

Dimension-free bounds for sums of independent matrices and simple tensors via the variational principle*

Nikita Zhivotovskiy[†]

Abstract

We consider the deviation inequalities for the sums of independent d by d random matrices, as well as rank one random tensors. Our focus is on the non-isotropic case and the bounds that do not depend explicitly on the dimension d , but rather on the effective rank. In an elementary and unified manner, we show the following results:

- A deviation bound for the sums of independent positive semi-definite matrices. This result complements the dimension-free bound of Koltchinskii and Lounici [Bernoulli, 23(1): 110–133, 2017] on the sample covariance matrix in the sub-Gaussian case.
- A new bound for truncated covariance matrices that is used to prove a dimension-free version of the bound of Adamczak, Litvak, Pajor and Tomczak-Jaegermann [Journal of Amer. Math. Soc., 23(2), 535–561, 2010] on the sample covariance matrix in the log-concave case.
- Dimension-free bounds for the operator norm of the sums of random tensors of rank one formed either by sub-Gaussian or by log-concave random vectors. This complements the result of Guédon and Rudelson [Adv. in Math., 208: 798–823, 2007].
- A non-isotropic version of the result of Alesker [Geom. Asp. of Funct. Anal., 77: 1–4, 1995] on the deviation of the norm of sub-exponential random vectors.
- A dimension-free lower tail bound for sums of positive semi-definite matrices with heavy-tailed entries, sharpening the bound of Oliveira [Prob. Th. and Rel. Fields, 166: 1175–1194, 2016].

Our approach is based on the duality formula between entropy and moment generating functions. In contrast to the known proofs of dimension-free bounds, we avoid Talagrand’s majorizing measure theorem, as well as generic chaining bounds for empirical processes. Some of our tools were pioneered by O. Catoni and co-authors in the context of robust statistical estimation.

Keywords: concentration inequalities; random matrices; sample covariance; effective rank; log-concave measures.

*This work was done when the author was at ETH Zürich, Department of Mathematics. He was funded in part by the ETH Foundations of Data Science (ETH-FDS).

[†]Department of Statistics, UC Berkeley. E-mail: zhivotovskiy@berkeley.edu

1 Introduction and main results

We study the non-asymptotic bounds for the sums of some independent random matrices as well as a closely related question of estimating the largest and smallest singular values of random matrices with independent rows. Assume that we are given a random n by d matrix A such that all of its rows $A_1^\top, \dots, A_n^\top$ are isotropic independent sub-Gaussian vectors (see the formal definitions below). We are interested in providing the upper and lower bounds on its singular values

$$s_1(A), s_2(A), \dots, s_d(A).$$

The question of upper bounding the largest singular values and lower bounding the smallest singular value is known to be essentially equivalent (see [50, Chapter 4]) to providing an upper bound on the operator norm of the difference between the sample covariance matrix formed by the rows A_i and the identity matrix. That is, one is interested in providing a high probability, non-asymptotic bound on

$$\left\| \frac{1}{n} \sum_{i=1}^n A_i A_i^\top - I_d \right\|. \quad (1.1)$$

Here and in what follows $\|\cdot\|$ stands for the operator norm of the matrix and for the Euclidean norm of the vector respectively. The latter question is also central in mathematical statistics, where one is interested in estimating the underlying covariance structure using the sample covariance matrix. One of the usual assumptions made when analyzing (1.1) is that the rows A_i^\top are isotropic and zero mean; that is, $\mathbf{E}A_i = 0$ and $\mathbf{E}A_i A_i^\top = I_d$, where in what follows I_d stands for the d by d identity matrix. The non-isotropic case can usually be reduced to the isotropic using a linear transformation. However, the problem is that in this case the bound on (1.1) will depend on the dimension d , whereas in many cases one expects that a dimension-free deviation bound is possible. The search for dimension-free bounds for sums of independent random matrices is motivated mainly by applications in statistics and data science, where it is usually assumed that the data lives on a low-dimensional manifold. Before providing our first result, we recall that for a random variable Y and $\alpha \in [1, 2]$, its ψ_α norm is defined as follows

$$\|Y\|_{\psi_\alpha} = \inf\{c > 0 : \mathbf{E} \exp(|Y|^\alpha / c^\alpha) \leq 2\}.$$

Using the standard convention, we say that $\|\cdot\|_{\psi_2}$ is the sub-Gaussian norm and $\|\cdot\|_{\psi_1}$ is the sub-exponential norm. We say that X is a sub-Gaussian random vector in \mathbb{R}^d if $\sup_{u \in S^{d-1}} \|\langle u, X \rangle\|_{\psi_2}$ is finite. A zero mean random vector X is isotropic, if $\mathbf{E}X X^\top = I_d$. Here and in what follows, S^{d-1} denotes the corresponding unit sphere and $\langle \cdot, \cdot \rangle$ is the standard inner product in \mathbb{R}^d . One of the central quantities appearing in this paper is the effective rank.

Definition 1.1. For a positive semi-definite matrix Σ define its effective rank as

$$\mathbf{r}(\Sigma) = \frac{\text{Tr}(\Sigma)}{\|\Sigma\|}.$$

The effective rank is always smaller than the matrix rank of Σ and, in particular, smaller than its dimensions. We also have $\mathbf{r}(I_d) = d$. Our first result is a general upper bound for sums of independent positive semi-definite d by d matrices satisfying the

sub-exponential norm equivalence assumption. This slightly generalizes the question of upper bounding (1.1), since we neither assume that the matrix is of rank one nor that the covariance matrix is identity.

Theorem 1.2 (A general rank version of Theorem 9 in [25]). *Assume that M_1, \dots, M_n are independent copies of a d by d positive semi-definite symmetric random matrix M with mean $\mathbf{E}M = \Sigma$. Let M satisfy for some $\kappa \geq 1$,*

$$\|x^\top Mx\|_{\psi_1} \leq \kappa^2 x^\top \Sigma x, \tag{1.2}$$

for all $x \in \mathbb{R}^d$. Then, for any $t > 0$, with probability at least $1 - \exp(-t)$, it holds that

$$\left\| \frac{1}{n} \sum_{i=1}^n M_i - \Sigma \right\| \leq 20\kappa^2 \|\Sigma\| \sqrt{\frac{4\mathbf{r}(\Sigma) + t}{n}}, \tag{1.3}$$

whenever $n \geq 4\mathbf{r}(\Sigma) + t$.

Remark 1.3. In the theorem above, we presented explicit constants. We place them to emphasize that, in contrast with existing dimension-free bounds, these constants are easy to obtain with the approach we follow. At the same time, little effort was made to get their optimal values.

Remark 1.4. In Section 3.2 we show that the same dimension-free bound holds under a weaker assumption (allowing heavy-tailed distributions), namely $\sqrt{\mathbf{E}(x^\top Mx)^2} \leq \kappa^2 x^\top \Sigma x$, but only for the lower tails of (1.3). This complements several known dimension-dependent lower tail bounds.

The norm equivalence assumption (1.2) is quite standard in the literature. As a matter of fact, Theorem 1.2 recovers one of the central results in high-dimensional statistics, as the following example shows.

Example 1.5 (The sample covariance matrix in the sub-Gaussian case [25]). The most natural application of Theorem 1.2 is when $M = XX^\top$ and X is zero mean *sub-Gaussian random vector* with a covariance matrix Σ . That is, there is $\kappa \geq 1$ such that for any $y \in \mathbb{R}^d$, it holds that

$$\|\langle y, X \rangle\|_{\psi_2} \leq \kappa \sqrt{y^\top \Sigma y}. \tag{1.4}$$

Using this line, for any $x \in \mathbb{R}^d$ we have $\|x^\top XX^\top x\|_{\psi_1} = \|x^\top X\|_{\psi_2}^2 \leq \kappa^2 x^\top \Sigma x$, which is sufficient for Theorem 1.2. This gives that, with probability at least $1 - \exp(-t)$, provided that $n \geq 4\mathbf{r}(\Sigma) + t$, it holds that

$$\left\| \frac{1}{n} \sum_{i=1}^n X_i X_i^\top - \Sigma \right\| = \sup_{y \in S^{d-1}} \left| \frac{1}{n} \sum_{i=1}^n \langle X_i, y \rangle^2 - \mathbf{E} \langle X, y \rangle^2 \right| \leq 20\kappa^2 \|\Sigma\| \sqrt{\frac{4\mathbf{r}(\Sigma) + t}{n}}, \tag{1.5}$$

where we used that for a symmetric d by d matrix A it holds that $\|A\| = \sup_{y \in S^{d-1}} y^\top A y$.

Recently, much attention has been paid to the dimension-free bound for the sample covariance matrix formed by sub-Gaussian random vectors. Although the dimension-dependent version of (1.5) (that is, where $\mathbf{r}(\Sigma)$ is replaced by d) follows from a simple discretization argument, the known approaches to obtaining the dimension-free bounds are quite technical and deserve a separate discussion:

- The bound of Theorem 1.2 is a general rank version¹ of the result of Koltchinskii and Lounici [25, Theorem 9]. Their proof is based on deep probabilistic results: the generic chaining tail bounds for quadratic processes and Talagrand’s majorizing

¹For the sake of technical simplicity, we work with finite dimensional vectors, whereas the results in [25] are formulated for a general Hilbert space.

measures theorem. In particular, this makes it difficult to provide any explicit constants using their approach. Some generalizations of the result on sample covariance matrices to positive semi-definite matrices are also considered in [53].

- Using the matrix deviation inequality of Liaw, Mehrabian, Plan, and Vershynin [30], Vershynin [50] gives an alternative proof of the bound of Koltchinskii and Lounici, but with the term κ^4 instead of κ^2 in (1.5). The dependence on κ in [30] has been recently improved in [21]. However, this improved (and optimal for some problems) result only leads to $\kappa^2 \log \kappa$ term in (1.5).
- Van Handel [48] gives an in-expectation version of (1.5) in the special case where X is a Gaussian random vector. Despite not using Talagrand’s majorizing measures theorem, their analysis is based on a Gaussian comparison theorem and does not cover the sub-Gaussian case. Note that in the Gaussian case, the in-expectation bound for the sample covariance matrix can be converted into an optimal high probability bound using one of the special concentration inequalities provided in [25, 1, 23].

