$ then $\sigma^2 = \mathbb{E} \varepsilon_1^2$ is no greater than d/n , and if in addition $T \geq d^2$ then all results in Sect. 18.3 below can be proved for this Bernoulli noise model too, see Remarks 18.5 and 18.6 for details.

18.2.3 Minimax Estimation Under the RIP

Assuming the matrix RIP to hold and Gaussian noise ε , one can show that the minimax risk for recovering a Hermitian rank k matrix is

$$\inf_{\hat{\theta}} \sup_{\theta \in R(k)} \mathbb{E}_{\theta} \|\hat{\theta} - \theta\|_F^2 \simeq \sigma^2 \frac{dk}{n}, \quad (18.25)$$

where \simeq denotes two-sided inequality up to universal constants.

For the upper bound one can use the nuclear norm minimisation procedure or matrix Dantzig selector from Candès and Plan [6], and needs n to be large enough so that the matrix RIP holds with $\tau_n(k) < c_0$ where c_0 is a small enough numerical constant. Such an estimator $\tilde{\theta}$ then satisfies, for every $\theta \in R(k)$ and those $n \in \mathbb{N}$ for which $\tau_n(k) < c_0$,

$$\|\tilde{\theta} - \theta\|_F^2 \leq D(\delta) \sigma^2 \frac{kd}{n}, \quad (18.26)$$

with probability greater than $1 - 2\delta$, and with the constant $D(\delta)$ depending on δ and also on c_0 (suppressed in the notation). Note that the results in Ref. [6] use a different scaling in sample size in their Theorem 2.4, but eq. (II.7) in that reference explains that this is just a question of renormalisation. The same result holds for randomly sampled ‘Pauli bases’, see Ref. [28] (and take note of the slightly different normalisation in the notation there, too), and also for the Bernoulli noise model from Sect. 18.2.2.2, see Ref. [12].

A key interpretation for quantum tomography applications is that, instead of having to measure all $n = d^2$ basis coefficients $\text{tr}(E_i \theta)$, $i = 1, \dots, d^2$, a number

$$n \approx kd \overline{\log d}$$

of *randomly* chosen basis measurements is sufficient to reconstruct θ in Frobenius norm loss (up to a small error). In situations where d is large compared to k such a gain can be crucial.

Remark 18.1 (Uniqueness) It is worth noting that in the absence of errors, so when $Y_0 = \mathcal{X}\theta_0$ in terms of the sampling operator of Eq. (18.4), the quantum shape constraint ensures that under a suitable RIP condition, only the single matrix θ_0 is compatible with the data. More specifically, let $Y_0 = \mathcal{X}\theta_0$ for some $\theta_0 \in \Theta_+$ of rank k , and assume that \mathcal{X} satisfies RIP with $\tau_n(4k) < \sqrt{2} - 1$. Then

$$\{\theta \in \Theta_+ : \mathcal{X}\theta = Y_0\} = \{\theta_0\}. \quad (18.27)$$

This is a direct consequence of Theorem 3.2 in Ref. [33], which states that if RIP is satisfied with $\tau_n(4k) < \sqrt{2} - 1$ and $Y_0 = \mathcal{X}\theta_0$, the unique solution of

$$\begin{aligned} & \operatorname{argmin} \quad \|\theta\|_{S_1} \\ & \text{subject to } \mathcal{X}\theta = Y_0 \end{aligned} \tag{18.28}$$

is given by θ_0 . If $\theta_0 \in \Theta_+$, then the minimisation can be replaced by (compare also Ref. [23]).

$$\begin{aligned} & \operatorname{argmin} \quad \operatorname{tr}(\theta) \\ & \text{subject to } \mathcal{X}\theta = Y_0, \theta \geq 0, \end{aligned} \tag{18.29}$$

giving rise to the above remark. This observation further signifies the role played by the quantum shape constraint.

18.3 Uncertainty Quantification for Low-Rank Matrix Recovery

We now turn to the problem of quantifying the uncertainty of estimators $\tilde{\theta}$ that satisfy the risk bound (18.26). In fact the confidence sets we construct could be used for any estimator of θ , but the conclusions are most interesting when used for minimax optimal estimators $\tilde{\theta}$. For the main flow of ideas we shall assume $\varepsilon = (\varepsilon_1, \dots, \varepsilon_n)^T \sim N(0, \sigma^2 I_n)$ but the results hold for the Bernoulli measurement model from Sect. 18.2.2.2 as well—this is summarised in Remark 18.5.

From a statistical point of view, we phrase the problem at hand as the one of constructing a confidence set for θ : a data-driven subset C_n of $\mathbb{M}_d(\mathbb{C})$ that is ‘centred’ at $\tilde{\theta}$ and that satisfies

$$\mathbb{P}_\theta(\theta \in C_n) \geq 1 - \alpha, \quad 0 < \alpha < 1,$$

for a chosen ‘coverage’ or significance level $1 - \alpha$, and such that the Frobenius norm diameter $|C_n|_F$ reflects the accuracy of estimation, that is, it satisfies, with high probability,

$$|C_n|_F^2 \approx \|\tilde{\theta} - \theta\|_F^2.$$

In particular such a confidence set provides, through its diameter $|C_n|_F$, a data-driven estimate of how well the algorithm has recovered the true matrix θ in Frobenius-norm loss, and in this sense provides a quantification of the uncertainty in the estimate.

In the situation of an experimentalist this can be used to decide sequentially whether more measurements should be taken (to improve the recovery rate), or whether a satisfactory performance has been reached. Concretely, if for some $\epsilon > 0$ a recovery level $\|\hat{\theta} - \theta\|_F \leq \epsilon$ is desired for an estimator $\hat{\theta}$, then assuming $\tilde{\theta}$ satisfies the minimax optimal risk bound dk/n from (18.26), we expect to need, ignoring constants,

$$\frac{dk}{n} < \epsilon^2 \text{ and hence at least } n > \frac{dk}{\epsilon^2}$$

measurements. Note that we also need the RIP to hold with $\tau_n(k)$ from (18.8) less than a small constant c_0 , which requires the same number of measurements, increased by a further poly-log factor of d (and independently of σ).

Since the rank k of θ remains unknown after estimation we cannot obviously guarantee that the recovery level ϵ has been reached after a given number of measurements. A confidence set C_n for $\tilde{\theta}$ provides such certificates with high probability, by checking whether $|C_n|_F \leq \epsilon$, and by continuing to take further measurements if not. The main goal is then to prove that a sequential procedure based on C_n does *not* require more than approximately

$$n > \frac{dk \overline{\log d}}{\epsilon^2}$$

samples (with high probability). We construct confidence procedures in the following subsections that work with at most as many measurements, for the designs from Condition 18.1.

18.3.1 Adaptive Sequential Sampling

Before we describe our confidence procedures, let us make the following definition, where we recall that $R(k)$ denotes the set of $d \times d$ Hermitian matrices of rank at most $k \leq d$.

Definition 18.1 Let $\epsilon > 0, \delta > 0$ be given constants. An algorithm \mathcal{A} returning a $d \times d$ matrix $\hat{\theta}$ after $\hat{n} \in \mathbb{N}$ measurements in model (18.3) is called an (ϵ, δ) —adaptive sampling procedure if, with \mathbb{P}_θ -probability greater than $1 - \delta$, the following properties hold for every $\theta \in R(k)$ and every $1 \leq k \leq d$:

$$\|\hat{\theta} - \theta\|_F \leq \epsilon, \tag{18.30}$$

and, for positive constants $C(\delta), \gamma$, the stopping time \hat{n} satisfies

$$\hat{n} \leq C(\delta) \frac{kd(\log d)^\gamma}{\epsilon^2}. \tag{18.31}$$

Such an algorithm provides recovery at given accuracy level ϵ with \hat{n} measurements of minimax optimal order of magnitude (up to a poly-log factor), and with probability greater than $1 - \delta$. The sampling algorithm is adaptive since it does not require the knowledge of k , and since the number of measurements required depends only on k and not on the ‘worst case’ rank d .

The construction of non-asymptotic confidence sets C_n for θ at any sample size n in the next subsections will imply that such algorithms exist for low rank matrix recovery problems. The main idea is to check sequentially, for a geometrically increasing number 2^m of samples, $m = 1, 2, \dots$, if the diameter $|C_{2^m}|_F$ of a confidence set exceeds ϵ . If this is not the case, the algorithm terminates. Otherwise one takes 2^{m+1} additional measurements and evaluates the diameter $|C_{2^{m+1}}|_F$. A precise description of the algorithm is given in the proof of the following theorem, which we detail for the case of ‘Pauli’ designs. The isotropic design case is discussed in Remark 18.9.

Theorem 18.1 *Consider observations in the model (18.3) under Condition 18.1(b) with $\theta \in \Theta_+$. Then an adaptive sampling algorithm in the sense of Definition 18.1 exists for any $\epsilon, \delta > 0$.*

Remark 18.2 (Dependence in σ of Definition 18.1 and Theorem 18.1) Definition 18.1 and Theorem 18.1 are stated for the case where the standard deviation of the noise σ is assumed to be bounded by an absolute constant. It is straight-forward to modify the proofs to obtain a version where the dependency of the constants on the variance is explicit. Indeed, under Condition 1(a), Theorem 18.1 continues to hold if Eq. (18.31) is replaced by

$$\hat{n} \leq C(\delta) \frac{\sigma^2 kd(\log d)^\gamma}{\epsilon^2}.$$

For the ‘Pauli design case’—Condition 1(b)—Eq. (18.31) can be modified to

$$\hat{n} \leq C(\delta) \left(\frac{\sigma^2 kd(\log d)^\gamma}{\epsilon^2} \vee \frac{d(\log d)^\gamma}{\epsilon^2} \right).$$

Remark 18.3 (Necessity of the Quantum Shape Constraint) Note that the assumption $\theta \in \Theta_+$ in the previous theorem is necessary (in the case of Pauli design): Else the example of $\theta = 0$ or $\theta = E_i$ —where E_i is an arbitrary element of the Pauli basis—demonstrates that the number of measurements has to be at least of order d^2 : otherwise with positive probability, E_i is not drawn at a fixed sample size. On this event, both the measurements and $\hat{\theta}$ coincide under the laws \mathbb{P}_0 and \mathbb{P}_{E_i} , so we cannot have $\|\hat{\theta} - 0\|_F < \epsilon$ and $\|\hat{\theta} - E_i\|_F < \epsilon$ simultaneously for every $\epsilon > 0$, disproving existence of an adaptive sampling algorithm. In fact, the crucial condition for Theorem 18.1 to work is that the nuclear norms $\|\theta\|_{S_1}$ are bounded by an absolute constant (here = 1), which is violated by $\|E_i\|_{S_1} = \sqrt{d}$.