Although the result of Theorem 1.2 can be proved using the machinery of generic chaining, our approach, based on the variational inequality and described in detail in Section 2, bypasses several technical steps appearing in the literature. Speaking informally, we use a smoothed version of the ε -net argument that allows us to properly capture the complexity of elliptic indexing sets without resorting to generic chaining. This extension will be key to our multilinear results, where the above mentioned tools are hard to apply.

Note that even though Example 1.5 is sharp in the rank one case (see the lower bound in [25]), the result of Theorem 1.2, due to its generality, can be suboptimal in other cases. For example, let A be a diagonal random matrix such that its diagonal elements are the same copy of the absolute value of a standard Gaussian random variable. In this case, Theorem 1.2 scales as $\sqrt{\frac{d+t}{n}}$, whereas the correct order is $\sqrt{\frac{t}{n}}$. We also remark that, at least in the rank one case, the bound of Theorem 1.2 is out of the scope of the so-called matrix concentration inequalities, since they provide additional logarithmic factors and suboptimal tails (see some related bounds in [47, 50, 23, 31]). We additionally refer to the recent work [7], where an in-expectation analog of (1.5) is shown in the Gaussian case, but with additional lower-order additive terms that are dimension dependent.

Motivated by the recent interest in random tensors [51, 18, 43, 9, 15], we show how our arguments can be extended to provide a multilinear extension of Theorem 1.2. That is, we are considering sums of independent random tensors of order higher than one and want to prove a bound similar to (1.5). Let us introduce this setup. Consider the simple (rank one) random symmetric tensor

$$X^{\otimes s} = \underbrace{X \otimes \dots \otimes X}_{s \text{ times}}$$

where X is a zero mean sub-Gaussian vector in \mathbb{R}^d and $X_1^{\otimes s}, \dots, X_n^{\otimes s}$ are its independent copies. We are interested in studying

$$\left\| \sum_{i=1}^n (X_i^{\otimes s} - \mathbf{E}X_i^{\otimes s}) \right\| = \sup_{v \in S^{d-1}} \left| \sum_{i=1}^n (\langle X_i, v \rangle^s - \mathbf{E}\langle X, v \rangle^s) \right|, \tag{1.6}$$

where $\|\cdot\|$ stands for the operator norm of the symmetric s -linear form. Here we used that for symmetric forms the expression is maximized by a single vector $v \in S^{d-1}$ (see e.g., [39, Section 2.3]).

The question of upper bounding (1.6) (with $\langle X_i, v \rangle^s$ usually replaced by the absolute value $|\langle X_i, v \rangle|^s$ in the right-hand side and non-integer values of s are allowed) is well studied [16, 19, 34, 3, 49, 35]. The results are usually of the following form: Assuming that X_1, \dots, X_n are i.i.d. copies of the isotropic vector satisfying certain regularity assumptions, one is interested in defining the smallest sample size n such that, with high probability,

$$\sup_{v \in S^{d-1}} \left| \frac{1}{n} \sum_{i=1}^n |\langle X_i, v \rangle|^s - \mathbf{E}|\langle X, v \rangle|^s \right| \leq \varepsilon.$$

The general form of the assumption (see [19, 49] and, in particular, [3, Theorem 4.2]) required to achieve this precision for some regular families of distributions is

$$n \geq C_s d^{s/2} \varepsilon^{-2} \log^{2+s}(2\varepsilon^{-2}), \tag{1.7}$$

where C_s depends only on s . Although the condition $n \geq C_s d^{s/2}$ is known to be optimal when ε is a constant [49], the dependence on ε is either suboptimal or not explicit in the existing results. In fact, a recent result of Mendelson [35] suggests that using a specific robust estimation procedure, it is possible to approximate the moments of the marginals for any $s \geq 1$ with ε being proportional to $\sqrt{\frac{d}{n} \log \frac{n}{d}}$. At the same time, the inequality (1.7) becomes vacuous in this regime whenever $s > 2$. Before we proceed, recall the following definition.

Definition 1.6. *The measure ν in \mathbb{R}^d is log-concave, if for any measurable subsets $A, B \subset \mathbb{R}^d$ and any $t \in [0, 1]$,*

$$\nu(tA + (1-t)B) \geq \nu(A)^t \nu(B)^{1-t},$$

whenever the set $tA + (1-t)B = \{ta + (1-t)b : a \in A, b \in B\}$ is measurable.

Our next result shows that provided that the sample size n is large enough, one can approximate the s -th integer moment of the marginals using their empirical counterparts with ε being proportional to $\|\Sigma\|^{s/2} \sqrt{\frac{\mathbf{r}(\Sigma)}{n}}$. This is the best possible approximation rate when $s = 2$ and X is a multivariate Gaussian random vector. Moreover, because of (1.7) this approximation rate was not previously achieved even in the isotropic case.

Theorem 1.7. *Let $s \geq 2$ be an integer. Assume that X_1, \dots, X_n are independent copies of a zero mean vector X that is either sub-Gaussian (1.4) or log-concave. There exist $c_s > 0$ that depends only on s and an absolute constant $c > 0$ such that the following holds. Assume that*

$$n \geq c_s (\mathbf{r}(\Sigma))^{s-1}. \tag{1.8}$$

Then, with probability at least $1 - cn \exp(-\sqrt{\mathbf{r}(\Sigma)})$, it holds that

$$\sup_{v \in S^{d-1}} \left| \frac{1}{n} \sum_{i=1}^n \langle X_i, v \rangle^s - \mathbf{E}\langle X, v \rangle^s \right| \leq C \|\Sigma\|^{s/2} \sqrt{\frac{\mathbf{r}(\Sigma)}{n}},$$

where $C = C(s, \kappa)$ in the sub-Gaussian case and $C = C(s)$ in the log-concave case.

Moreover, in the sub-Gaussian case if $n \geq c_s (\mathbf{r}(\Sigma))^{s+1}$, then the same bound holds with probability at least $1 - cn \exp(-\mathbf{r}(\Sigma))$.

To simplify our proofs, we focus only on the tensor case; in particular, we consider the integer values of s . In Theorem 1.7 we require that either $cn \exp(-\mathbf{r}(\Sigma)) < 1$ or $cn \exp(-\sqrt{\mathbf{r}(\Sigma)}) < 1$. These assumptions can be dropped by slightly inflating our upper bound (see also [3, Remark 4.3]). It is likely that in the sub-Gaussian case one can extend our arguments, namely, the decoupling-chaining argument discussed below, so that the assertion holds with probability $1 - cn \exp(-\mathbf{r}(\Sigma))$ whenever $n \geq c_s (\mathbf{r}(\Sigma))^{s-1}$. Indeed, we

know by (1.5) that this is the case at least when $s = 2$. We preferred a shorter proof instead of a more accurate estimate of the tail.

In the log-concave case when $s = 2$, Theorem 1.7 complements the renowned result of Adamczak, Litvak, Pajor and Tomczak-Jaegermann [3]. The main advantage of our result is the explicit dependence on the effective rank, similar to the sample covariance bound of Koltchinskii and Lounici in the Gaussian case [25]. Our next result sharpens the tail estimate in this specific case and coincides with the best known bound in the isotropic case.

Theorem 1.8. *Assume that X_1, \dots, X_n are independent copies of a zero-mean, log-concave random vector X with covariance Σ . There are absolute constants $c_1, c_2, c_3 > 0$ such that the following holds. We have, with probability at least $1 - c_1 \exp(-(\mathbf{r}(\Sigma)n)^{1/4}) - 2 \exp(-\mathbf{r}(\Sigma))$,*

$$\left\| \frac{1}{n} \sum_{i=1}^n X_i X_i^\top - \Sigma \right\| \leq c_2 \|\Sigma\| \sqrt{\frac{\mathbf{r}(\Sigma)}{n}},$$

whenever $n \geq c_3 \mathbf{r}(\Sigma)$.

In both proofs we combine the variational inequality approach with the decoupling-chaining argument developed in [3, 45]. We remark that the adaptation of the latter argument to the non-isotropic case is quite straightforward. Our analysis improves the bound on the *spread part* [45, Section 9.4] whose analysis in the isotropic case combines the ε -net argument and the Bernstein inequality. Finally, observe that since the Gaussian distribution is log-concave, the rate of convergence $\|\Sigma\| \sqrt{\frac{\mathbf{r}(\Sigma)}{n}}$ is the best possible due to the lower bound in [25].

There is a version of Theorem 1.8, which follows from our proof with minimal changes. For the reader's convenience we also present an explicit tail bound.

Proposition 1.9. *Assume that X_1, \dots, X_n are independent copies of a zero-mean random vector with covariance Σ such that for some $\kappa \geq 1$ and all $y \in \mathbb{R}^d$, it holds that*

$$\|\langle X, y \rangle\|_{\psi_1} \leq \kappa \sqrt{y^\top \Sigma y} \quad \text{and} \quad \max_i \|X_i\| \leq \kappa \sqrt{\|\Sigma\| (\mathbf{r}(\Sigma)n)^{\frac{1}{4}}} \quad \text{almost surely.}$$

There are absolute constants $c_1, c_2, c_3 > 0$ such that the following holds. For any $t \geq 0$ we have, with probability at least $1 - c_1 \exp(-t)$,

$$\left\| \frac{1}{n} \sum_{i=1}^n X_i X_i^\top - \Sigma \right\| \leq c_2 \kappa^2 \|\Sigma\| \left(\sqrt{\frac{\mathbf{r}(\Sigma) + t}{n}} + \frac{t^2}{n} \right),$$

whenever $n \geq c_3 (\mathbf{r}(\Sigma) + t)$.

In Section 3 we provide two additional results: a bound on the deviation of the norm of a sub-exponential random vector and a lower tail version of Theorem 1.2. Finally, as a part of the proof of Theorem 1.7, we provide a simple proof of the bound by Hsu, Kakade and Zhang [20] on the deviation of the norm of a sub-Gaussian random vector.

2 An approach based on the variational equality

Our approach will be based on the following duality relation (see [8, Corollary 4.14]): for a probability space (Θ, μ) and any measurable function g such that $\mathbf{E}_\mu \exp(g(\theta)) < \infty$, it holds that

$$\log \mathbf{E}_\mu \exp(g(\theta)) = \sup_{\rho \ll \mu} (\mathbf{E}_\rho g(\theta) - \mathcal{KL}(\rho, \mu)), \tag{2.1}$$

where the supremum is taken with respect to all measures absolutely continuous with respect to μ and

$$\mathcal{KL}(\rho, \mu) = \int \log \left(\frac{d\rho}{d\mu} \right) d\rho$$

denotes the Kullback-Leibler divergence between ρ and μ . The equality (2.1) is used in the proof of the additivity of entropy [29, Proposition 5.6] and in the transportation method for proving concentration inequalities [8, Chapter 8]. A useful corollary of the variational equality is the following lemma (see e.g., [12, Proposition 2.1] and discussions therein).

Lemma 2.1. *Assume that X_i are i.i.d. random variables defined on some measurable space \mathcal{X} . Assume also that Θ (called the parameter space) is a subset of \mathbb{R}^p . Let $f : \mathcal{X} \times \Theta \rightarrow \mathbb{R}$ be such that $\mathbf{E}_X \exp(f(X, \theta)) < \infty$ almost surely. Let μ be a distribution (called prior) on Θ and let ρ be any distribution (called posterior) on Θ such that $\rho \ll \mu$. Then, with probability at least $1 - \exp(-t)$, simultaneously for all such ρ we have*

$$\frac{1}{n} \sum_{i=1}^n \mathbf{E}_\rho f(X_i, \theta) \leq \mathbf{E}_\rho \log(\mathbf{E}_X \exp(f(X, \theta))) + \frac{\mathcal{KL}(\rho, \mu) + t}{n},$$

where θ is distributed according to ρ .

Proof. We sketch the proof for the sake of completeness. Let $g : \mathcal{X}^n \times \Theta \rightarrow \mathbb{R}$ in (2.1) be equal to $\sum_{i=1}^n f(X_i, \theta) - n \log \mathbf{E}_X \exp(f(X, \theta))$. Let \mathbf{E} denote the expectation with respect to the i.i.d. sample X_1, \dots, X_n . Using successively (2.1), Fubini's theorem and independence of X_1, \dots, X_n , we have

$$\begin{aligned} & \mathbf{E} \exp \sup_{\rho \ll \mu} \left(\mathbf{E}_\rho \sum_{i=1}^n f(X_i, \theta) - \mathbf{E}_\rho n \log \mathbf{E}_X \exp(f(X, \theta)) - \mathcal{KL}(\rho, \mu) \right) \\ &= \mathbf{E} \mathbf{E}_\mu \exp \left(\sum_{i=1}^n f(X_i, \theta) - n \log \mathbf{E}_X \exp(f(X, \theta)) \right) \\ &= \mathbf{E}_\mu (\mathbf{E}_X \exp(f(X, \theta)) - \log \mathbf{E}_X \exp(f(X, \theta)))^n = 1. \end{aligned}$$

By Markov's inequality for any random variable Y the identity $\mathbf{E} \exp(Y) = 1$ implies that $Y < t$, with probability at least $1 - \exp(-t)$. The claim follows by taking

$$Y = \sup_{\rho \ll \mu} \left(\mathbf{E}_\rho \sum_{i=1}^n f(X_i, \theta) - \mathbf{E}_\rho n \log \mathbf{E}_X \exp(f(X, \theta)) - \mathcal{KL}(\rho, \mu) \right). \quad \square$$

Remark 2.2. In Lemma 2.1 we assumed that $\mathbf{E}_X \exp(f(X, \theta)) < \infty$ for all $\theta \in \Theta$. However, this does not imply that $f(X, \theta)$ is integrable with respect to ρ . If it is not the case, one can conventionally take the cases where $\mathbf{E}_\rho f(X, \theta)$ is infinite into account, so that the inequality of Lemma 2.1 still holds. For more details see [11, Appendix A].

Our analysis is inspired by the application of (2.1) and Lemma 2.1 in the works of Catoni and co-authors [6, 5, 10, 11, 12] on robust mean and covariance estimation as well as by the work of Oliveira [40] on the lower tails of sample covariance matrices under minimal assumptions. This approach is usually called the PAC-Bayesian method in the literature. In robust mean estimation, one is making minimal distributional assumptions (for example, by considering heavy-tailed distributions) aiming to estimate the mean of the random variable/vector/matrix using the estimators that necessarily differ from the sample mean (see [10, 37, 33, 12, 36, 14, 41, 22, 35] and the recent survey [32]). Our aim is somewhat different: We work with sums of independent random matrices and multilinear forms. It is important to note that except the recent works of Catoni and Giulini [17, 12], statistical guarantees based on (2.1) are dimension-dependent. A detailed technical comparison with these papers is deferred to Section 4.

2.1 Motivating examples: matrices with isotropic sub-Gaussian rows and the Gaussian complexity of ellipsoids

To motivate (and illustrate) the application of the variational equality (2.1) in the context of high-dimensional probability, we first show how Lemma 2.1 can be used to recover the standard bound on the largest and smallest singular values of the the n by d random matrix A having independent, mean zero, isotropic ($\mathbf{E}A_iA_i^\top = I_d$) sub-Gaussian rows. In this case, (1.4) can be rewritten as

$$\sup_{y \in S^{d-1}} \|\langle y, A_i \rangle\|_{\psi_2} \leq \kappa.$$

In view of [50, Lemma 4.1.5], it is enough to show the following statement.

Proposition 2.3 ([50, Theorem 4.6.1]). *Let A be an n by d random matrix whose rows A_i^\top are independent, mean zero, sub-gaussian isotropic random vectors. We have for any $t \geq 0$, with probability at least $1 - \exp(-t)$,*

$$\left\| \frac{1}{n} \sum_{i=1}^n A_i A_i^\top - I_d \right\| \leq 52\kappa^2 \sqrt{\frac{d+t}{n}},$$

whenever $n \geq d + t$.

The standard way of proving Proposition 2.3 uses an ε -net argument combined with the Bernstein inequality in terms of the $\|\cdot\|_{\psi_1}$ norm and the union bound. We demonstrate that if the prior μ and the posterior ρ are correctly chosen, then Lemma 2.1 recovers the same bound without directly exploiting a discretization argument.

Remark 2.4. The condition $n \geq d + t$ does not appear in [50, Theorem 4.6.1]. As a result, this bound contains an additional additive term scaling as $\frac{d+t}{n}$. In our regime when $n \geq d + t$, this term is naturally dominated by $\sqrt{\frac{d+t}{n}}$. In some sense, we only captured the sub-Gaussian regime in the deviation bound. This regime is arguably the most interesting when considering statistical estimation problems.

Our analysis requires the following standard result. Since we need a version with an explicit constant, we reproduce these lines for the sake of completeness.

Lemma 2.5. *Let Y be a zero mean random variable. Then for any λ such that $|\lambda| \leq \frac{1}{2\|Y\|_{\psi_1}}$,*

$$\mathbf{E} \exp(\lambda Y) \leq \exp(4\lambda^2 \|Y\|_{\psi_1}^2).$$

Proof. First, by Markov’s inequality and any $t \geq 0$, it holds that

$$\Pr(|Y| \geq t) = \Pr(\exp(|Y|/\|Y\|_{\psi_1}) \geq \exp(t/\|Y\|_{\psi_1})) \leq 2 \exp(-t/\|Y\|_{\psi_1}).$$

In the following lines, we assume without loss of generality that $\|Y\|_{\psi_1} = 1$. We have for $p \geq 1$,

$$\mathbf{E}|Y|^p = \int_{t=0}^{\infty} \Pr(|Y|^p \geq t) dt = \int_{t=0}^{\infty} \Pr(|Y| \geq t^{1/p}) p t^{p-1} dt \leq 2 \int_{t=0}^{\infty} \exp(-t) p t^{p-1} dt = 2p! . \quad (2.2)$$

Finally, when $|\lambda| \leq 1/2$ by Taylor’s expansion and since $\mathbf{E}Y = 0$, we have

$$\mathbf{E} \exp(\lambda Y) = 1 + \sum_{p=2}^{\infty} \frac{\lambda^p \mathbf{E}|Y|^p}{p!} \leq 1 + 2\mathbf{E} \sum_{p=2}^{\infty} \lambda^p = 1 + \frac{2\lambda^2}{1-\lambda} \leq 1 + 4\lambda^2 \leq \exp(4\lambda^2).$$

The claim follows. □

Proof of Proposition 2.3. Fix $\varepsilon > 0$. Our aim is to choose μ and ρ . Let

$$\Theta = (1 + \varepsilon)B^d \times (1 + \varepsilon)B^d,$$

where B^d is a unit ball in \mathbb{R}^d . Choose μ to be a product of two uniform measures each defined on $(1 + \varepsilon)B^d$. For $u, v \in S^{d-1}$ let $\rho_{u,v}$ be a product of uniform distributions on the balls $\{x \in \mathbb{R}^d : \|x - u\| \leq \varepsilon\}$ and $\{x \in \mathbb{R}^d : \|x - v\| \leq \varepsilon\}$. Observe that both balls belong to $(1 + \varepsilon)B^d$. Because of this, if (θ, ν) is distributed according to $\rho_{u,v}$, we have $\mathbf{E}_{\rho_{u,v}}(\theta, \nu) = (u, v)$. By the additivity of \mathcal{KL} -divergence for product measures and the formula for the volume of the d -dimensional ball, we have

$$\mathcal{KL}(\rho_{u,v}, \mu) = 2 \int_{\|x-u\| \leq \varepsilon} \log \left(\frac{(1 + \varepsilon)^d}{\varepsilon^d} \right) \frac{1}{\text{vol}(\{x \in \mathbb{R}^d : \|x - v\| \leq \varepsilon\})} dx = 2d \log \left(\frac{1 + \varepsilon}{\varepsilon} \right),$$

where $\text{vol}(S)$ denotes the volume of the set S . Fix $\lambda \in \mathbb{R}$ and consider the random variable $\lambda \theta^\top A_i A_i^\top \nu$, where (θ, ν) is distributed according to $\rho_{u,v}$. We want to plug this random variable into Lemma 2.1. Observe that conditionally on (θ, ν) , we have, using $\|YZ\|_{\psi_1} \leq \|Y\|_{\psi_2} \|Z\|_{\psi_2}$,

$$\|\theta^\top A_i A_i^\top \nu\|_{\psi_1} \leq \|\theta\| \|\nu\| \sup_{y \in S^{d-1}} \|\langle y, A_i \rangle\|_{\psi_2}^2 \leq (1 + \varepsilon)^2 \kappa^2,$$

where the last inequality follows from the fact that $\theta, \nu \in (1 + \varepsilon)B^d$ almost surely. Conditionally on (θ, ν) , combining the triangle and Jensen's inequalities, we have

$$\|\theta^\top A_i A_i^\top \nu - \theta^\top \mathbf{E} A_i A_i^\top \nu\|_{\psi_1} \leq 2 \|\theta^\top A_i A_i^\top \nu\|_{\psi_1}.$$