18.3.2 A Non-asymptotic Confidence Set Based on Unbiased Risk Estimation and Sample-Splitting

We suppose that we have two samples at hand, the first being used to construct an estimator $\tilde{\theta}$, such as the one from (18.26). We freeze $\tilde{\theta}$ and the first sample in what follows and all probabilistic statements are under the distribution \mathbb{P}_θ of the second sample Y , X of size $n \in \mathbb{N}$, conditional on the value of $\tilde{\theta}$. We define the following residual sum of squares statistic (recalling that σ^2 is known):

$$\hat{r}_n = \frac{1}{n} \|Y - \mathcal{X}\tilde{\theta}\|_F^2 - \sigma^2,$$

which satisfies $\mathbb{E}_\theta \hat{r}_n = \|\theta - \tilde{\theta}\|_F^2$ as is easily seen (see the proof of Theorem 18.2 below). Given $\alpha > 0$, let $\xi_{\alpha, \sigma}$ be quantile constants such that

$$\Pr \left(\sum_{i=1}^n (\varepsilon_i^2 - 1) > \xi_{\alpha, \sigma} \sqrt{n} \right) = \alpha \quad (18.32)$$

(these constants converge to the quantiles of a fixed normal distribution as $n \rightarrow \infty$), let $z_\alpha = \log(3/\alpha)$ and, for $z \geq 0$ a fixed constant to be chosen, define the confidence set

$$C_n = \left\{ v \in \mathbb{H}_d(\mathbb{C}) : \|v - \tilde{\theta}\|_F^2 \leq 2 \left(\hat{r}_n + z \frac{d}{n} + \frac{\bar{z} + \xi_{\alpha/3, \sigma}}{\sqrt{n}} \right) \right\}, \quad (18.33)$$

where

$$\bar{z}^2 = \bar{z}^2(\alpha, d, n, \sigma, v) = z_{\alpha/3} \sigma^2 \max(3\|v - \tilde{\theta}\|_F^2, 4zd/n).$$

Note that in the ‘quantum shape constraint’ case we can always bound $\|v - \tilde{\theta}\|_F \leq 2$ which gives a confidence set that is easier to compute and of only marginally larger overall diameter. In many important situations, however, the quantity \bar{z}/\sqrt{n} is of smaller order than $1/\sqrt{n}$, and the more complicated expression above is preferable.

It is not difficult to see (using that $x^2 \lesssim y + x/\sqrt{n}$ implies $x^2 \lesssim y + 1/n$) that the square Frobenius norm diameter of this confidence set is, with high probability, of order

$$\|C_n\|_F^2 \lesssim \|\tilde{\theta} - \theta\|_F^2 + \frac{zd + z_{\alpha/3}}{n} + \frac{\xi_{\alpha/3, \sigma}}{\sqrt{n}}. \quad (18.34)$$

Whenever $d \geq \sqrt{n}$ —so as long as at most $n \leq d^2$ measurements have been taken—the deviation terms are of smaller order than kd/n , and hence C_n has minimax optimal expected squared diameter whenever the estimator $\tilde{\theta}$ is minimax optimal as

in (18.26). Improvements for $d < \sqrt{n}$, corresponding to $n > d^2$ measurements, will be discussed in the next subsections.

The following result shows that C_n is an honest confidence set for arbitrary $d \times d$ matrices (without any rank constraint). Note that the result is non-asymptotic—it holds for every $n \in \mathbb{N}$.

Theorem 18.2 *Let $\theta \in \mathbb{H}_d(\mathbb{C})$ be arbitrary and let \mathbb{P}_θ be the distribution of Y, X from model (18.3).*

(a) *Assume Condition 18.1(a) and let C_n be given by (18.33) with $z = 0$. We then have for every $n \in \mathbb{N}$ that*

$$\mathbb{P}_\theta(\theta \in C_n) \geq 1 - \frac{2\alpha}{3} - 2e^{-cn}$$

where c is a numerical constant. In the case of standard Gaussian design, $c = 1/24$ is admissible.

(b) *Assume Condition 18.1(b), let C_n be given by (18.33) with $z > 0$ and assume also that $\theta \in \Theta_+$ and $\tilde{\theta} \in \Theta_+$ (that is, both satisfy the ‘quantum shape constraint’). Then for every $n \in \mathbb{N}$,*

$$\mathbb{P}_\theta(\theta \in C_n) \geq 1 - \frac{2\alpha}{3} - 2e^{-C(K)z}$$

where, for K the coherence constant of the basis,

$$C(K) = \frac{1}{(16 + 8/3)K^2}.$$

In Part (a), if we want to control the coverage probability at level $1 - \alpha$, n needs to be large enough so that the third deviation term is controlled at level $\alpha/3$. In the Gaussian design case with $\alpha = 0.05$, $n \geq 100$ is sufficient, for smaller sample sizes one can reduce the coverage level. The bound in (b) is entirely non-asymptotic (using the quantum constraint) for suitable choices of z . Also note that the quantile constants z, z_α, ξ_α all scale at least as $O(\log(1/\alpha))$ in the desired coverage level $\alpha \rightarrow 0$.

Remark 18.4 (Dependence of the Confidence Set’s Diameter on K (Pauli Design) and σ) Note that in the case of the Pauli design from Condition 1(b), the confidence set’s diameter depends on K only through the potential dependence of $\|\theta - \tilde{\theta}\|_F^2$ on K —the constants involved in the construction of \tilde{C}_n and on the bound on its diameter do not depend on K . On the other hand, the coverage probability of the confidence set depends on K , see Theorem 18.2, (b).

In this paper we assume that σ is a universal constant, and so as such it does not appear in Eqs. (18.33) and (18.34). It can however be interesting to investigate the

dependence in σ . In the case of isotropic design from Condition 1(a), we could set

$$C_n = \left\{ v \in \mathbb{H}_d(\mathbb{C}) : \|v - \tilde{\theta}\|_F^2 \leq 2 \left(\hat{r}_n + zc\sigma^2 \frac{d}{n} + \sigma^2 \frac{\bar{z} + \xi_{\alpha/3, \sigma}}{\sqrt{n}} \right) \right\},$$

(where σ^2 could be replaced by twice the plug-in estimator of σ^2 , using $\hat{\theta}$) and one would get

$$\mathbb{E}_\theta |C_n|_F^2 \lesssim \|\tilde{\theta} - \theta\|_F^2 + \sigma^2 \frac{zd + z_{\alpha/3}}{n} + \sigma^2 \frac{\xi_{\alpha/3, \sigma}}{\sqrt{n}},$$

and Theorem 18.2 also holds by introducing minor changes in the proof. In the case of the Pauli design from Condition 1(b), we could set

$$C_n = \left\{ v \in \mathbb{H}_d(\mathbb{C}) : \|v - \tilde{\theta}\|_F^2 \leq 2 \left(\hat{r}_n + zc \frac{d}{n} + \sigma^2 \frac{\bar{z} + \xi_{\alpha/3, \sigma}}{\sqrt{n}} \right) \right\},$$

(where σ^2 could be replaced by twice the plug-in estimator of σ^2 , using $\hat{\theta}$) and one would get

$$\mathbb{E}_\theta |C_n|_F^2 \lesssim \|\tilde{\theta} - \theta\|_F^2 + \frac{zd + z_{\alpha/3}}{n} + \sigma^2 \frac{\xi_{\alpha/3, \sigma}}{\sqrt{n}},$$

and Theorem 18.2 also holds by introducing minor changes in the proof. In this case we do not get a full dependence in σ as in the isotropic design case from Condition 1(a). However if $k^2 d \lesssim n$, we could also obtain a result similar to the one for the Gaussian design, using part (c) of Lemma 18.1.

Remark 18.5 (Bernoulli Noise) Theorem 18.2(b) holds as well for the Bernoulli measurement model from Sect. 18.2.2.2 with $T \geq d^2$, with slightly different constants in the construction of C_n and the coverage probabilities. See Remark 18.10 after the proof of Theorem 18.2(b) below. The modified quantile constants z , z_α , ξ_α still scale as $O(\sqrt{1/\alpha})$ in the desired coverage level $\alpha \rightarrow 0$, and hence the adaptive sampling Theorem 18.1 holds for such noise too, if the number T of preparations of the quantum state exceeds d^2 .

Remark 18.6 (Unknown Variance) The above confidence set C_n can be constructed with $\tilde{r}_n = \frac{1}{n} \|Y - \mathcal{X}\tilde{\theta}\|^2$ replacing \hat{r}_n —so without requiring knowledge of σ —if an a priori bound $\sigma^2 \leq vd/n$ is available, with v a known constant. An example of such a situation was discussed at the end of Sect. 18.2.2.2 above in quantum tomography problems: when $T \geq n$, the constant z should be increased by v in the construction of C_n , and the coverage proof goes through as well by compensating for the centring at $\mathbb{E}\varepsilon_i^2 = \sigma^2$ by the additional deviation constant v .

Remark 18.7 (Anisotropic Design Instead of Condition 1(a)) It is also interesting to consider the case of anisotropic design. This case is not very different, when it comes

to confidence sets, than isotropic design, as long as the variance-covariance matrix of the anisotropic sub-Gaussian design is such that the ratio of its largest eigenvalue with the smallest eigenvalue is bounded. Lemma 18.1(a), which quantifies the effect of the design, would change as follows: There exist constants $c_-, c_+, c > 0$ that depend only on the variance-covariance matrix of the anisotropic sub-Gaussian design and that are such that

$$\Pr\left(c_-\|\vartheta\|_F^2 \leq \frac{1}{n}\|\mathcal{X}\vartheta\|^2 \leq c_+\|\vartheta\|_F^2\right) \geq 1 - 2e^{-cn}.$$

Using this instead of the inequality in Lemma 18.1(a) in the proof of Theorem 18.2, part (a) leads to a similar result as Theorem 18.2, part (a).

18.3.3 Improvements When $d \leq \sqrt{n}$

The confidence set from Theorem 18.2 is optimal whenever the desired performance of $\|\theta - \tilde{\theta}\|_F^2$ is no better than of order $1/\sqrt{n}$. From a minimax point of view we expect $\|\theta - \tilde{\theta}\|_F^2$ to be of order kd/n for low rank $\theta \in R(k)$. In absence of knowledge about $k \geq 1$ the confidence set from Theorem 18.2 can hence be guaranteed to be optimal whenever $d \geq \sqrt{n}$, corresponding to the important regime $n \leq d^2$ for sequential sampling algorithms. Refinements for measurement scales $n \geq d^2$ are also of interest—we present two optimal approaches in this subsection for the designs from Condition 18.1.

18.3.3.1 Isotropic Design and U -Statistics

Consider first isotropic i.i.d design from Condition 18.1(a), and an estimator $\tilde{\theta}$ based on an initial sample of size n (all statements that follow are conditional on that sample). Collect another n samples to perform the uncertainty quantification step. Define the U -statistic

$$\hat{R}_n = \frac{2}{n(n-1)} \sum_{i < j} \sum_{m,k} (Y_i X_{m,k}^i - \tilde{\theta}_{m,k})(Y_j X_{m,k}^j - \tilde{\theta}_{m,k}) \quad (18.35)$$

whose \mathbb{E}_θ -expectation, conditional on $\tilde{\theta}$, equals $\|\theta - \tilde{\theta}\|_F^2$ in view of

$$\mathbb{E} Y_i X_{m,k}^i = \mathbb{E} \sum_{m',k'} X_{m',k'}^i X_{m,k}^i \theta_{m',k'} = \theta_{m,k}.$$

Define

$$C_n = \left\{ v \in \mathbb{H}_d(\mathbb{R}) : \|v - \tilde{\theta}\|_F^2 \leq \hat{R}_n + z_{\alpha,n} \right\} \quad (18.36)$$

where

$$z_{\alpha,n} = \frac{C_1 \|\theta - \tilde{\theta}\|_F}{\sqrt{n}} + \frac{C_2 d}{n}$$

and $C_1 \geq \zeta_1 \|\theta\|_F$, $C_2 \geq \zeta_2 \|\theta\|_F^2$ with ζ_i constants depending on α, σ . Note that if $\theta \in \Theta_+$ then $\|\theta\|_F \leq 1$ can be used as an upper bound. In practice the constants ζ_i can be calibrated by Monte Carlo simulations (see the implementation section below), or chosen based on concentration inequalities for U -statistics (see Ref. [14], Theorem 4.4.8). This confidence set has expected diameter

$$\mathbb{E}_\theta |C_n|_F^2 \lesssim \|\tilde{\theta} - \theta\|_F^2 + \frac{C_1 + C_2 d}{n},$$

and hence is compatible with any minimax recover rate $\|\tilde{\theta} - \theta\|_F^2 \lesssim kd/n$ from (18.26), where $k \geq 1$ is now arbitrary. For suitable choices of ζ_i we now show that C_n also has non-asymptotic coverage.

Theorem 18.3 *Assume Condition 18.1(a), and let C_n be as in (18.36). For every $\alpha > 0$ we can choose $\zeta_i(\alpha) = O(\sqrt{1/\alpha})$, $i = 1, 2$, large enough so that for every $n \in \mathbb{N}$ we have*

$$\mathbb{P}_\theta(\theta \in C_n) \geq 1 - \alpha.$$

Remark 18.8 (Dependence of the Confidence Set's Diameter on σ) As what was noted in Remark 18.4, Theorem 18.3 does not make explicit the dependence on σ , which is assumed to be (bounded by) an universal constant. In order to take the dependence on σ into account, we could replace $z_{\alpha,n}$ in Eq. (18.36) by $\frac{C_1 \|\theta - \tilde{\theta}\|_F}{\sqrt{n}} + \sigma^2 \frac{C_2 d}{n}$ (where σ^2 could be replaced by twice the plug-in estimator of σ^2 , using $\hat{\theta}$), and we would get

$$\mathbb{E}_\theta |C_n|_F^2 \lesssim \|\tilde{\theta} - \theta\|_F^2 + \sigma^2 \frac{C_1 + C_2 d}{n},$$

and Theorem 18.3 also holds by introducing minor changes in the proof.

18.3.3.2 Re-averaging Basis Elements When $d \leq \sqrt{n}$

Consider the setting of Condition 18.1(b) where we sample uniformly at random from a (scaled) basis $\{dE_1, \dots, dE_d\}$ of $\mathbb{M}_d(\mathbb{C})$. When $d \leq \sqrt{n}$ we are taking

$n \geq d^2$ measurements, and there is no need to sample at *random* from the basis as we can measure each individual coefficient, possibly even multiple times. Repeatedly sampling a basis coefficient $\text{tr}(E_k\theta)$ leads to a reduction of the variance of the measurement by averaging. More precisely, when taking $n = md^2$ measurements for some (for simplicity integer) $m \geq 1$, and if $(Y_{k,l} : l = 1, \dots, m)$ are the measurements Y_i corresponding to the basis element $E_k, k \in \{1, \dots, d^2\}$, we can form averaged measurements

$$Z_k = \frac{1}{\sqrt{m}} \sum_{l=1}^m Y_{k,l} = \sqrt{md} \langle E_k, \theta \rangle_F + \epsilon_k, \quad \epsilon_k = \frac{1}{\sqrt{m}} \sum_{l=1}^m \epsilon_{k,l} \sim N(0, \sigma^2).$$

We can then define the new measurement vector $\tilde{Z} = (\tilde{Z}_1, \dots, \tilde{Z}_{d^2})^T$ (using also $m = n/d^2$)

$$\tilde{Z}_k = Z_k - \sqrt{n} \langle \tilde{\theta}, E_k \rangle = \sqrt{n} \langle E_k, \theta - \tilde{\theta} \rangle_F + \epsilon_k, \quad k = 1, \dots, d^2$$

and the statistic

$$\hat{R}_n = \frac{1}{n} \|\tilde{Z}\|_{\mathbb{R}^{d^2}}^2 - \frac{\sigma^2 d^2}{n}$$

which estimates $\|\theta - \tilde{\theta}\|_F^2$ with precision

$$\begin{aligned} \hat{R}_n - \|\theta - \tilde{\theta}\|_F^2 &= \frac{2}{\sqrt{n}} \sum_{k=1}^{d^2} \epsilon_k \langle E_k, \theta - \tilde{\theta} \rangle_F + \frac{1}{n} \sum_{k=1}^{d^2} (\epsilon_k^2 - \mathbb{E}\epsilon_k^2) \\ &= O_P \left(\frac{\sigma \|\theta - \tilde{\theta}\|_F}{\sqrt{n}} + \frac{\sigma^2 d}{n} \right). \end{aligned}$$

Hence, for z_α the quantiles of a $N(0, 1)$ distribution and $\xi_{\alpha,\sigma}$ as in (18.32) with d^2 replacing n there, we can define a confidence set

$$\bar{C}_n = \left\{ v \in \mathbb{H}_d(\mathbb{C}) : \|v - \tilde{\theta}\|_F^2 \leq \hat{R}_n + \frac{z_{\alpha/2} \sigma \|\theta - \tilde{\theta}\|_F}{\sqrt{n}} + \frac{\xi_{\alpha/2, \sigma d}}{n} \right\} \quad (18.37)$$

which has non-asymptotic coverage

$$\mathbb{P}_\theta(\theta \in \bar{C}_n) \geq 1 - \alpha$$

for every $n \in \mathbb{N}$, by similar (in fact, since Lemma 18.1 is not needed, simpler) arguments as in the proof of Theorem 18.2 below. The expected diameter of \bar{C}_n is

by construction

$$\mathbb{E}_\theta |\bar{C}_n|_F^2 \lesssim \|\theta - \tilde{\theta}\|_F^2 + \frac{\sigma^2 d}{n}, \quad (18.38)$$

now compatible with *any* rate of recovery kd/n , $1 \leq k \leq d$.

18.3.4 A Confidence Set in Trace Norm Under Quantum Shape Constraints

The confidence sets from the previous subsections are all valid in the sense that they contain information about the recovery of θ by $\tilde{\theta}$ in Frobenius norm $\|\cdot\|_F$. It is of interest to obtain results in stronger norms, such as for instance the nuclear norm $\|\cdot\|_{S_1}$, which is particularly meaningful for quantum tomography problems since it then corresponds to the total variation distance on the set of ‘probability density matrices’. In fact, since

$$\frac{1}{2} \|\theta - \tilde{\theta}\|_{S_1} = \sup_{\|X\|_{op}=1} \text{tr} \left(X(\theta - \tilde{\theta}) \right), \quad (18.39)$$

the nuclear norm has a clear interpretation in terms of the maximum probability with which two quantum states can be distinguished by arbitrary measurements.

The absence of the ‘Hilbert space geometry’ induced by the relationship of the Frobenius norm to the inner product $\langle \cdot, \cdot \rangle_F$ makes this problem significantly harder, both technically and from an information-theoretic point of view. In particular it appears that the quantum shape constraint $\theta \in \Theta_+$ is crucial to obtain any results whatsoever, and for the theoretical results presented here it will be more convenient to perform an asymptotic analysis where $\min(n, d) \rightarrow \infty$ (with o , O -notation to be understood accordingly).

Instead of Condition 18.1 we shall now consider any design (X^1, \dots, X^n) in model (18.3) that satisfies the matrix RIP (18.7) with

$$\tau_n(k) = c \sqrt{kd \frac{\overline{\log}(d)}{n}}. \quad (18.40)$$

As discussed above, this covers in particular the designs from Condition 18.1. We shall still use the convention discussed before Condition 18.1 that θ and the matrices X^i are such that $\text{tr}(X^i \theta)$ is always real-valued.

In contrast to the results from the previous section we shall now assume a minimal low rank constraint on the parameter space:

Condition 18.2 $\theta \in R^+(k) := R(k) \cap \Theta_+$ for some k satisfying

$$k \sqrt{\frac{d \log d}{n}} = o(1),$$

This in particular implies that the RIP holds with $\tau_n(k) = o(1)$. Given this minimal rank constraint $\theta \in R^+(k)$, we now show that it is possible to construct a confidence set C_n that adapts to any low rank $1 \leq k_0 < k$. Here we may choose $k = d$ but note that this forces $n \gg d^2$ (for Condition 18.2 to hold with $k = d$).

We assume that there exists an estimator $\tilde{\theta}_{\text{Pilot}}$ that satisfies, uniformly in $R(k_0)$ for any $k_0 \leq k$ and for n large enough,

$$\|\tilde{\theta}_{\text{Pilot}} - \theta\|_F^2 \leq D\sigma^2 \frac{k_0 d}{n} := \frac{r_n^2(k_0)}{4} \quad (18.41)$$

where $D = D(\delta)$ depends on δ , and where so-defined r_n will be used frequently below. Such estimators exist as has already been discussed before (18.26). We shall in fact require a little more, namely the following oracle inequality: for any k and any matrix S of rank $k \leq d$, with high probability and for n large enough,

$$\|\tilde{\theta}_{\text{Pilot}} - \theta\|_F \lesssim \|\theta - S\|_F + r_n(k), \quad (18.42)$$

which in fact implies (18.41). Such inequalities exist assuming the RIP and Condition 18.2, see, e.g., Theorem 2.8 in Ref. [6]. Starting from $\tilde{\theta}_{\text{Pilot}}$ one can construct (see Theorem 18.5 below) an estimator that recovers $\theta \in R(k)$ in nuclear norm at rate $k\sqrt{d/n}$, which is again optimal from a minimax point of view, even under the quantum constraint (as discussed, e.g., in Ref. [24]). We now construct an adaptive confidence set for θ centred at a suitable projection of $\tilde{\theta}_{\text{Pilot}}$ onto Θ_+ .

In the proof of Theorem 18.4 below we will construct estimated eigenvalues ($\hat{\lambda}_j, j = 1, \dots, d$) of θ (see after Lemma 18.3). Given those eigenvalues and $\tilde{\theta}_{\text{Pilot}}$, we choose \hat{k} to equal the smallest integer $\leq d$ such that there exists a rank \hat{k} matrix $\tilde{\theta}'$ for which

$$\|\tilde{\theta}' - \tilde{\theta}_{\text{Pilot}}\|_F \leq r_n(\hat{k}) \text{ and } 1 - \sum_{J \leq \hat{k}} \hat{\lambda}_J \leq 2\hat{k}\sqrt{d/n}$$

is satisfied. Such \hat{k} exists with high probability (since the inequalities are satisfied for the true θ and λ_j 's, as our proofs imply). Define next $\hat{\vartheta}$ to be the $\langle \cdot, \cdot \rangle_F$ -projection of $\tilde{\theta}_{\text{Pilot}}$ onto

$$R^+(2\hat{k}) := R(2\hat{k}) \cap \Theta_+$$

and note that, since $2\hat{k} \geq \hat{k}$,

$$\|\tilde{\theta}_{\text{Pilot}} - \hat{v}\|_F = \|\tilde{\theta}_{\text{Pilot}} - R^+(2\hat{k})\|_F \leq \|\tilde{\theta}_{\text{Pilot}} - \tilde{\theta}'\|_F \leq r_n(\hat{k}). \quad (18.43)$$

Finally define, for C a constant chosen below,

$$C_n = \left\{ v \in \Theta_+ : \|v - \hat{v}\|_{S_1} \leq C\sqrt{\hat{k}}r_n(\hat{k}) \right\}. \quad (18.44)$$

Theorem 18.4 *Assume Condition 18.2 for some $1 \leq k \leq d$, and let $\delta > 0$ be given. Assume that with probability greater than $1 - 2\delta/3$, (a) the RIP (18.7) holds with $\tau_n(k)$ as in (18.40) and (b) there exists an estimator $\tilde{\theta}_{\text{Pilot}}$ for which (18.42) holds. Then we can choose $C = C(\delta)$ large enough so that, for C_n as in the last display,*

$$\liminf_{\min(n,d) \rightarrow \infty} \inf_{\theta \in R^+(k)} \mathbb{P}_\theta(\theta \in C_n) \geq 1 - \delta.$$

Moreover, uniformly in $R^+(k_0)$, $1 \leq k_0 \leq k$, and with \mathbb{P}_θ -probability greater than $1 - \delta$,

$$|C_n|_{S_1} \lesssim \sqrt{k_0}r_n(k_0).$$

Theorem 18.4 should mainly serve the purpose of illustrating that the quantum shape constraint allows for the construction of an optimal trace norm confidence set that adapts to the unknown low rank structure. Implementation of C_n is not straightforward so Theorem 18.4 is mostly of theoretical interest. Let us also observe that in full generality a result like Theorem 18.4 cannot be proved *without* the quantum shape constraint. This follows from a careful study of certain hypothesis testing problems (combined with lower bound techniques for confidence sets as in Refs. [19, 30]). Precise results are subject of current research and will be reported elsewhere.