Further, since $\mathbf{E} A_i A_i^\top = I_d$ we have by Lemma 2.5, conditionally on (θ, ν) ,

$$\mathbf{E} \exp(\lambda \theta^\top A_i A_i^\top \nu) \leq \exp(\lambda \theta^\top \nu + 16\lambda^2(1 + \varepsilon)^4 \kappa^4),$$

whenever $\lambda \leq \frac{1}{4(1 + \varepsilon)^2 \kappa^2}$. Therefore, since $\mathbf{E}_{\rho_{u,v}} \theta^\top A_i A_i^\top \nu = u^\top A_i A_i^\top v$ and $\mathbf{E}_{\rho_{u,v}} \theta^\top \nu = u^\top v$ Lemma 2.1 gives that simultaneously for all $u, v \in S^{d-1}$, with probability at least $1 - \exp(-t)$,

$$\frac{1}{n} \sum_{i=1}^n \lambda u^\top A_i A_i^\top v - \lambda u^\top v \leq 16\lambda^2 \kappa^4 (1 + \varepsilon)^4 + \frac{2d \log((1 + \varepsilon)/\varepsilon) + t}{n}. \tag{2.3}$$

We choose $\varepsilon = \frac{1}{\sqrt{e-1}}$ to guarantee that $2d \log((1 + \varepsilon)/\varepsilon) = d$. Then, taking

$$\lambda = \frac{1}{4(1 + \varepsilon)^2 \kappa^2} \sqrt{\frac{d + t}{n}},$$

we require $n \geq d + t$. Simplifying (2.3) for this choice of parameters, we prove the claim. \square

Another motivating fact is that (2.1) correctly reflects the Gaussian complexity of the ellipsoid. It is well-known that for the ellipsoids, the Dudley integral argument does not give an optimal bound, while the generic chaining does (see [45, Chapter 2.5]). Although one can instead directly use the Cauchy-Schwarz inequality, it is easy to show that the variational equality (2.1) captures the same bound.

Example 2.6 (The Gaussian complexity of ellipsoids via the variational equality). Let Z be a standard normal random vector. Let Σ be a positive semi-definite d by d matrix. It hold that

$$\mathbf{E} \sup_{v \in \Sigma^{1/2} S^{d-1}} \langle Z, v \rangle \leq \sqrt{\text{Tr}(\Sigma)}.$$

Proof. Set $\Theta = \mathbb{R}^d$ and let $\beta > 0$. Let the prior distribution μ be a multivariate Gaussian distribution with mean zero and covariance $\beta^{-1}\Sigma$. For $v \in \Sigma^{1/2}S^{d-1}$ let the distribution ρ_v be a multivariate Gaussian distribution with mean v and covariance $\beta^{-1}\Sigma$. By the standard formula, we have

$$\mathcal{KL}(\rho_v, \mu) = \beta/2.$$

Let θ be distributed according to ρ_v . By Jensen's inequality for any $\lambda > 0$,

$$\mathbf{E} \sup_{v \in \Sigma^{1/2}S^{d-1}} \langle Z, v \rangle = \mathbf{E} \sup_{v \in \Sigma^{1/2}S^{d-1}} \mathbf{E}_{\rho_v} \langle Z, \theta \rangle \leq \frac{1}{\lambda} \log \mathbf{E} \sup_{v \in \Sigma^{1/2}S^{d-1}} \exp(\lambda \mathbf{E}_{\rho_v} \langle Z, \theta \rangle).$$

By the line of the proof of Lemma 2.1, we have

$$\mathbf{E} \sup_{v \in \Sigma^{1/2}S^{d-1}} \exp(\mathbf{E}_{\rho_v} \lambda \langle Z, \theta \rangle - \mathbf{E}_{\rho_v} \log \mathbf{E} \exp(\lambda \langle Z, \theta \rangle)) - \beta/2 \leq 1.$$

Since Z is a standard normal random vector, it holds that

$$\mathbf{E}_{\rho_v} \log \mathbf{E} \exp(\lambda \langle Z, \theta \rangle) = \mathbf{E}_{\rho_v} \lambda^2 \|\theta\|^2 / 2 = \lambda^2 (\|v\|^2 + \beta^{-1} \text{Tr}(\Sigma)) / 2.$$

Combining previous inequalities and simplifying, we have

$$\mathbf{E} \sup_{v \in \Sigma^{1/2}S^{d-1}} \langle Z, v \rangle \leq \inf_{\lambda, \beta > 0} \lambda (\|\Sigma\| + \beta^{-1} \text{Tr}(\Sigma)) / 2 + \beta / (2\lambda) = \sqrt{\text{Tr}(\Sigma)}.$$

The claim follows. □

Observe that the proof explicitly uses a bound on the expected squared norm of a multivariate normal vector (not for the norm of Z though), which is closer to a more “algebraic” approach based on the Cauchy-Schwarz inequality, whereas the generic chaining is a “geometric” approach; we refer to [45, Chapter 2.5] for a detailed discussion of the Gaussian complexity of ellipsoids.

2.2 Proof of Theorem 1.2

In view of Proposition 2.3 and Example 2.6, a natural idea is to use the uniform distribution for ρ and μ on ellipsoids induced by the structure of the matrix Σ . It appears that working with ellipsoids directly is quite tedious. To avoid these technical problems, we work with the non-isotropic truncated Gaussian distribution. Throughout the proof, we assume without loss of generality that Σ is invertible. If it is not the case, the distribution of M lives almost surely in a lower-dimensional subspace. We can project on this subspace and continue the proof without changes. Fix $\beta > 0$. Let

$$\Theta = \mathbb{R}^d \times \mathbb{R}^d,$$

and choose the prior distribution μ on Θ as the product of two multivariate Gaussian distributions in \mathbb{R}^d both with mean zero and covariance matrix $\beta^{-1}\Sigma$. For $u, v \in \Sigma^{1/2}S^{d-1}$ let the posterior distribution $\rho_{u,v}$ be defined as follows. For $r > 0$ consider the density function f_u in \mathbb{R}^d given by

$$f_u(x) = \frac{1}{p(2\pi)^{d/2} \sqrt{\det(\beta^{-1}\Sigma)}} \exp\left(-\frac{(x-u)^\top \beta \Sigma^{-1} (x-u)}{2}\right) \mathbf{1}(\|x-u\| \leq r), \quad (2.4)$$

where $p > 0$ is a normalization constant. That is, the distribution defined by f_u is a multivariate normal distribution restricted to the ball $\{x \in \mathbb{R}^d : \|x-u\| \leq r\}$. Our distribution $\rho_{u,v}$ on Θ is now defined as a product of two distributions given by f_u and f_v

respectively. Observe that since f_u is symmetric around u (that is, for any $y \in \mathbb{R}^d$, we have $f_u(u + y) = f_u(u - y)$), we have for (θ, ν) distributed according to $\rho_{u,v}$,

$$\mathbf{E}_{\rho_{u,v}}(\theta, \nu) = (\mathbf{E}_{\rho_u}\theta, \mathbf{E}_{\rho_v}\nu) = (u, v), \tag{2.5}$$

where ρ_u and ρ_v denote the marginals of $\rho_{u,v}$.

Let us now compute the Kullback-Leibler divergence between $\rho_{u,v}$ and μ . Let g denote the density function of a multivariate Gaussian distribution with mean zero and covariance $\beta^{-1}\Sigma$. By the additivity of the Kullback-Leibler divergence for product measures, we have

$$\mathcal{KL}(\rho_{u,v}, \mu) = \int \log\left(\frac{f_u(x)}{g(x)}\right) f_u(x) dx + \int \log\left(\frac{f_v(x)}{g(x)}\right) f_v(x) dx.$$

Both terms are now analyzed similarly. For θ distributed according to ρ_u ,

$$\begin{aligned} \int \log\left(\frac{f_u(x)}{g(x)}\right) f_u(x) dx &= \mathbf{E}_{\rho_u} \log\left(\frac{1}{p} \exp\left(\frac{-(\theta - u)^\top \beta \Sigma^{-1} (\theta - u) + \theta^\top \beta \Sigma^{-1} \theta}{2}\right)\right) \\ &= \log\left(\frac{1}{p}\right) + \mathbf{E}_{\rho_u} \left(\frac{u^\top \beta \Sigma^{-1} \theta + \theta^\top \beta \Sigma^{-1} u - u^\top \beta \Sigma^{-1} u}{2}\right) \\ &= \log\left(\frac{1}{p}\right) + \frac{u^\top \beta \Sigma^{-1} u}{2} = \log\left(\frac{1}{p}\right) + \frac{\beta}{2}, \end{aligned}$$

where in the last line we used $u \in \Sigma^{1/2} S^{d-1}$. Let Z be a random vector having a multivariate Gaussian distribution with mean zero and covariance $\beta^{-1}\Sigma$. By (2.4) and using the translation $u \rightarrow 0$, we have $p = \Pr(\|Z\| \leq r)$. By Markov's inequality we have $\Pr(\|Z\| > r) \leq \mathbf{E}\|Z\|^2/r^2 = \beta^{-1} \text{Tr}(\Sigma)/r^2$. We choose

$$r = \sqrt{2\beta^{-1} \text{Tr}(\Sigma)}$$

and get $p \geq 1/2$. Therefore, we have $\log(1/p) \leq \log 2$. Finally, for this choice of r ,

$$\mathcal{KL}(\rho_{u,v}, \mu) \leq 2 \log 2 + \beta, \tag{2.6}$$

For $\lambda \in \mathbb{R}$ we want to plug the function $\lambda \theta^\top \Sigma^{-1/2} M \Sigma^{-1/2} \nu$ into Lemma 2.1, where (θ, ν) is distributed according to $\rho_{u,v}$. By (2.5) we have

$$\frac{1}{n} \sum_{i=1}^n \mathbf{E}_{\rho_{u,v}} \lambda \theta^\top \Sigma^{-1/2} M_i \Sigma^{-1/2} \nu = \frac{1}{n} \sum_{i=1}^n \lambda u^\top \Sigma^{-1/2} M_i \Sigma^{-1/2} v. \tag{2.7}$$