18.4 Simulation Experiments

In order to illustrate the methods from this paper, we present some numerical simulations. The setting of the experiments is as follows: A random matrix $\eta \in \mathbb{M}_d(\mathbb{C})$ of norm $\|\eta\|_F = R^{1/2}$ is generated according to two distinct procedures that we will specify later, and the observations are

$$\bar{Y}_i = \text{tr}(X^i \eta) + \varepsilon_i.$$

where the ε_i are i.i.d. Gaussian of mean 0 and variance 1. The observations are reparametrised so that η represents the ‘estimation error’ $\theta - \hat{\theta}$, and we investigate

how well the statistics

$$\hat{r}_n = \frac{1}{n} \|\bar{Y}\| - 1 \text{ and } \hat{R}_n = \frac{2}{n(n-1)} \sum_{i < j} \sum_{m,k} \bar{Y}_i X_{m,k}^i \bar{Y}_j X_{m,k}^j$$

estimate the ‘accuracy of estimation’ $\|\eta\|_F^2 = \|\theta - \hat{\theta}\|_F^2$, conditional on the value of $\hat{\theta}$. We will choose η in order to illustrate two extreme cases: a first one where the nuclear norm $\|\eta\|_{S_1}$ is ‘small’, corresponding to a situation where the quantum constraint is fulfilled; and a second one where the nuclear norm is large, corresponding to a situation where the quantum constraint is *not* fulfilled. More precisely we generate the parameter η in two ways:

- ‘Random Dirac’ case: set a single entry (with position chosen at random on the diagonal) of η to $R^{1/2}$, and all the other coordinates equal to 0.
- ‘Random Pauli’ case: Set η equal to a Pauli basis element chosen uniformly at random and then multiplied by $R^{1/2}$.

The designs that we consider are the Gaussian design, and the Pauli design, described in Condition 1. We perform experiments with $d = 32$, $R \in \{0.1, 1\}$ and

$$n \in \{100, 200, 500, 1000, 2000, 5000\}.$$

Note that $d^2 = 1024$, so that the first four choices of n correspond to the important regime $n < d^2$. Our results are plotted as a function of the number n of samples in Figs. 18.2, 18.3, 18.4, and 18.5. The solid red and blue curves are the median errors of the normalised estimation errors

$$\frac{\sqrt{\hat{R}_n - R}}{R^{1/2}}, \quad \text{and} \quad \frac{\sqrt{\hat{r}_n - R}}{R^{1/2}},$$

after 1000 iterations, and the dotted lines are respectively, the (two-sided) 90% quantiles. We also report (see Tables 18.1, 18.2, 18.3, and 18.4) how well the confidence sets based on these estimates of the norm perform in terms of coverage probabilities, and of diameters. The diameters are computed as

$$\left(\hat{R}_n + \frac{C_{\text{UStat}} d}{n} + \frac{C'_{\text{UStat}} \hat{R}_n^{1/2}}{\sqrt{n}} \right)^{1/2},$$

for the U-Statistic approach and

$$\left(\hat{r}_n + \frac{C_{\text{RSS}}}{\sqrt{n}} + \frac{C'_{\text{RSS}} \hat{r}_n^{1/2}}{\sqrt{n}} \right)^{1/2},$$




Fig. 18.2 Gaussian design, and random Dirac (a single entry, chosen at random, is non-zero on the diagonal) η , with $R = 0.1$ (left picture) and $R = 1$ (right picture)




Fig. 18.3 Gaussian design, and random Pauli η , with $R = 0.1$ (left picture) and $R = 1$ (right picture)




Fig. 18.4 Pauli design, and random Dirac (a single entry, chosen at random, is non-zero on the diagonal) η , with $R = 0.1$ (left picture) and $R = 1$ (right picture)




Fig. 18.5 Pauli design, and random Pauli η , with $R = 0.1$ (left picture) and $R = 1$ (right picture)

Table 18.1 Gaussian design, and random Dirac (a single entry, chosen at random, is non-zero on the diagonal) η , with $R = 0.1$ (left table) and $R = 1$ (right table)

n	$R = 0.1$						$R = 1$					
	100	200	500	1000	2000	5000	100	200	500	1000	2000	5000
Coverage U-Stat	0.97	0.98	0.99	1.00	1.00	1.00	0.93	0.96	0.97	0.98	0.98	0.98
Diameter U-Stat	1.10	0.64	0.34	0.24	0.18	0.14	2.43	1.84	1.44	1.27	1.17	1.10
Coverage RSS	0.97	0.97	0.98	0.98	0.98	0.98	0.99	0.99	0.99	0.99	0.99	0.99
Diameter RSS	0.38	0.31	0.23	0.19	0.16	0.14	1.69	1.49	1.32	1.22	1.16	1.10

Table 18.2 Gaussian design, and random Pauli η , with $R = 0.1$ (left table) and $R = 1$ (right table)

n	$R = 0.1$						$R = 1$					
	100	200	500	1000	2000	5000	100	200	500	1000	2000	5000
Coverage U-Stat	0.98	0.98	0.99	0.99	1.0	1.0	0.93	0.95	0.97	0.98	0.98	0.98
Diameter U-Stat	1.10	0.62	0.34	0.24	0.18	0.14	2.40	1.83	1.43	1.27	1.18	1.10
Coverage RSS	0.98	0.98	0.97	0.97	0.97	0.97	0.99	0.99	0.99	0.99	1.00	1.00
Diameter RSS	0.39	0.31	0.23	0.19	0.17	0.14	1.71	1.49	1.31	1.22	1.16	1.10

Table 18.3 Pauli design, and random Dirac (a single entry, chosen at random, is non-zero on the diagonal) η , with $R = 0.1$ (left table) and $R = 1$ (right table)

n	$R = 0.1$						$R = 1$					
	100	200	500	1000	2000	5000	100	200	500	1000	2000	5000
Coverage U-Stat	0.97	0.98	0.98	0.99	0.98	0.98	0.85	0.54	0.69	0.69	0.70	0.71
Diameter U-Stat	1.10	0.63	0.34	0.24	0.18	0.14	2.28	1.87	1.43	1.26	1.18	1.10
Coverage RSS	0.96	0.96	0.96	0.96	0.97	0.97	0.88	0.89	0.88	0.88	0.88	0.88
Diameter RSS	0.39	0.29	0.23	0.19	0.16	0.14	1.70	1.50	1.30	1.21	1.16	1.10

Table 18.4 Pauli design, and random Pauli η , with $R = 0.1$ (left table) and $R = 1$ (right table)

n	$R = 0.1$						$R = 1$					
	100	200	500	1000	2000	5000	100	200	500	1000	2000	5000
Coverage U-Stat	0.97	0.97	0.96	0.86	0.65	0.58	0.82	0.22	0.25	0.27	0.30	0.37
Diameter U-Stat	1.09	0.57	0.34	0.25	0.18	0.15	2.45	2.09	1.33	1.38	1.19	1.09
Coverage RSS	0.93	0.86	0.77	0.77	0.77	0.77	0.12	0.19	0.40	0.63	0.56	0.53
Diameter RSS	0.38	0.29	0.22	0.19	0.16	0.14	1.71	1.56	1.31	1.26	1.14	1.08

for the RSS approach, where we have chosen $C_{\text{UStat}} = 2.5$, $C_{\text{RSS}} = 1$ and $C'_{\text{UStat}} = C_{\text{RSS}} = 6$ for all experiments—calibrated to a 95% coverage level.

From these numerical results, several observations can be made:

- In Gaussian random designs, the results are insensitive to the nature of η (see Figs. 18.2 and 18.3 and Tables 18.1 and 18.2). This is not surprising since the Gaussian design is ‘isotropic’.
- For Pauli designs with the quantum constraint (see Fig. 18.4 and Table 18.3) the RSS method works quite well even for small sample sizes. But the U-Stat method is not very reliable—indeed we see no empirical evidence that Theorem 18.3 should also hold true for Pauli design.
- For Pauli design and when the quantum shape constraint is *not* satisfied our methods cease to provide reliable results (see Fig. 18.5 and in particular Table 18.4). Indeed, when the matrix η is chosen itself as a random Pauli (which is the hardest signal to detect under Pauli design) both the RSS and the U-Stat approach perform poorly. The confidence set are not honest anymore, which is in line with the theoretical limitations we observe in Theorem 18.2. Figure 18.5 illustrates that the methods do not detect the signal, since the norm of η is largely under-evaluated for small sample sizes. These limitations are less pronounced when $n \geq d^2$. In this case one could use alternatively the re-averaging approach from Sect. 18.3.3.2 (not investigated in the simulations) to obtain honest results without the quantum shape constraint.

18.5 Proofs

18.5.1 Proof of Theorem 18.1

Proof Before we define the algorithm and prove the result, a few preparatory remarks are required: Our sequential procedure will be implemented in $m = 1, 2, \dots, T$ potential steps, in each of which $2 \cdot 2^m = 2^{m+1}$ measurements are taken. The arguments below will show that we can restrict the search to at most

$$T = O(\log(d/\epsilon))$$

steps. We also note that from the discussion after (18.7)—in particular since $c = c(\delta)$ from (18.8) is $O(1/\delta^2)$ —a simple union bound over $m \leq T$ implies that the RIP holds with probability $\geq 1 - \delta'$, some $\delta' > 0$, *simultaneously* for every $m \leq T$ satisfying $2^m \geq c'kd \log d$, and with $\tau_{2^m}(k) < c_0$, where c' is a constant that depends on δ', c_0 only. The maximum over $T = O(\log(d/\epsilon))$ terms is absorbed in a slightly enlarged poly-log term. Hence, simultaneously for all such sample sizes $2^m, m \leq T$, a nuclear norm regulariser exists that achieves the optimal rate from (18.26) with $n = 2^m$ and for every $k \leq d$, with probability greater than $1 - \delta/3$. Projecting this estimator onto Θ_+ changes the Frobenius error only by a universal multiplicative

constant (arguing as in (18.43) below), and we denote by $\tilde{\theta}_{2^m} \in \Theta_+$ the resulting estimator computed from a sample of size 2^m .

We now describe the algorithm at the m -th step: Split the 2^{m+1} observations into two halves and use the first subsample to construct $\tilde{\theta}_{2^m} \in \Theta_+$ satisfying (18.26) with \mathbb{P}_θ -probability $\geq 1 - \delta/3$. Then use the other 2^m observations to construct a confidence set C_{2^m} for θ centred at $\tilde{\theta}_{2^m}$: if $2^m < d^2$ we take C_{2^m} from (18.33) and if $2^m \geq d^2$ we take C_{2^m} from (18.37)—in both cases of non-asymptotic coverage at least $1 - \alpha$, $\alpha = \delta/(3T)$. If $|C_{2^m}|_F \leq \epsilon$ we terminate the procedure ($m = \hat{m}$, $\hat{n} = 2^{\hat{m}+1}$, $\hat{\theta} = \tilde{\theta}_{2^{\hat{m}}}$), but if $|C_{2^m}|_F > \epsilon$ we repeat the above procedure with $2 \cdot 2^{m+1} = 2^{m+1+1}$ new measurements, etc., until the algorithm terminates, in which case we have used

$$\sum_{m \leq \hat{m}} 2^{m+1} \lesssim 2^{\hat{m}} \approx \hat{n}$$

measurements in total.