It is only left to compute $\mathbf{E}_{\rho_{u,v}} \log(\mathbf{E} \exp(\lambda \theta^\top \Sigma^{-1/2} M \Sigma^{-1/2} \nu))$. Conditionally on (θ, ν) , we have as in the proof of Proposition 2.3

$$\begin{aligned} &\|\theta^\top \Sigma^{-1/2} M \Sigma^{-1/2} \nu - \theta^\top \Sigma^{-1/2} \mathbf{E} M \Sigma^{-1/2} \nu\|_{\psi_1} \\ &\leq 2 \|\theta^\top \Sigma^{-1/2} M \Sigma^{-1/2} \nu\|_{\psi_1} \\ &\leq \|\theta^\top \Sigma^{-1/2} M \Sigma^{-1/2} \theta\|_{\psi_1} + \|\nu^\top \Sigma^{-1/2} M \Sigma^{-1/2} \nu\|_{\psi_1} \\ &\leq \kappa^2 (\|\theta\|^2 + \|\nu\|^2), \end{aligned}$$

where both the $\|\cdot\|_{\psi_1}$ norm and the expectation are considered with respect to the distribution of M , and the second line uses the Cauchy-Schwarz inequality. Taking again the expectation with respect to M only, we have by Lemma 2.5

$$\mathbf{E} \exp(\lambda \theta^\top \Sigma^{-1/2} M \Sigma^{-1/2} \nu) \leq \exp(\lambda \theta^\top \Sigma^{-1/2} \mathbf{E} M \Sigma^{-1/2} \nu) \exp(4\lambda^2 \kappa^4 (\|\theta\|^2 + \|\nu\|^2)^2),$$

provided that $|\lambda| \leq \frac{1}{2\kappa^2(\|\theta\|^2 + \|\nu\|^2)}$. Observe that by our choice of r , we have almost surely

$$\max\{\|\theta\|^2, \|\nu\|^2\} \leq (\sqrt{\|\Sigma\|} + r)^2 = (\sqrt{\|\Sigma\|} + \sqrt{2\beta^{-1} \text{Tr}(\Sigma)})^2.$$

Let us choose $\beta = 2\mathbf{r}(\Sigma)$. Thus, by (2.6), (2.7) and Lemma 2.1 we have for any fixed λ such that $|\lambda| \leq \frac{1}{16\kappa^2\|\Sigma\|}$ simultaneously for all $u, v \in S^{d-1}$,

$$\frac{1}{n} \sum_{i=1}^n \lambda u^T M v \leq \lambda u^T \Sigma v + 64\lambda^2 \kappa^4 \|\Sigma\|^2 + \frac{4\mathbf{r}(\Sigma) + t}{n},$$

where we used $2 \log 2 + 2\mathbf{r}(\Sigma) \leq 4\mathbf{r}(\Sigma)$. We choose $\lambda = \frac{1}{16\kappa^2\|\Sigma\|} \sqrt{\frac{4\mathbf{r}(\Sigma) + t}{n}}$ and finish the proof. \square

2.3 Proofs of Theorem 1.7 and Theorem 1.8

We first present some auxiliary results, then we prove Theorem 1.8 and Theorem 1.7. The technique of the proof combines the analysis of Theorem 1.2 with a careful truncation argument. We also use the decoupling-chaining argument to control the large components in the sums. In the last part, we are mainly adapting the previously known techniques.

We need the following result, which is similar to the deviation inequality appearing in [20]. We provide a simple proof based on the variational equality (2.1).

Lemma 2.7. *Assume that X is a zero mean κ -sub-Gaussian random vector (1.4). Then, with probability at least $1 - \exp(-t)$,*

$$\|X\|^2 \leq 36\kappa^2 \left(\text{Tr}(\Sigma)/2 + \sqrt{2t \text{Tr}(\Sigma)\|\Sigma\|} + t\|\Sigma\| \right). \tag{2.8}$$

Remark 2.8. We will be frequently using the following relaxation of the bounds (2.8):

$$\|X\|^2 \leq 36\kappa^2 (\text{Tr}(\Sigma) + 2t\|\Sigma\|).$$

Proof. Observe that $\|X\| = \sup_{v \in S^{d-1}} \langle X, v \rangle$. Thus, we upper bound $\langle X, v \rangle$ uniformly over the sphere. Set $\Theta = \mathbb{R}^d$. Let the prior distribution μ be a multivariate Gaussian distribution with mean zero and covariance $\beta^{-1}I_d$. For $v \in S^{d-1}$ let ρ_v be a multivariate Gaussian distribution with mean v and covariance $\beta^{-1}I_d$. By the standard formula, we have

$$\mathcal{KL}(\rho_v, \mu) = \beta/2.$$

Our function is $\lambda \langle X, \theta \rangle$, where θ is distributed according to ρ_v . To apply Lemma 2.1 (with $n = 1$) we only need to compute $\mathbf{E}_{\rho_v} \log(\mathbf{E}_X \exp(\lambda \langle X, \theta \rangle))$. Conditionally on θ , by the sub-Gaussian assumption, we have

$$\log \mathbf{E}_X \exp(\lambda \langle X, \theta \rangle) \leq 9\kappa^2 \lambda^2 \theta^T \Sigma \theta, \tag{2.9}$$

where to get the explicit constant, one should keep track of the constant factors in the implications of [50, Proposition 2.5.2]. We have

$$\mathbf{E}_{\rho_v} (\theta^T \Sigma \theta) = v^T \Sigma v + \beta^{-1} \text{Tr}(\Sigma).$$

Therefore, for any $\lambda > 0$, simultaneously for all $v \in S^{d-1}$, we have, with probability at least $1 - \exp(-t)$,

$$\langle X, v \rangle \leq 9\kappa^2 \lambda (\beta^{-1} \text{Tr}(\Sigma) + \|\Sigma\|) + \lambda^{-1}(\beta/2 + t).$$

Choosing $\lambda = \sqrt{\frac{\beta/2+t}{9\kappa^2(\beta^{-1} \text{Tr}(\Sigma) + \|\Sigma\|)}}$ and $\beta = \sqrt{\frac{2t \text{Tr}(\Sigma)}{\|\Sigma\|}}$, we prove the claim. \square

Remark 2.9. The leading constant in Lemma 2.7 can be made optimal if we assume a bound on the moment generating function as in the Gaussian case. That is, if instead of (2.9), we have, conditionally on θ ,

$$\log \mathbf{E}_X \exp(\lambda \langle X, \theta \rangle) \leq \lambda^2 \theta^T \Sigma \theta / 2,$$

then, optimizing with respect to λ and β , one can show that, with probability at least $1 - \exp(-t)$, it holds that

$$\|X\|^2 \leq \text{Tr}(\Sigma) + 2\sqrt{2t \text{Tr}(\Sigma) \|\Sigma\|} + 2t\|\Sigma\|.$$

This is what one can achieve if X is a zero mean Gaussian vector, in which case one can use the Gaussian concentration inequality (see [8, Example 5.7]). This observation appears (implicitly) in the works of Catoni and co-authors.

Define the truncation function

$$\psi(x) = \begin{cases} x, & \text{for } x \in [-1, 1]; \\ \text{sign}(x), & \text{for } |x| > 1. \end{cases} \tag{2.10}$$

It is a symmetric function such that for all $x \in \mathbb{R}$,

$$\psi(x) \leq \log(1 + x + x^2).$$

The following bound will be important in our analysis.

Lemma 2.10. *Let ψ be as above and let Z be a square integrable random variable. We have*

$$\psi(\mathbf{E}Z) \leq \mathbf{E} \log(1 + Z + Z^2) + \min\{1, \mathbf{E}Z^2/6\}.$$

Moreover, for any $a > 0$, it holds that

$$\mathbf{E} \log(1 + Z + Z^2) + a \mathbf{E} \min\{1, Z^2/6\} \leq \mathbf{E} \log \left(1 + Z + \left(1 + \frac{(7 + \sqrt{6})(\exp(a) - 1)}{6} \right) Z^2 \right).$$

Proof. In the proof we use the following fact. The function

$$x \mapsto \log(1 + x + x^2) + x^2/6$$

is convex. Observe also that if $0 \leq t \leq a$, then

$$\exp(t) \leq 1 + \frac{t(e^a - 1)}{a}. \tag{2.11}$$

We proceed with the following lines

$$\begin{aligned} \psi(\mathbf{E}Z) &= \min\{\psi(\mathbf{E}Z), 1\} \\ &\leq \min\{\psi(\mathbf{E}Z) + (\mathbf{E}Z)^2/6, 1\} \\ &\leq \min\{\log(1 + \mathbf{E}Z + (\mathbf{E}Z)^2) + (\mathbf{E}Z)^2/6, 1\} \\ &\leq \min\{\mathbf{E}(\log(1 + Z + Z^2) + Z^2/6), 1\} \quad (\text{By Jensen's inequality}) \\ &\leq \mathbf{E} \log(1 + Z + Z^2) + \min\{1, \mathbf{E}Z^2/6\}. \end{aligned}$$

For the second inequality we have

$$\begin{aligned} &\mathbf{E} \log(1 + Z + Z^2) + a \mathbf{E} \min\{1, Z^2/6\} \\ &= \mathbf{E} \log \left((1 + Z + Z^2) \exp(\min\{a, aZ^2/6\}) \right) \\ &\leq \mathbf{E} \log \left((1 + Z + Z^2) (1 + \min\{1, Z^2/6\} (e^a - 1)) \right) \quad \text{by (2.11)}. \end{aligned}$$

Further, we have

$$\begin{aligned} (1 + Z + Z^2) \min\{1, Z^2/6\} &\leq (Z^2/6 + Z^2 + Z \min\{1, Z^2/6\}) \\ &\leq (7/6 + 1/\sqrt{6})Z^2, \end{aligned}$$

where we used that if $|Z| \leq \sqrt{6}$, then $Z \min\{1, Z^2/6\} \leq |Z|^3/6 \leq Z^2/\sqrt{6}$. Otherwise, if $|Z| \geq \sqrt{6}$, then $Z \leq Z^2/\sqrt{6}$. These computations conclude the proof. \square

This allows us to prove the following uniform bound.

Lemma 2.11. Assume that X_1, \dots, X_n are independent copies of a zero-mean random vector X in \mathbb{R}^d with covariance Σ . Let s be an integer and assume that X satisfies that for some $\eta \geq 1$ and any $y \in \mathbb{R}^d$,

$$(\mathbf{E}|\langle y, X \rangle|^{2s})^{\frac{1}{2s}} \leq \eta \sqrt{y^T \Sigma y}. \tag{2.12}$$

Fix the truncation level

$$\lambda = \sqrt{\frac{\mathbf{r}(\Sigma) + t}{n\eta^{2s}\|\Sigma\|^s}}.$$

Then there is $C_s > 0$ that depends only on s such that for any $t > 0$, with probability at least $1 - 2 \exp(-t)$, it holds that

$$\sup_{v \in S^{d-1}} \left| \frac{1}{n\lambda} \sum_{i=1}^n \psi(\lambda \langle v, X_i \rangle^s) - \mathbf{E} \langle v, X \rangle^s \right| \leq C_s \eta^s \|\Sigma\|^{s/2} \sqrt{\frac{\mathbf{r}(\Sigma) + t}{n}}.$$

Proof. Fix $\beta > 0$ and let $\Theta = (\mathbb{R}^d)^s$. Choose the prior distribution μ on Θ as the product of s multivariate Gaussian distributions in \mathbb{R}^d with mean zero and covariance $\beta^{-1}I_d$. For $v \in S^{d-1}$ let the posterior distribution $\rho_{v, \dots, v}$ be defined as the product of s multivariate Gaussian distributions in \mathbb{R}^d each with mean v and covariance $\beta^{-1}I_d$. For $(\theta_1, \dots, \theta_s)$ distributed according to $\rho_{v, \dots, v}$, we have

$$\mathbf{E}_{\rho_{v, \dots, v}}(\theta_1, \dots, \theta_s) = (v, \dots, v), \tag{2.13}$$

In what follows, we use the simplifying notation $\rho = \rho_{v, \dots, v}$ and the marginals of ρ will be denoted by ρ_k for all $k = 1, \dots, s$. Using the additivity of the Kullback-Leibler divergence for product measures, we get

$$\mathcal{KL}(\rho_{v, \dots, v}, \mu) = s\beta/2.$$

Fix $\lambda > 0$ and let $\beta = \mathbf{r}(\Sigma)$. Our plan will be to bound $\psi(\lambda \langle v, X \rangle^s)$ by

$$\mathbf{E}_\rho \log \left(1 + \lambda \prod_{k=1}^s \langle X, \theta_k \rangle + c_s \lambda^2 \prod_{k=1}^s \langle X, \theta_k \rangle^2 \right),$$

where $c_s > 0$ depends only on s and some other terms that do not depend on v ; then we apply Lemma 2.1. This idea is related to the influence function approach used by Catoni [10, 11] in the context of robust estimation. Using (2.13) and the first part of Lemma 2.10, we have

$$\begin{aligned} \psi(\lambda \langle v, X \rangle^s) &= \psi \left(\lambda \mathbf{E}_\rho \prod_{k=1}^s \langle X, \theta_k \rangle \right) \\ &\leq \mathbf{E}_\rho \log \left(1 + \lambda \prod_{k=1}^s \langle X, \theta_k \rangle + \lambda^2 \prod_{k=1}^s \langle X, \theta_k \rangle^2 \right) + \min \left\{ 1, \lambda^2 \mathbf{E}_\rho \prod_{k=1}^s \langle X, \theta_k \rangle^2 / 6 \right\}. \end{aligned}$$

We start with the last summand. Conditionally on X , the following holds:

$$\mathbf{E}_\rho \prod_{k=1}^s \langle X, \theta_k \rangle^2 = \prod_{k=1}^s \mathbf{E}_{\rho_k} \langle X, \theta_k \rangle^2 = (\langle X, v \rangle^2 + \beta^{-1} \|X\|^2)^s \leq 2^{s-1} (\langle X, v \rangle^{2s} + \beta^{-s} \|X\|^{2s}).$$

Therefore, we have

$$\min \left\{ 1, \lambda^2 \mathbf{E}_\rho \prod_{k=1}^s \langle X, \theta_k \rangle^2 / 6 \right\} \leq \min \{ 1, 2^{s-1} \lambda^2 \langle X, v \rangle^{2s} / 6 \} + \min \{ 1, 2^{s-1} \lambda^2 \beta^{-s} \|X\|^{2s} / 6 \}.$$

We observe that conditionally on X , for each k , the distribution of $\langle X, \theta_k \rangle$ is Gaussian with mean $\langle X, v \rangle$. Since it is symmetric, we have that $\Pr(\langle X, \theta_k \rangle^2 \geq \langle X, v \rangle^2) \geq \frac{1}{2}$. Thus, with probability at least $\frac{1}{2^s}$, it holds that $\prod_{k=1}^s \langle X, \theta_k \rangle^2 \geq \langle X, v \rangle^{2s}$. This observation implies that

$$\min \{ 1, 2^{s-1} \lambda^2 \langle X, v \rangle^{2s} / 6 \} \leq 2^s \mathbf{E}_\rho \min \left\{ 1, 2^{s-1} \lambda^2 \prod_{k=1}^s \langle X, \theta_k \rangle^2 / 6 \right\}.$$

Choosing $a = 2^s \max\{1, 2^{s-1}/6\}$ in the second part of Lemma 2.10, we have that there is $c_s > 1$ depending only on s such that

$$\begin{aligned} \mathbf{E}_\rho \log \left(1 + \lambda \prod_{k=1}^s \langle X, \theta_k \rangle + \lambda^2 \prod_{k=1}^s \langle X, \theta_k \rangle^2 \right) + 2^s \mathbf{E}_\rho \min \left\{ 1, 2^{s-1} \lambda^2 \prod_{k=1}^s \langle X, \theta_k \rangle^2 / 6 \right\} \\ \leq \mathbf{E}_\rho \log \left(1 + \lambda \prod_{k=1}^s \langle X, \theta_k \rangle + c_s \lambda^2 \prod_{k=1}^s \langle X, \theta_k \rangle^2 \right). \end{aligned}$$

We plug $f(X, \theta_1, \dots, \theta_s) = \log \left(1 + \lambda \prod_{k=1}^s \langle X, \theta_k \rangle + c_s \lambda^2 \prod_{k=1}^s \langle X, \theta_k \rangle^2 \right)$ into Lemma 2.1. Using $\log(1 + y) \leq y$ for $y \geq -1$, we have

$$\begin{aligned} \mathbf{E}_\rho \log \mathbf{E} \exp(f(X, \theta_1, \dots, \theta_s)) &\leq \lambda \mathbf{E}_\rho \mathbf{E} \prod_{k=1}^s \langle X, \theta_k \rangle + c_s \lambda^2 \mathbf{E}_\rho \mathbf{E} \prod_{k=1}^s \langle X, \theta_k \rangle^2 \\ &= \lambda \mathbf{E} \langle X, v \rangle^s + c_s \lambda^2 \mathbf{E}_\rho \mathbf{E} \prod_{k=1}^s \langle X, \theta_k \rangle^2, \end{aligned}$$

We need to upper bound the last term. Applying Hölder's inequality and since θ_k are independent for $k = 1, \dots, s$, we have

$$\mathbf{E}_\rho \mathbf{E} \prod_{k=1}^s \langle X, \theta_k \rangle^2 \leq \mathbf{E}_\rho \prod_{k=1}^s (\mathbf{E} \langle X, \theta_k \rangle^{2s})^{\frac{1}{s}} \leq \mathbf{E}_\rho \prod_{k=1}^s \left(\eta \sqrt{\theta_k^\top \Sigma \theta_k} \right)^2 \leq (2\eta^2 \|\Sigma\|)^s, \quad (2.14)$$

where we used that for our choice $\beta = \mathbf{r}(\Sigma)$,

$$\mathbf{E}_{\rho_k} \theta_k^\top \Sigma \theta_k \leq \|\Sigma\| + \beta^{-1} \text{Tr}(\Sigma) = 2\|\Sigma\|.$$

Thus, Lemma 2.1 implies, that with probability at least $1 - \exp(-t)$, for all $v \in S^{d-1}$,

$$\begin{aligned} \frac{1}{n\lambda} \sum_{i=1}^n \mathbf{E}_\rho \log \left(1 + \lambda \prod_{k=1}^s \langle X, \theta_k \rangle + c_s \lambda^2 \prod_{k=1}^s \langle X, \theta_k \rangle^2 \right) \\ \leq \mathbf{E} \langle X, v \rangle^s + c_s \lambda (2\eta^2 \|\Sigma\|)^s + \frac{s\mathbf{r}(\Sigma) + 2t}{2n\lambda}. \end{aligned}$$