To analyse this algorithm, recall that the quantile constants z, z_α, ξ_α appearing in the confidence sets (18.33) and (18.37) for our choice of $\alpha = \delta/(3T)$ grow at most as $O(\log(1/\alpha)) = O(\log T) = o(\log d)$. In particular in view of (18.26) and (18.34) or (18.38) the algorithm necessarily stops at a ‘maximal sample size’ $n = 2^{T+1}$ in which the squared Frobenius risk of the maximal model ($k = d$) is controlled at level ϵ . Such $T \in \mathbb{N}$ is $O(\log(d/\epsilon))$ and depends on $\sigma, d, \epsilon, \delta$, hence can be chosen by the experimenter.

To prove that this algorithms works we show that the event

$$\left\{ \|\hat{\theta} - \theta\|_F^2 > \epsilon^2 \right\} \cup \left\{ \hat{n} > \frac{C(\delta)kd(\log d)^\gamma}{\epsilon^2} \right\} = A_1 \cup A_2$$

has probability at most $2\delta/3$ for large enough $C(\delta), \gamma$. By the union bound it suffices to bound the probability of each event separately by $\delta/3$. For the first: Since \hat{n} has been selected we know $|C_{\hat{n}}|_F \leq \epsilon$ and since $\hat{\theta} = \tilde{\theta}_{\hat{n}}$ the event A_1 can only happen when $\theta \notin C_{\hat{n}}$. Therefore

$$\mathbb{P}_\theta(A_1) \leq \mathbb{P}_\theta(\theta \notin C_{\hat{n}}) \leq \sum_{m=1}^T \mathbb{P}_\theta(\theta \notin C_{2^m}) \leq \delta \frac{T}{3T} = \frac{\delta}{3}.$$

For A_2 , whenever $\theta \in R(k)$ and for all $m \leq T$ for which $2^m \geq c'kd\overline{\log d}$, we have, as discussed above, from (18.34) or (18.38) and (18.26) that

$$\mathbb{E}_\theta |C_{2^m}|_F^2 \leq D' \frac{kd \log T}{2^m},$$

where D' is a constant. In the last inequality the expectation is taken under the distribution of the sample used for the construction of C_{2^m} , and it holds on the event

on which $\tilde{\theta}_{2^m}$ realises the risk bound (18.26). Then let $C(\delta)$, γ be large enough so that $C(\delta)kd(\log d)^\gamma/\epsilon^2 \geq c'kd\overline{\log d}$ and let $m_0 \in \mathbb{N}$ be the smallest integer such that

$$2^{m_0} > \frac{C(\delta)kd(\log d)^\gamma}{\epsilon^2}.$$

Then, for $C(\delta)$ large enough and since $T = O(\log(d/\epsilon))$,

$$\mathbb{P}_\theta \left(\hat{n} > \frac{C(\delta)kd(\log d)^\gamma}{\epsilon^2} \right) \leq \mathbb{P}_\theta \left(|C_{2^{m_0}}|_F^2 > \epsilon^2 \right) \leq \frac{\mathbb{E}_\theta |C_{2^{m_0}}|_F^2}{\epsilon^2} \leq \frac{D' \log T}{C(\delta)(\log d)^\gamma} < \delta/3,$$

by Markov's inequality, completing the proof. \blacksquare

Remark 18.9 (Isotropic Sampling) The proof above works analogously for isotropic designs as defined in Condition 18.1a). When $2^m \geq d^2$, we replace the confidence set (18.37) in the above proof by the confidence set from (18.36). Assuming also that $\|\theta\|_F \leq M$ for some fixed constant M , we can construct a similar upper bound for T and the above proof applies directly (with T of slighter larger but still small enough order). Instead of assuming an upper bound on $\|\theta\|_F$ one can simply continue using the confidence set (18.33) also when $2^m \geq d^2$, in which case one has the slightly worse bound

$$\hat{n} \leq C(\delta) \max \left(\frac{kd\overline{\log d}}{\epsilon^2}, \frac{1}{\epsilon^4} \right)$$

for the number of measurements required.

18.5.2 Proof of Theorem 18.2

Proof By Lemma 18.1 below with $\vartheta = \tilde{\theta} - \theta$ the \mathbb{P}_θ -probability of the complement of the event

$$\mathcal{E} = \left\{ \left| \frac{1}{n} \|\mathcal{X}(\tilde{\theta} - \theta)\|^2 - \|\tilde{\theta} - \theta\|_F^2 \right| \leq \max \left(\frac{\|\theta - \tilde{\theta}\|_F^2}{2}, \frac{zd}{n} \right) \right\}$$

is bounded by the deviation terms $2e^{-cn}$ and $2e^{-C(K)z}$, respectively (note $z = 0$ in Case (a)). We restrict to this event in what follows. We can decompose

$$\hat{r}_n = \frac{1}{n} \|\mathcal{X}(\tilde{\theta} - \theta)\|^2 + \frac{2}{n} \langle \varepsilon, \mathcal{X}(\theta - \tilde{\theta}) \rangle + \frac{1}{n} \sum_{i=1}^n (\varepsilon_i^2 - \mathbb{E}\varepsilon_i^2) = A + B + C.$$

Since $\mathbb{P}(Y + Z < 0) \leq \mathbb{P}(Y < 0) + \mathbb{P}(Z < 0)$ for any random variables Y, Z we can bound the probability

$$\mathbb{P}_\theta(\theta \notin C_n, \mathcal{E}) = \mathbb{P}_\theta \left(\left\{ \frac{1}{2} \|\theta - \tilde{\theta}\|_F^2 > A + B + C + \frac{zd}{n} + \frac{\bar{z} + \xi_{\alpha/3, \sigma}}{\sqrt{n}} \right\}, \mathcal{E} \right)$$

by the sum of the following probabilities

$$I := \mathbb{P}_\theta \left(\left\{ \frac{1}{2} \|\theta - \tilde{\theta}\|_F^2 > \frac{1}{n} \|\mathcal{X}(\tilde{\theta} - \theta)\|^2 + \frac{zd}{n} \right\}, \mathcal{E} \right),$$

$$II := \mathbb{P}_\theta \left(\left\{ -\frac{1}{\sqrt{n}} \langle \varepsilon, \mathcal{X}(\theta - \tilde{\theta}) \rangle > \bar{z} \right\}, \mathcal{E} \right),$$

$$III := \mathbb{P}_\theta \left(-\frac{1}{\sqrt{n}} \sum_{i=1}^n (\varepsilon_i^2 - \mathbb{E}\varepsilon_i^2) > \xi_{\alpha/3, \sigma} \right).$$

The first probability I is bounded by

$$\begin{aligned} & \mathbb{P}_\theta \left(\left\{ -\frac{1}{n} \|\mathcal{X}(\tilde{\theta} - \theta)\|^2 + \|\theta - \tilde{\theta}\|_F^2 > \frac{1}{2} \|\theta - \tilde{\theta}\|_F^2 + \frac{zd}{n} \right\}, \mathcal{E} \right) \\ & \leq \mathbb{P}_\theta \left(\left\{ \left| \frac{1}{n} \|\mathcal{X}(\tilde{\theta} - \theta)\|^2 - \|\tilde{\theta} - \theta\|_F^2 \right| > \max \left(\frac{\|\theta - \tilde{\theta}\|_F^2}{2}, \frac{zd}{n} \right) \right\}, \mathcal{E} \right) = 0 \end{aligned}$$

About term II : Conditional on \mathcal{X} the variable $\frac{1}{\sqrt{n}} \langle \varepsilon, \mathcal{X}(\theta - \tilde{\theta}) \rangle$ is centred Gaussian with variance $(\sigma^2/n) \|\mathcal{X}(\theta - \tilde{\theta})\|^2$. The standard Gaussian tail bound then gives by definition of \bar{z} , and conditional on \mathcal{X} ,

$$\begin{aligned} & \leq \exp\{-\bar{z}^2/2(\sigma^2/n)\|\mathcal{X}(\theta - \tilde{\theta})\|^2\} \\ & = \exp \left\{ -\frac{z_{\alpha/3} \max(3\|\theta - \tilde{\theta}\|_F^2, 4zd/n)}{2\|\mathcal{X}(\theta - \tilde{\theta})\|^2/n} \right\} \leq \exp\{-z_{\alpha/3}\} = \alpha/3 \end{aligned}$$

since, on the event \mathcal{E} ,

$$\max(3\|\theta - \tilde{\theta}\|_F^2, 4zd/n) \geq (2/n)\|\mathcal{X}(\theta - \tilde{\theta})\|^2.$$

The overall bound for II follows from integrating the last but one inequality over the distribution of X . Term III is bounded by $\alpha/3$ by definition of $\xi_{\alpha, \sigma}$. ■

Remark 18.10 (Modification of the Proof for Bernoulli Errors) If instead of Gaussian errors we work with the error model from Sect. 18.2.2.2, we require a modified treatment of the terms II, III in the above proof. For the pure noise term III

we modify the quantile constants slightly to $\xi_{\alpha,\sigma} = \sqrt{(1/\alpha)}$. If the number T of preparations satisfies $T \geq 4d^2$ then Chebyshev's inequality and (18.24) give

$$\mathbb{P}_\theta \left(\left| \frac{1}{\sqrt{n}} \sum_{i=1}^n (\varepsilon_i^2 - \mathbb{E}\varepsilon_i^2) \right| > \xi_{\alpha/3,\sigma} \right) \leq \frac{\alpha}{3n} \sum_{i=1}^n \mathbb{E}\varepsilon_i^4 \leq \frac{\alpha}{3} \frac{4d^2}{T} \leq \frac{\alpha}{3}.$$

For the 'cross term' we have likewise with $z_\alpha = \sqrt{1/\alpha}$ and $a_i = (\mathcal{X}(\theta - \tilde{\theta}))_i$ that, on the event \mathcal{E} ,

$$\mathbb{P}_\varepsilon \left(\left\{ -\frac{1}{\sqrt{n}} \langle \varepsilon, \mathcal{X}(\theta - \tilde{\theta}) \rangle > \bar{z} \right\}, \mathcal{E} \right) \leq \frac{1}{n\bar{z}^2} \mathbb{E}_\varepsilon \left(\sum_{i=1}^n \varepsilon_i a_i 1_{\mathcal{E}} \right)^2 \leq \frac{d}{T\bar{z}^2} \frac{\|\mathcal{X}(\theta - \tilde{\theta})\|^2}{n} 1_{\mathcal{E}} \leq \alpha/3,$$

just as at the end of the proof of Theorem 18.2, so that coverage follows from integrating the last inequality w.r.t. the distribution of X . The scaling $T \approx d^2$ is similar to the one discussed in Theorem 3 in Ref. [12].

Lemma 18.1

(a) For isotropic design from Condition 18.1(a) and any fixed matrix $\vartheta \in \mathbb{H}_d(\mathbb{C})$ we have, for every $n \in \mathbb{N}$,

$$\Pr \left(\left| \frac{1}{n} \|\mathcal{X}\vartheta\|^2 - \|\vartheta\|_F^2 \right| > \frac{\|\vartheta\|_F^2}{2} \right) \leq 2e^{-cn}.$$

In the standard Gaussian design case we can take $c = 1/24$.

(b) In the 'Pauli basis' case from Condition 18.1(b) we have for any fixed matrix $\vartheta \in \mathbb{H}_d(\mathbb{C})$ satisfying the Schatten-1-norm bound $\|\vartheta\|_{S_1} \leq 2$ and every $n \in \mathbb{N}$,

$$\Pr \left(\left| \frac{1}{n} \|\mathcal{X}\vartheta\|^2 - \|\vartheta\|_F^2 \right| > \max \left(\frac{\|\vartheta\|_F^2}{2}, z \frac{d}{n} \right) \right) \leq 2 \exp \{-C(K)z\}$$

where $C(K) = 1/[(16 + 8/3)K^2]$, and where K is the coherence constant of the basis.