By the above computations we have on the same event

$$\begin{aligned} & \frac{1}{n\lambda} \sum_{i=1}^n \psi(\lambda \langle v, X_i \rangle^s) \\ & \leq \mathbf{E} \langle X, v \rangle^s + c_s \lambda (2\eta^2 \|\Sigma\|)^s + \frac{s\mathbf{r}(\Sigma) + 2t}{2n\lambda} + \frac{1}{n\lambda} \sum_{i=1}^n \min \left\{ 1, \frac{2^{s-1} \lambda^2 \|X_i\|^{2s}}{6\beta^s} \right\}. \end{aligned}$$

It is left to control the last sum. Denote $Y = \min \left\{ 1, \frac{2^{s-1} \lambda^2 \|X\|^{2s}}{6\beta^s} \right\}$ and let Y_1, \dots, Y_n be independent copies of Y . Observe that $Y \in [0, 1]$, which implies $\text{Var}(Y) \leq \mathbf{E}Y$. By the standard Bernstein inequality [8, Corollary 2.11], with probability at least $1 - \exp(-t)$,

$$\frac{1}{n\lambda} \sum_{i=1}^n Y_i \leq \frac{1}{\lambda} \left(\mathbf{E}Y + \sqrt{\frac{2\mathbf{E}Yt}{n}} + \frac{2t}{3n} \right) \leq \frac{1}{\lambda} \left(2\mathbf{E}Y + \frac{7t}{6n} \right).$$

Finally, denoting the standard basis in \mathbb{R}^d by e_1, \dots, e_d and using the triangle inequality, we have

$$\left(\mathbf{E} \|X\|^{2s} \right)^{\frac{1}{s}} = \left(\mathbf{E} \left(\sum_{j=1}^d \langle e_j, X \rangle^2 \right)^s \right)^{\frac{1}{s}} \leq \sum_{j=1}^d \left(\mathbf{E} \langle e_j, X \rangle^{2s} \right)^{\frac{1}{s}} \leq \eta^2 \text{Tr}(\Sigma),$$

which implies $\mathbf{E}Y \leq 2^{s-1} \eta^{2s} \lambda^2 \|\Sigma\|^s / 6$. Combining these estimates and taking our choice of λ into account, the one-sided bound follows.

To finish the proof, we need to get a two-sided bound. Since the function ψ is symmetric, it is enough to consider $\rho_{v, \dots, v, -v}$ instead of $\rho_{v, \dots, v}$ as the posterior distribution, for which the same analysis holds. The claim follows. \square

2.3.1 Bounds on the norms of random vectors

First, consider the sub-Gaussian case. By [50, Proposition 2.5.2] the sub-Gaussian vector X satisfies for all $q \geq 1$ and any $y \in \mathbb{R}^d$,

$$\left(\mathbf{E} |\langle y, X \rangle|^q \right)^{1/q} \leq 3\sqrt{q} \|\langle y, X \rangle\|_{\psi_2} \leq 3\kappa \sqrt{q y^T \Sigma y}. \tag{2.15}$$

Choosing $q = 2s$, we obtain that when applying Lemma 2.11 we may choose $\eta = 3\kappa\sqrt{2s}$.

The norm of the sub-Gaussian vector can easily be analyzed using Lemma 2.7. Indeed, by the union bound, with probability at least $1 - n \exp(-t)$, it holds that

$$\max_i \|X_i\| \leq \left(36\kappa^2 (\text{Tr}(\Sigma) + 2t\|\Sigma\|) \right)^{\frac{1}{2}} \leq \kappa \sqrt{108 \text{Tr}(\Sigma)}, \tag{2.16}$$

whenever $t \leq \mathbf{r}(\Sigma)$.

Second, we consider the log-concave case. We use Borell's characterization of log-concave distributions (see e.g., [3, Lemma 2.3]): If X is a zero mean random vector in \mathbb{R}^d with a log-concave distribution, then for any $y \in S^{d-1}$,

$$\|\langle y, X \rangle\|_{\psi_1} \leq \kappa \sqrt{y^T \Sigma y}, \tag{2.17}$$

where κ is a universal constant. We use the symbol κ in the proof but absorb it by the generic constant in our final statements. By (2.2) and (2.17) we have for any $y \in S^{d-1}$ and $q \geq 1$,

$$\mathbf{E} |\langle y, X \rangle|^q \leq 2q! \|\langle y, X \rangle\|_{\psi_1}^q \leq 2q^q \|\langle y, X \rangle\|_{\psi_1}^q \leq 2(\kappa q)^q (y^T \Sigma y)^{q/2}. \tag{2.18}$$

Choosing $q = 2s$, we obtain that in Lemma 2.11 we may choose $\eta = 4\kappa s$. The second component is the inequality of Paouris [42], written in the following form [2, Theorem 2]: If X is a random vector in \mathbb{R}^d with a log-concave distribution, then for any $q \geq 1$,

$$(\mathbf{E}\|X\|^q)^{1/q} \leq C(\mathbf{E}\|X\| + \sup_{y \in S^{d-1}} (\mathbf{E}|\langle y, X \rangle|^q)^{1/q}), \tag{2.19}$$

where $C > 0$ is a universal constant. For a zero mean random vector X we have $\mathbf{E}\|X\| \leq \sqrt{\text{Tr}(\Sigma)}$. By (2.18) we have $\mathbf{E}|\langle y, X \rangle|^q \leq 2(\kappa q)^q \|\Sigma\|^{q/2}$. Thus, by Markov's inequality together with (2.19), with probability at least $1 - \exp(-t)$, it holds that $\|X\| \leq eC(\sqrt{\text{Tr}(\Sigma)} + 2\kappa t \sqrt{\|\Sigma\|})$. By the union bound, with probability at least $1 - n \exp(-t)$,

$$\max_i \|X_i\| \leq eC(\sqrt{\text{Tr}(\Sigma)} + 2\kappa t \sqrt{\|\Sigma\|}). \tag{2.20}$$

2.3.2 Proof of Theorem 1.8

We start with an upper tail. Let $\lambda > 0$ be a fixed truncation level. We can write the following decomposition similar to the one used in [3]:

$$\begin{aligned} & \sup_{v \in S^{d-1}} \left(\frac{1}{n} \sum_{i=1}^n \langle v, X_i \rangle^2 - \mathbf{E}\langle v, X \rangle^2 \right) \\ & \leq \sup_{v \in S^{d-1}} \frac{1}{n} \sum_{i=1}^n \langle v, X_i \rangle^2 \mathbf{1}(\lambda \langle v, X_i \rangle^2 > 1) + \sup_{v \in S^{d-1}} \left(\frac{1}{n\lambda} \sum_{i=1}^n \psi(\lambda \langle v, X_i \rangle^2) - \mathbf{E}\langle v, X \rangle^2 \right). \end{aligned} \tag{2.21}$$

The analysis of the second term will follow from Lemma 2.11. Indeed, for $\eta = 8\kappa$, defining

$$\lambda = \frac{1}{\|\Sigma\|} \sqrt{\frac{2\mathbf{r}(\Sigma)}{n(8\kappa)^4}},$$

which will be our choice throughout the proof, we have, with probability at least $1 - 2 \exp(-\mathbf{r}(\Sigma))$,

$$\sup_{v \in S^{d-1}} \left(\frac{1}{n\lambda} \sum_{i=1}^n \psi(\lambda \langle v, X_i \rangle^2) - \mathbf{E}\langle v, X \rangle^2 \right) \leq c\kappa^2 \|\Sigma\| \sqrt{\frac{\mathbf{r}(\Sigma)}{n}},$$

where $c > 0$ is an absolute constant. Since κ is an absolute constant in the log-concave case, we will sometimes absorb it by other absolute constants.

The analysis of large summands can be done via a well-known decoupling-chaining argument. This argument leads to the following result.

Lemma 2.12 (Proposition 9.4.2 in [45]). *Assume that Y_1, \dots, Y_n are independent copies of a random vector Y in \mathbb{R}^d such that for all $x \in S^{d-1}$, it holds that*

$$\|\langle x, Y \rangle\|_{\psi_1} \leq 1.$$

There are absolute constants $c_1, c_2 > 0$ such that the following holds. For any $t > 0$, with probability at least $1 - c_1 \exp(-t)$, we have uniformly over $k = 1, \dots, n$,

$$\sup_{x \in S^{d-1}} \sup_{\substack{I \subseteq \{1, \dots, n\} \\ |I| \leq k}} \left(\sum_{i \in I} \langle x, Y_i \rangle^2 \right)^{\frac{1}{2}} \leq c_2 \left(\max_i \|Y_i\| + \sqrt{k} \log \left(\frac{en}{k} \right) + t \right).$$

Our first observation is that the above bound does not depend on d . Moreover, the distribution of Y is not necessarily isotopic. Thus, we can adapt this result to our case. Recall that by (2.17) for a zero mean, log-concave vector X , we have for any $y \in S^{d-1}$, $\|\langle y, X \rangle\|_{\psi_1} \leq \kappa \sqrt{\|\Sigma\|}$. Thus, denoting $Y = \frac{1}{\kappa \sqrt{\|\Sigma\|}} X$, we have $\|\langle x, Y \rangle\|_{\psi_1} \leq 1$ for all $x \in S^{d-1}$. By Lemma 2.12, with probability at least $1 - c_1 \exp(-t)$, it holds that

$$\sup_{y \in S^{d-1}} \sup_{\substack{I \subseteq \{1, \dots, n\}, \\ |I| \leq k}} \left(\sum_{i \in I} \langle y, X_i \rangle^2 \right)^{\frac{1}{2}} \leq c_2 \max_i \|X_i\| + c_2 \kappa \sqrt{\|\Sigma\|} \left(\sqrt{k} \log \left(\frac{en}{k} \right) + t \right). \quad (2.22)$$