(c) In the 'Pauli basis' case from Condition 18.1(b) we have for any fixed matrix $\vartheta \in \mathbb{H}_d(\mathbb{C})$ such that the rank of ϑ is smaller than $2k$ and every $n \in \mathbb{N}$,

$$\Pr \left(\left| \frac{1}{n} \|\mathcal{X}\vartheta\|^2 - \|\vartheta\|_F^2 \right| > \max \left(\frac{\|\vartheta\|_F^2}{2}, z \frac{d}{n} \right) \right) \leq 2 \exp \left\{ -\frac{n}{17K^2 k^2 d} \right\}.$$

Proof We first prove the isotropic case. From (18.5) we see

$$\Pr \left(\left| \frac{1}{n} \|\mathcal{X}\vartheta\|^2 - \|\vartheta\|_F^2 \right| > \|\vartheta\|_F^2/2 \right) = \Pr \left(\left| \sum_{i=1}^n (Z_i^2 - \mathbb{E}Z_i^2) / \|\vartheta\|_F^2 \right| > n/2 \right)$$

where the $Z_i/\|\vartheta\|_F$ are sub-Gaussian random variables. Then the $Z_i^2/\|\vartheta\|_F^2$ are sub-exponential and we can apply Bernstein's inequality (Prop. 4.1.8 in Ref. [14]) to the last probability. We give the details for the Gaussian case and derive explicit constants. In this case $g_i := Z_i/\|\vartheta\|_F \sim N(0, 1)$ so the last probability is bounded, using Theorem 4.1.9 in Ref. [14], by

$$\Pr\left(\left|\sum_{i=1}^n (g_i^2 - 1)\right| > \frac{n}{2}\right) \leq 2 \exp\left\{-\frac{n^2/4}{4n + 2n}\right\},$$

and the result follows.

Under Condition 18.1(b), if we write $D = \max(n\|\vartheta\|_F^2/2, zd)$ we can reduce likewise to bound the probability in question by

$$\Pr\left(\left|\sum_{i=1}^n (Y_i - \mathbb{E}Y_1)\right| > D\right)$$

where the $Y_i = |\text{tr}(X^i \vartheta)|^2$ are i.i.d. bounded random variables. Using $\|E_i\|_{op} \leq K/\sqrt{d}$ from Condition 18.1(b) and the quantum constraint $\|\vartheta\|_F \leq \|\vartheta\|_{S_1} \leq 2$ we can bound

$$|Y_i| \leq d^2 \max_i \|E_i\|_{op}^2 \|\vartheta\|_{S_1}^2 \leq 4K^2 d := U$$

as well as

$$\mathbb{E}Y_i^2 \leq U \mathbb{E}|Y_i| \leq 4K^2 d \|\vartheta\|_F^2 := s^2.$$

Bernstein's inequality for bounded variables (e.g., Theorem 4.1.7 in Ref. [14]) applies to give the bound

$$2 \exp\left\{-\frac{D^2}{2ns^2 + \frac{2}{3}UD}\right\} \leq 2 \exp\{-C(K)z\},$$

after some basic computations, by distinguishing the two regimes of $D = n\|\vartheta\|_F^2/2 \geq zd$ and $D = zd \geq n\|\vartheta\|_F^2/2$.

Finally for (c), using the same reasoning as above and using $\|E_i\|_{op} \leq K/\sqrt{d}$ from Condition 18.1(b) and the fact that the estimator is also of rank less than k , we have $\|\vartheta\|_F \leq \|\vartheta\|_{S_1} \leq \sqrt{2k}\|\vartheta\|_F$ we can bound

$$|Y_i| \leq d^2 \max_i \|E_i\|_{op}^2 \|\vartheta\|_{S_1}^2 \leq 2K^2 kd \|\vartheta\|_F^2 := \tilde{U}$$

as well as

$$\mathbb{E}Y_i^2 \leq \tilde{U}\mathbb{E}|Y_i| \leq 2K^2dk^2\|\vartheta\|_F^4 := \tilde{s}^2.$$

Bernstein's inequality for bounded variables (e.g., Theorem 4.1.7 in Ref. [14]) applies to give the bound

$$2 \exp \left\{ -\frac{D^2}{2n\tilde{s}^2 + \frac{2}{3}\tilde{U}D} \right\} \leq 2 \exp \left\{ -\frac{n}{17K^2k^2d} \right\},$$

after some basic computations. ■

18.5.3 Proof of Theorem 18.3

Proof Since $\mathbb{E}_\theta \hat{R}_n = \|\theta - \tilde{\theta}\|_F^2$ we have from Chebyshev's inequality

$$\begin{aligned} \mathbb{P}_\theta(\theta \notin C_n) &\leq \mathbb{P}_\theta \left(|\hat{R}_n - \mathbb{E}\hat{R}_n| > z_{\alpha,n} \right) \\ &\leq \frac{\text{Var}_\theta(\hat{R}_n - \mathbb{E}\hat{R}_n)}{z_{\alpha,n}^2}. \end{aligned}$$

Now $U_n = \hat{R}_n - \mathbb{E}_\theta \hat{R}_n$ is a centred U-statistic and has Hoeffding decomposition $U_n = 2L_n + D_n$ where

$$L_n = \frac{1}{n} \sum_{i=1}^n \sum_{m,k} (Y_i X_{m,k}^i - \mathbb{E}_\theta[Y_i X_{m,k}^i])(\Theta_{m,k} - \tilde{\Theta}_{m,k})$$

is the linear part and

$$D_n = \frac{2}{n(n-1)} \sum_{i < j} \sum_{m,k} (Y_i X_{m,k}^i - \mathbb{E}_\theta[Y_i X_{m,k}^i])(Y_j X_{m,k}^j - \mathbb{E}[Y_j X_{m,k}^j])$$

the degenerate part. We note that L_n and D_n are orthogonal in $L^2(\mathbb{P}_\theta)$.

The linear part can be decomposed into

$$L_n = L_n^{(1)} + L_n^{(2)}$$

where

$$L_n^{(1)} = \frac{1}{n} \sum_{i=1}^n \sum_{m,k} \left(\sum_{m',k'} X_{m',k'}^i X_{m,k}^i \Theta_{m',k'} - \Theta_{m,k} \right) (\Theta_{m,k} - \tilde{\Theta}_{m,k})$$

and

$$L_n^{(2)} = \frac{1}{n} \sum_{i=1}^n \varepsilon_i \sum_{m,k} X_{m,k}^i (\Theta_{m,k} - \tilde{\Theta}_{m,k}).$$

Now by the i.i.d. assumption we have

$$\text{Var}_\theta(L_n^{(2)}) = \sigma^2 \frac{\|\tilde{\theta} - \theta\|_F^2}{n}.$$

Moreover, by transposing the indices m, k and m', k' in an arbitrary way into single indices $M = 1, \dots, d^2, K = 1, \dots, d^2, d^2 = p$, respectively, basic computations given before eq. (28) in Ref. [30] imply that the variance of the second term is bounded by

$$\text{Var}_\theta(L_n^{(1)}) \leq \frac{c \|\theta - \tilde{\theta}\|_F^2 \|\theta\|_F^2}{n}$$

where c is a constant that depends only on $\mathbb{E}X_{1,1}^4$ (which is finite since the $X_{1,1}$ are sub-Gaussian in view of Condition 18.1(a)). Moreover, the degenerate term satisfies

$$\text{Var}_\theta(D_n) \leq c \frac{d}{n^2} \|\theta\|_F^4$$

in view of standard U -statistic computations leading to eq. (6.6) in Ref. [21], with $d^2 = p$, and using the same transposition of indices as before. This proves coverage by choosing the constants in the definition of $z_{\alpha,n}$ large enough. ■

18.5.4 Proof of Theorem 18.4

We prove the result for symmetric matrices with real entries—the case of Hermitian matrices requires only minor (mostly notational) adaptations.

Given the estimator $\tilde{\theta}_{\text{Pilot}}$, we can easily transform it into another estimator $\tilde{\theta}$ for which the following is true.

Theorem 18.5 *There exists an estimator $\tilde{\theta}$ that satisfies, uniformly in $\theta \in R(k)$, for any $k \leq d$ and with \mathbb{P}_θ -probability greater than $1 - 2\delta/3$,*

$$\|\tilde{\theta} - \theta\|_F \leq r_n(k),$$

as well as,

$$\tilde{\theta} \in R(k),$$

and then also

$$\|\tilde{\theta} - \theta\|_{S_1} \leq \sqrt{2kr_n(k)}.$$

Proof Let $\tilde{\theta}_{\text{Pilot}}$ and let $\tilde{\theta}$ be the element of $R(d)$ with smallest rank k' such that

$$\|\tilde{\theta}_{\text{Pilot}} - \tilde{\theta}\|_F^2 \leq \frac{r_n^2(k')}{4}.$$

Such $\tilde{\theta}$ exists and has rank $\leq k$, with probability $\geq 1 - 2\delta/3$, since $\theta \in R(k)$ satisfies the above inequality in view of (18.41). The $\|\cdot\|_F^2$ -loss of $\tilde{\theta}$ is no larger than $r_n(k)$ by the triangle inequality

$$\|\tilde{\theta} - \theta\|_F \leq \|\tilde{\theta} - \tilde{\theta}_{\text{Pilot}}\|_F + \|\tilde{\theta}_{\text{Pilot}} - \theta\|_F,$$

and this completes the proof of the third claim in view of (18.2). \blacksquare

The rest of the proof consists of three steps: The first establishes some auxiliary empirical process type results, which are then used in the second step to construct a sufficiently good simultaneous estimate of the eigenvalues of θ . In Step III the coverage of the confidence set is established.

18.5.4.1 Step I

Let $\theta \in R^+(k) = R(k) \cap \Theta_+$ and let $\tilde{\theta}$ be the estimator from Theorem 18.5. Then with probability $\geq 1 - 2\delta/3$, and if $\eta = \tilde{\theta} - \theta$, we have

$$\|\eta\|_F^2 \leq r_n^2(k) \quad \forall \theta \in R^+(k), \quad (18.45)$$

and that

$$\eta \in R(2k).$$

For the rest of the proof we restrict in what follows to the event of probability greater than or equal to $1 - 2\delta/3$ described by (a) and (b) in the hypothesis of the theorem.

Write $Y'_i = Y_i - \text{tr}(X^i \tilde{\theta})$ for the ‘new observations’

$$Y'_i = \text{tr}(X^i \eta) + \varepsilon_i, \quad i = 1, \dots, n.$$

For any $d \times d'$ matrix V we set

$$\tilde{\gamma}_\eta(V) = V^T \left(\frac{1}{n} \sum_{i=1}^n X^i Y'_i \right) V$$

which estimates

$$\gamma_\eta(V) = V^T \eta V.$$

Let now U be any unit vector in \mathbb{R}^d . Then in the above notation ($d' = 1$) we can write

$$\begin{aligned} \tilde{\gamma}_\eta(U) &= \frac{1}{n} \sum_{i=1}^n \sum_{m,m' \leq d} U_m U_{m'} X_{m,m'}^i Y_i' \\ &= \frac{1}{n} \sum_{i=1}^n \sum_{m,m' \leq d} U_m U_{m'} X_{m,m'}^i (\text{tr}(X^i \eta) + \varepsilon_i) \\ &= \frac{1}{n} \sum_{i=1}^n \sum_{m,m' \leq d} U_m U_{m'} X_{m,m'}^i \left(\sum_{k,k' \leq d} X_{k,k'}^i \eta_{k,k'} + \varepsilon_i \right). \end{aligned}$$

If \mathbb{U} denotes the $d \times d$ matrix UU^T , the last quantity can be written as

$$\frac{1}{n} \langle \mathcal{X} \mathbb{U}, \mathcal{X} \eta \rangle + \frac{1}{n} \langle \mathcal{X} \mathbb{U}, \varepsilon \rangle.$$

We can hence bound, for $\mathcal{S} = \{U \in \mathbb{R}^d : \|U\|_2 = 1\}$

$$\begin{aligned} &\sup_{\eta \in R(2k), \|\eta\|_F \leq r_n(k), U \in \mathcal{S}} |\tilde{\gamma}_\eta(U) - \gamma_\eta(U)| \\ &\leq \sup_{\eta \in R(2k), \|\eta\|_F \leq r_n(k), U \in \mathcal{S}} \left| \frac{1}{n} \langle \mathcal{X} \mathbb{U}, \mathcal{X} \eta \rangle - \langle \mathbb{U}, \eta \rangle \right| + \sup_{U \in \mathcal{S}} \left| \frac{1}{n} \langle \mathcal{X} \mathbb{U}, \varepsilon \rangle \right|. \end{aligned}$$

Lemma 18.2 *The right hand side on the last inequality is, with probability greater than $1 - \delta$, of order*

$$v_n := O \left(r_n(k) \tau_n(k) + \sqrt{\frac{d}{n}} \right).$$

Proof The first term in the bound corresponds to the first supremum on the right hand side of the last inequality, and follows directly from the matrix RIP (and Lemma 18.4). For the second term we argue conditionally on the values of \mathcal{X} and on the event for which the matrix RIP is satisfied. We bound the supremum of the Gaussian process

$$\mathbb{G}_\varepsilon(U) := \frac{1}{\sqrt{n}} \langle \mathcal{X} \mathbb{U}, \varepsilon \rangle \sim N(0, \|\mathcal{X} \mathbb{U}\|^2/n)$$

indexed by elements U of the unit sphere \mathcal{S} of \mathbb{R}^d , which satisfies the metric entropy bound

$$\log N(\delta, \mathcal{S}, \|\cdot\|) \lesssim d \log(A/\delta)$$

by a standard covering argument. Moreover $\mathbb{U} = UU^T \in R(1)$ and hence for any pair of vectors $U, \bar{U} \in \mathcal{S}$ we have that $\mathbb{U} - \bar{\mathbb{U}} \in R(2)$. From the RIP we deduce for every fixed $U, \bar{U} \in \mathcal{S}$ that

$$\begin{aligned} \frac{1}{n} \|\mathcal{X}\mathbb{U} - \mathcal{X}\bar{\mathbb{U}}\|^2 &= \|\mathbb{U} - \bar{\mathbb{U}}\|_F^2 \left(1 + \frac{\frac{1}{n} \|\mathcal{X}(\mathbb{U} - \bar{\mathbb{U}})\|^2 - \|\mathbb{U} - \bar{\mathbb{U}}\|_F^2}{\|\mathbb{U} - \bar{\mathbb{U}}\|_F^2} \right) \\ &\leq (1 + \tau_n(2)) \|\mathbb{U} - \bar{\mathbb{U}}\|_F^2 \leq C \|U - \bar{U}\|^2 \end{aligned}$$

since $\tau_n(2) = O(1)$ and since

$$\|\mathbb{U} - \bar{\mathbb{U}}\|_F^2 = \sum_{m,m'} (U_m U_{m'} - \bar{U}_m \bar{U}_{m'})^2 = \sum_{m,m'} (U_m U_{m'} - U_m \bar{U}_{m'} + U_m \bar{U}_{m'} - \bar{U}_m \bar{U}_{m'})^2 \leq 2 \|U - \bar{U}\|^2.$$

Hence any δ -covering of \mathcal{S} in $\|\cdot\|$ induces a δ/C covering of \mathcal{S} in the intrinsic covariance $d_{\mathbb{G}_\varepsilon}$ of the (conditional on \mathcal{X}) Gaussian process \mathbb{G}_ε , i.e.,

$$\log N(\delta, \mathcal{S}, d_{\mathbb{G}_\varepsilon}) \lesssim d \log(A'/\delta)$$

with constants independent of X . By Dudley's metric entropy bound (e.g., Ref. [14]) applied to the conditional Gaussian process we have for $d > 0$ some constant

$$\mathbb{E} \sup_{U \in \mathcal{S}} |\mathbb{G}_\varepsilon(U)| \lesssim \int_0^d \sqrt{\log N(\delta, \mathcal{S}, d_{\mathbb{G}_\varepsilon})} d\delta \lesssim \sqrt{d}$$

and hence we deduce that

$$\mathbb{E}_\varepsilon \sup_{U \in \mathcal{S}} \frac{1}{n} |\langle \mathcal{X}U, \varepsilon \rangle| = \mathbb{E}_\varepsilon \frac{1}{\sqrt{n}} \sup_{U \in \mathcal{S}} |\mathbb{G}_\varepsilon(U)| \lesssim \sqrt{\frac{d}{n}} \quad (18.46)$$

with constants independent of X , so that the result follows from applying Markov's inequality. \blacksquare

18.5.4.2 Step II

Define the estimator

$$\hat{\theta}' = \tilde{\theta} + \frac{1}{n} \sum_{i=1}^n X^i Y_i' = \tilde{\theta} + \tilde{\gamma}_\eta(I_d).$$

Then we can write, using $U^T \tilde{\gamma}_\eta(I_d)U = \tilde{\gamma}_\eta(U)$,

$$\begin{aligned} U^T \hat{\theta}'U - U^T \theta U &= U^T (\tilde{\theta} + \tilde{\gamma}_\eta(I_d))U - U^T (\tilde{\theta} + \eta)U \\ &= \tilde{\gamma}_\eta(U) - \gamma_\eta(U), \end{aligned}$$

and from the previous lemma we conclude, for any unit vector U that with probability $\geq 1 - \delta$,

$$|U^T \hat{\theta}'U - U^T \theta U| \leq v_n.$$

Let now $\hat{\theta}$ be any symmetric positive definite matrix such that

$$|U^T \hat{\theta}U - U^T \hat{\theta}'U| \leq v_n.$$

Such a matrix exists, for instance $\theta \in R^+(k)$, and by the triangle inequality we also have

$$|U^T \hat{\theta}U - U^T \theta U| \leq 2v_n. \tag{18.47}$$

Lemma 18.3 *Let M be a symmetric positive definite $d \times d$ matrix with eigenvalues λ_j 's ordered such that $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_d$. For any $j \leq d$ consider an arbitrary collection of j orthonormal vectors $\mathcal{V}_j = (V^\iota : 1 \leq \iota \leq j)$ in \mathbb{R}^d . Then we have*

$$(a) \quad \lambda_{j+1} \leq \sup_{U \in \mathcal{S}, U \perp \text{span}(\mathcal{V}_j)} U^T M U,$$

and

$$(b) \quad \sum_{\iota \leq j} \lambda_\iota \geq \sum_{\iota \leq j} (V^\iota)^T M V^\iota.$$

Let \hat{R} be the rotation that diagonalises $\hat{\theta}$ such that $\hat{R}^T \hat{\theta} \hat{R} = \text{diag}(\hat{\lambda}_j : j = 1, \dots, d)$ ordered such that $\hat{\lambda}_j \geq \hat{\lambda}_{j+1} \forall j$. Moreover let R be the rotation that does the same for θ and its eigenvalues λ_j . We apply the previous lemma with $M = \hat{\theta}$ and \mathcal{V} equal to the column vectors $r_\iota : \iota \leq l - 1$ of R to obtain, for any fixed $l \leq j \leq d$,

$$\hat{\lambda}_l \leq \sup_{U \in \mathcal{S}, U \perp \text{span}(r_{\iota, \iota \leq l-1})} U^T \hat{\theta} U, \tag{18.48}$$

and also that

$$\sum_{l \leq j} \hat{\lambda}_l \geq \sum_{l \leq j} r_l^T \hat{\theta} r_l. \tag{18.49}$$

From (18.47) we deduce, that

$$\hat{\lambda}_l \leq \sup_{U \in \mathcal{S}, U \perp \text{span}(r_\iota, \iota \leq j-1)} U^T \theta U + 2v_n = \lambda_j + 2v_n \quad \forall l \leq j,$$

as well as

$$\sum_{l \leq j} \hat{\lambda}_l \geq \sum_{l \leq j} r_l^T \theta r_l - 2jv_n = \sum_{l \leq j} \lambda_l - 2jv_n,$$

with probability $\geq 1 - \delta$. Combining these bounds we obtain

$$\left| \sum_{l \leq j} \hat{\lambda}_l - \sum_{l \leq j} \lambda_l \right| \leq 2jv_n, \quad j \leq d. \quad (18.50)$$

18.5.4.3 Step III

We show that the confidence sets covers the true parameter on the event of probability $\geq 1 - \delta$ on which Steps I and II are valid, and for the constant C chosen large enough.

Let $\Pi = \Pi_{R^+(2\hat{k})}$ be the projection operator onto $R^+(2\hat{k})$. We have

$$\|\hat{\vartheta} - \theta\|_{S_1} \leq \|\hat{\vartheta} - \Pi\theta\|_{S_1} + \|\Pi\theta - \theta\|_{S_1}.$$

We have, using (18.50) and Lemma 18.5 below

$$\begin{aligned} \|\Pi\theta - \theta\|_{S_1} &= \sum_{J > 2\hat{k}} \lambda_J = 1 - \sum_{J \leq 2\hat{k}} \lambda_J \\ &\leq 1 - \sum_{J \leq 2\hat{k}} \hat{\lambda}_J + 4\hat{k}v_n \\ &\leq 6v_n\hat{k} \leq (C/2)\sqrt{\hat{k}}r_n(\hat{k}) \end{aligned}$$

for C large enough.

Moreover, using the oracle inequality (18.42) with $S = \Pi\theta$ and (18.43),

$$\begin{aligned} \|\hat{\vartheta} - \Pi\theta\|_{S_1} &\leq \sqrt{4\hat{k}}\|\hat{\vartheta} - \Pi\theta\|_F \\ &\leq \sqrt{4\hat{k}}(\|\hat{\vartheta} - \theta\|_F + \|\Pi\theta - \theta\|_F) \\ &\leq \sqrt{4\hat{k}}(\|\hat{\vartheta} - \tilde{\theta}_{\text{Pilot}}\|_F + \|\tilde{\theta}_{\text{Pilot}} - \theta\|_F + \|\Pi\theta - \theta\|_F) \\ &\lesssim \sqrt{\hat{k}}(r_n(\hat{k}) + \|\Pi\theta - \theta\|_F). \end{aligned}$$

We finally deal with the approximation error: Note

$$\|\Pi\theta - \theta\|_F^2 = \sum_{l>\hat{k}} \lambda_l^2 \leq \max_{l>2\hat{k}} |\lambda_l| \sum_{l>2\hat{k}} |\lambda_l|.$$

By (18.50) we know that

$$\sum_{l>\hat{k}} \lambda_l = 1 - \sum_{l\leq\hat{k}} \lambda_l \leq 1 - \sum_{l\leq\hat{k}} \hat{\lambda}_l + 2v_n\hat{k} \leq 4v_n\hat{k}.$$

Hence out of the λ_l 's with indices $l > \hat{k}$ there have to be less than \hat{k} coefficients which exceed $4v_n$. Since the eigenvalues are ordered this implies that the λ_l 's with indices $l > 2\hat{k}$ are all less than or equal to $4v_n$, and hence the quantity in the last but one display is bounded by (since $\hat{k} < 2\hat{k}$), using again (18.50) and the definition of \hat{k} ,

$$4v_n \left(1 - \sum_{l\leq\hat{k}} |\lambda_l|\right) \lesssim v_n \left(1 - \sum_{l\leq\hat{k}} |\hat{\lambda}_l|\right) + \hat{k}v_n^2 \lesssim v_n^2\hat{k} \lesssim \sqrt{\hat{k}}r_n(\hat{k}).$$

Overall we get the bound

$$\|\hat{\vartheta} - \Pi\theta\|_{S_1} \lesssim \hat{k}v_n \lesssim (C/2)\sqrt{\hat{k}}r_n(\hat{k})$$

for C large enough, which completes the proof of coverage of C_n by collecting the above bounds. The diameter bound follows from $\hat{k} \leq k$ (in view of the defining inequalities of \hat{k} being satisfied, for instance, for $\hat{\theta}' = \theta$, whenever $\theta \in R^+(k_0)$.)

18.6 Auxiliary Results

18.6.1 Proof of Lemma 18.3

- (a) Consider the subspaces $E = \text{span}((V^l)_{l\leq j})^\perp$ and $F = \text{span}((e_l)_{l\leq j+1})$ of \mathbb{R}^d , where the e_l 's are the eigenvectors of the $d \times d$ matrix M corresponding to eigenvalues λ_j . Since $\dim(E) + \dim(F) = (d - j) + j + 1 = d + 1$, we know that $E \cap F$ is not empty and there is a vectorial sub-space of dimension 1 in the intersection. Take $U \in E \cap F$ such that $\|U\| = 1$. Since $U \in F$, it can be written as

$$U = \sum_{l=1}^{j+1} u_l e_l$$

for some coefficients u_ι . Since the e_ι 's are orthogonal eigenvectors of the symmetric matrix M we necessarily have

$$MU = \sum_{\iota=1}^{j+1} \lambda_\iota u_\iota e_\iota,$$

and thus

$$U^T MU = \sum_{\iota=1}^{j+1} \lambda_\iota u_\iota^2.$$

Since the λ_ι 's are all non-negative and ordered in decreasing absolute value, one has

$$U^T MU = \sum_{\iota=1}^{j+1} \lambda_\iota u_\iota^2 \geq \lambda_{j+1} \sum_{\iota=1}^{j+1} u_\iota^2 = \lambda_{j+1} \|U\|^2 = \lambda_{j+1}.$$

Taking the supremum in U yields the result.

- (b) For each $\iota \leq j$, let us write the decomposition of V^ι on the basis of eigenvectors ($e_l : l \leq d$) of M as

$$V^\iota = \sum_{l \leq d} v_l^\iota e_l.$$

Since the (e_l) are the eigenvectors of M we have

$$\sum_{\iota \leq j} (V^\iota)^T M V^\iota = \sum_{\iota \leq j} \sum_{l=1}^d \lambda_l (v_l^\iota)^2,$$

where $\sum_{l=1}^d (v_l^\iota)^2 = 1$ and $\sum_{\iota \leq j} (v_l^\iota)^2 \leq 1$, since the V^ι are orthonormal. The last expression is maximised in $(v_l^\iota)_{\iota \leq j, 1 \leq l \leq d}$ and under these constraints, when $v_l^\iota = 1$ and $v_l^\iota = 0$ if $\iota \neq l$ (since the (λ_ι) are in decreasing order), and this gives

$$\sum_{\iota \leq j} (V^\iota)^T M V^\iota \leq \sum_{\iota \leq j} \lambda_\iota.$$

18.6.2 Some Further Lemmas

Lemma 18.4 *Under the RIP (18.7) we have for every $1 \leq k \leq d$ that, with probability at least $1 - \delta$,*

$$\sup_{A, B \in R(k)} \left| \frac{\frac{1}{n} \langle \mathcal{X}A, \mathcal{X}B \rangle - \langle A, B \rangle_F}{\|A\|_F \|B\|_F} \right| \leq 10\tau_n(k). \quad (18.51)$$

Proof The matrix RIP can be written as

$$\sup_{A \in R(k)} \left| \frac{\langle \mathcal{X}A, \mathcal{X}A \rangle}{n \langle A, A \rangle_F} - 1 \right| = \frac{|\langle A, (n^{-1}M - \mathbb{I})A \rangle_F|}{\langle A, A \rangle_F} \leq \tau_n(k), \quad (18.52)$$

for a suitable $M \in \mathbb{H}_{d^2}(\mathbb{C})$. The above bound then follows from applying the Cauchy-Schwarz inequality to

$$\frac{1}{n} \langle \mathcal{X}A, \mathcal{X}B \rangle - \langle A, B \rangle_F = \langle A, (n^{-1}M - \mathbb{I})B \rangle_F. \quad (18.53)$$

■

The following lemma can be proved by basic linear algebra, and is left to the reader.

Lemma 18.5 *Let $M \geq 0$ with positive eigenvalues $(\lambda_j)_j$ ordered in decreasing order. Denote with $\Pi_{R^+(j-1)}$ the projection onto $R^+(j-1) = R(j-1) \cap \Theta_+$. Then for any $2 \leq j \leq d$ we have*

$$\sum_{j' \geq j} \lambda_{j'} = \|M - \Pi_{R^+(j-1)}M\|_{S_1}.$$

Acknowledgements The work of A. Carpentier was done when she was in Cambridge and is partially supported by the Deutsche Forschungsgemeinschaft (DFG) Emmy Noether grant MuSyAD (CA 1488/1-1), by the DFG-314838170, GRK 2297 MathCoRe, by the DFG GRK 2433 DAEDALUS (384950143/GRK2433), and by the DFG CRC 1294 ‘Data Assimilation’, Project A03, and by the UFA-DFH through the French-German Doktorandenkolleg CDFA 01-18. D. Gross acknowledges support by the DFG (SPP1798 CoSIP), Germany’s Excellence Strategy—Cluster of Excellence Matter and Light for Quantum Computing (ML4Q) EXC 2004/1—390534769, and the ARO under contract W911NF-14-1-0098 (Quantum Characterization, Verification, and Validation). J. Eisert was supported by the DFG, CRC 183, Project B01, by the DFG GRK DAEDALUS, the DFG SPP1798 CoSIP, the ERC, and the Templeton Foundation.

References

1. L. Artiles, R. Gill, M. Guta, An invitation to quantum tomography. *J. R. Stat. Soc.* **67**, 109 (2005)
2. K. Audenaert, S. Scheel, Quantum tomographic reconstruction with error bars: a Kalman filter approach. *New J. Phys.* **11**(2), 023028 (2009)
3. P.J. Bickel, Y. Ritov, A.B. Tsybakov, Simultaneous analysis of lasso and Dantzig selector. *Ann. Stat.* **37**(4), 1705–1732 (2009)
4. R. Blume-Kohout, Robust error bars for quantum tomography (2012). arXiv:1202.5270
5. P. Bühlmann, S. van de Geer, *Statistics for High-Dimensional Data: Methods, Theory and Applications* (Springer, Berlin, 2011)
6. E.J. Candès, Y. Plan, Tight oracle inequalities for low-rank matrix recovery from a minimal number of noisy random measurements. *IEEE Trans. Inform. Theory* **57**(4), 2342–2359 (2011)
7. E.J. Candès, T. Tao, The Dantzig selector: statistical estimation when p is much larger than n . *Ann. Stat.* **35**(6), 2313–2351 (2007)
8. A. Carpentier, A. Kim, An iterative hard thresholding estimator for low rank matrix recovery with explicit limiting distribution (2015). arxiv preprint 1502.04654
9. A. Carpentier, O. Klopp, M. Löffler, R. Nickl, et al., Adaptive confidence sets for matrix completion. *Bernoulli* **24**(4A), 2429–2460 (2018)
10. M. Christandl, R. Renner, Reliable quantum state tomography. *Phys. Rev. Lett.* **109**, 120403 (2012)
11. R. De Wolf, A brief introduction to Fourier analysis on the Boolean cube. *Theory Comput. Libr. Grad. Surv.* **1**, 1–20 (2008)
12. S.T. Flammia, D. Gross, Y.-K. Liu, J. Eisert, Quantum tomography via compressed sensing: error bounds, sample complexity and efficient estimators. *New J. Phys.* **14**(9), 095022 (2012)
13. E. Giné, R. Nickl, Confidence bands in density estimation. *Ann. Stat.* **38**, 1122–1170 (2010)
14. E. Giné, R. Nickl, *Mathematical Foundations of Infinite-Dimensional Statistical Models* (Cambridge University Press, Cambridge, 2015)
15. D. Gross, Recovering low-rank matrices from few coefficients in any basis. *IEEE Trans. Inf. Theory* **57**(3), 1548–1566 (2011)
16. D. Gross, Y.-K. Liu, S.T. Flammia, S. Becker, J. Eisert, Quantum state tomography via compressed sensing. *Phys. Rev. Lett.* **105**(15), 150401 (2010)
17. M. Guta, T. Kypraios, I. Dryden, Rank-based model selection for multiple ions quantum tomography. *New J. Phys.* **14**, 105002 (2012)
18. H. Haeffner, W. Haensel, C.F. Roos, J. Benhelm, D.C. al kar, M. Chwalla, T. Koerber, U.D. Rapol, M. Riebe, P.O. Schmidt, C. Becher, O. Gühne, W. Dür, R. Blatt, Scalable multi-particle entanglement of trapped ions. *Nature* **438**, 643 (2005)
19. M. Hoffmann, R. Nickl, On adaptive inference and confidence bands. *Ann. Stat.* **39**, 2382–2409 (2011)
20. A.S. Holevo, *Statistical Structure of Quantum Theory* (Springer, Berlin, 2001)
21. Y.I. Ingster, A.B. Tsybakov, N. Verzelen, Detection boundary in sparse regression. *Electron. J. Stat.* **4**, 1476–1526 (2010)
22. A. Javanmard, A. Montanari, Confidence intervals and hypothesis testing for high-dimensional regression. *J. Mach. Learn. Res.* **15**(1), 2869–2909 (2014)
23. A. Kaley, R.L. Kosut, I.H. Deutsch, Informationally complete measurements from compressed sensing methodology. arXiv:1502.00536
24. V. Koltchinskii, Von Neumann entropy penalization and low-rank matrix estimation. *Ann. Stat.* **39**(6), 2936–2973 (2011)
25. V. Koltchinskii, K. Lounici, A.B. Tsybakov, Nuclear-norm penalization and optimal rates for noisy low-rank matrix completion. *Ann. Stat.* **39**(5), 2302–2329 (2011)
26. H. Leeb, B.M. Pötscher, Can one estimate the conditional distribution of post-model-selection estimators? *Ann. Stat.* **34**(5), 2554–2591 (2006)

27. U. Leonhardt, *Measuring the Quantum State of Light* (Cambridge University Press, Cambridge, 2005)
28. Y.-K. Liu, Universal low-rank matrix recovery from Pauli measurements. *Adv. Neural Inf. Process. Syst.* **24**, 1638–1646 (2011)
29. V. Mnih, C. Szepesvári, J.Y. Audibert, Empirical bernstein stopping, in *Proceedings of the 25th International Conference on Machine Learning* (ACM, New York, 2008), pp. 672–679
30. R. Nickl, S. van de Geer, Confidence sets in sparse regression. *Ann. Stat.* **41**(6), 2852–2876 (2013)
31. M.A. Nielsen, I.L. Chuang, *Quantum Computation and Quantum Information* (Cambridge University Press, Cambridge, 2000)
32. A. Peres, *Quantum Theory* (Springer, Berlin, 1995)
33. B. Recht, M. Fazel, P.A. Parrilo, Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM Rev.* **52**, 471 (2010)
34. J. Shang, H.K. Ng, A. Sehwat, X. Li, B.-G. Englert, Optimal error regions for quantum state estimation. *New J. Phys.* **15**(12), 123026 (2013)
35. K. Temme, F. Verstraete, Quantum chi-squared and goodness of fit testing. *J. Math. Phys.* **56**(1), 012202 (2015)
36. S. van de Geer, P. Bühlmann, Y. Ritov, R. Dezeure, On asymptotically optimal confidence regions and tests for high-dimensional models. *Ann. Stat.* **42**(3), 1166–1202 (2014)