Throughout the proof we choose $t = (\mathbf{r}(\Sigma)n)^{1/4}$. From now on we can follow the arguments in [3, 45] with several modifications needed to take the effective rank into account. Observe that by (2.20), with probability at least

$$1 - n \exp(-2(\mathbf{r}(\Sigma)n)^{1/4}) \geq 1 - \exp(-(\mathbf{r}(\Sigma)n)^{1/4}),$$

it holds that

$$\max_i \|X_i\| \leq eC(1 + 4\kappa) (\|\Sigma\|^2 \mathbf{r}(\Sigma)n)^{1/4}, \quad (2.23)$$

where we used that $\sqrt{\text{Tr}(\Sigma)} = \sqrt{\|\Sigma\| \mathbf{r}(\Sigma)} \leq (\|\Sigma\|^2 \mathbf{r}(\Sigma)n)^{1/4}$ for $n \geq \mathbf{r}(\Sigma)$. Denote

$$m = \sup_{v \in S^{d-1}} |\{i \in \{1, \dots, n\} : \lambda \langle v, X_i \rangle^2 > 1\}|.$$

Case 1 On the event where (2.22) and (2.23) hold, we have that $m \leq \frac{\mathbf{r}(\Sigma)}{\log^2 \left(\frac{e^2 n}{\mathbf{r}(\Sigma)} \right)}$. Observe that $m \mapsto \sqrt{m} \log \left(\frac{e^2 n}{m} \right)$ is increasing for $1 \leq m \leq n$ and since $n \geq \mathbf{r}(\Sigma)$,

$$\sqrt{m} \log \left(\frac{en}{m} \right) \leq \sqrt{\mathbf{r}(\Sigma)} + \sqrt{\mathbf{r}(\Sigma)} \frac{\log \log^2 \left(\frac{e^2 n}{\mathbf{r}(\Sigma)} \right)}{\log \left(\frac{e^2 n}{\mathbf{r}(\Sigma)} \right)} \leq 2\sqrt{\mathbf{r}(\Sigma)} \leq 2(n\mathbf{r}(\Sigma))^{1/4}.$$

Thus, on the same event (2.22) and (2.23) imply

$$\begin{aligned} \sup_{v \in S^{d-1}} \sum_{i=1}^n \langle v, X_i \rangle^2 \mathbf{1}(\lambda \langle v, X_i \rangle^2 > 1) &\leq \sup_{v \in S^{d-1}} \sup_{\substack{I \subseteq \{1, \dots, n\}, \\ |I| \leq m}} \left(\sum_{i \in I} \langle v, X_i \rangle^2 \right) \\ &\leq c_2^2 \left(\max_i \|X_i\| + 3\kappa \sqrt{\|\Sigma\|} \left((n\mathbf{r}(\Sigma))^{1/4} \right) \right)^2 \\ &\leq c_3 \|\Sigma\| \sqrt{n\mathbf{r}(\Sigma)}, \end{aligned}$$

where $c_3 > 0$ is an absolute constant.

Case 2 Otherwise, we assume $m > \frac{\mathbf{r}(\Sigma)}{\log^2 \left(\frac{e^2 n}{\mathbf{r}(\Sigma)} \right)}$. By (2.22), (2.23) and the union bound, we have, with probability at least $1 - (1 + c_1) \exp(-(\mathbf{r}(\Sigma)n)^{1/4})$,

$$\sqrt{m\lambda^{-1}} \leq ec_2 C(1 + 2\kappa) (\|\Sigma\|^2 \mathbf{r}(\Sigma)n)^{1/4} + c_2 \kappa \sqrt{\|\Sigma\|} \left(\sqrt{m} \log \left(\frac{en}{m} \right) + (\mathbf{r}(\Sigma)n)^{1/4} \right). \quad (2.24)$$

Assume that the sample size n is large enough, so that $\lambda^{-1} \geq 4 \left(c_2 \kappa \sqrt{\|\Sigma\|} \log \left(\frac{en}{m} \right) \right)^2$. Solving (2.24) we get, in particular, that

$$m \log^2 \left(\frac{en}{m} \right) \leq \frac{c_4 \sqrt{\mathbf{r}(\Sigma)n}}{4c_2^2 \kappa^2} \quad \text{and} \quad c_2 \kappa \sqrt{\|\Sigma\|} \sqrt{m} \log \left(\frac{en}{m} \right) \leq c_5 (\|\Sigma\|^2 \mathbf{r}(\Sigma)n)^{1/4}, \quad (2.25)$$

where $c_4, c_5 > 0$ are absolute constants. Combining (2.22), (2.25) and dividing both sides by n , we show that for some $c_6 > 0$, on the corresponding event

$$\sup_{v \in S^{d-1}} \frac{1}{n} \sum_{i=1}^n \langle v, X_i \rangle^2 \mathbf{1}(\lambda \langle v, X_i \rangle^2 > 1) \leq \sup_{v \in S^{d-1}} \sup_{\substack{I \subseteq \{1, \dots, n\}, \\ |I| \leq m}} \left(\frac{1}{n} \sum_{i \in I} \langle v, X_i \rangle^2 \right) \leq c_6 \|\Sigma\| \sqrt{\frac{\mathbf{r}(\Sigma)}{n}}.$$

It is only left to check that $\lambda^{-1} \geq 4 \left(c_2 \kappa \sqrt{\|\Sigma\|} \log \left(\frac{en}{m} \right) \right)^2$, which is

$$\sqrt{\frac{n(8\kappa)^4}{2\mathbf{r}(\Sigma)}} \geq 4 \left(c_2 \kappa \log \left(\frac{en}{m} \right) \right)^2. \tag{2.26}$$

Since we assumed that $m > \frac{\mathbf{r}(\Sigma)}{\log^2 \left(\frac{e^2 n}{\mathbf{r}(\Sigma)} \right)}$, one can easily show that if $n \geq c_7 \mathbf{r}(\Sigma)$ for some $c_7 > 0$, then the inequality (2.26) holds. Indeed, $x \mapsto \sqrt{x}$ grows faster than $x \mapsto \log^2(x)$. We proved the upper tail bound.

Although one can prove the lower tail bound completely analogously, we discuss an alternative argument. For the same value of λ , using the definition of the truncation function and since $\langle v, X \rangle^2 \geq 0$, we have, with probability at least $1 - 2 \exp(-\mathbf{r}(\Sigma))$,

$$\begin{aligned} \sup_{v \in S^{d-1}} \left(\mathbf{E} \langle v, X \rangle^2 - \frac{1}{n} \sum_{i=1}^n \langle v, X_i \rangle^2 \right) &\leq \sup_{v \in S^{d-1}} \left(\mathbf{E} \langle v, X \rangle^2 - \frac{1}{n\lambda} \sum_{i=1}^n \psi(\lambda \langle v, X_i \rangle^2) \right) \\ &\leq c\kappa^2 \|\Sigma\| \sqrt{\frac{\mathbf{r}(\Sigma)}{n}}. \end{aligned}$$

Applying the union bound we finish the proof. We discuss a related argument in Section 3.2. □

2.3.3 Proof of Proposition 1.9

There are only a few changes compared to the proof of Theorem 1.8. First, an analog of the norm bound (2.23) is implied by our assumption. When applying Lemma 2.11 to the second term in (2.21) we choose

$$\lambda = \frac{1}{\|\Sigma\|} \sqrt{\frac{\mathbf{r}(\Sigma) + t}{n(8\kappa)^4}}.$$

Our second modification is that we consider the cases: $m \leq \frac{\mathbf{r}(\Sigma) + t}{\log^2 \left(\frac{e^2 n}{\mathbf{r}(\Sigma) + t} \right)}$ and $m > \frac{\mathbf{r}(\Sigma) + t}{\log^2 \left(\frac{e^2 n}{\mathbf{r}(\Sigma) + t} \right)}$. In the first case $\sqrt{m} \log \left(\frac{en}{m} \right) \leq 2\sqrt{\mathbf{r}(\Sigma) + t}$. In this case, with probability at least $1 - c_1 \exp(-t)$, it holds that

$$\sup_{v \in S^{d-1}} \sum_{i=1}^n \langle v, X_i \rangle^2 \mathbf{1}(\lambda \langle v, X_i \rangle^2 > 1) \leq c_2^2 \kappa^2 \|\Sigma\| \left((\mathbf{r}(\Sigma)n)^{\frac{1}{4}} + 2\sqrt{\mathbf{r}(\Sigma) + t} + t \right)^2,$$

Otherwise, if $m > \frac{\mathbf{r}(\Sigma) + t}{\log^2 \left(\frac{e^2 n}{\mathbf{r}(\Sigma) + t} \right)}$ following the same lines, we show that

$$\sup_{v \in S^{d-1}} \frac{1}{n} \sum_{i=1}^n \langle v, X_i \rangle^2 \mathbf{1}(\lambda \langle v, X_i \rangle^2 > 1) \leq c_3 \kappa^2 \|\Sigma\| \left(\sqrt{\mathbf{r}(\Sigma)n} + t^2 \right),$$

where $c_3 > 0$ is an absolute constant. When $m > \frac{\mathbf{r}(\Sigma)+t}{\log^2\left(\frac{e^2 n}{\mathbf{r}(\Sigma)+t}\right)}$, the condition (2.26) will be rewritten as

$$\sqrt{\frac{n(8\kappa)^4}{\mathbf{r}(\Sigma)+t}} \geq 4 \left(c_2 \kappa \log\left(\frac{en}{m}\right) \right)^2,$$

which implies that it is sufficient to take $n \geq c_4(\mathbf{r}(\Sigma)+t)$. Combining these bounds as in the proof of Theorem 1.8, we prove the claim. \square

2.3.4 Proof of Theorem 1.7

The log-concave case The proof for the log-concave case is quite similar to the proof of Theorem 1.8. Let us start with an upper tail. Let $\lambda_1 > 0$ be a fixed truncation level. As above, we write the following decomposition:

$$\begin{aligned} & \sup_{v \in S^{d-1}} \left(\frac{1}{n} \sum_{i=1}^n \langle v, X_i \rangle^s - \mathbf{E} \langle v, X \rangle^s \right) \\ & \leq \sup_{v \in S^{d-1}} \frac{1}{n} \sum_{i=1}^n |\langle v, X_i \rangle|^s \mathbf{1}(\lambda_1 |\langle v, X_i \rangle|^s > 1) + \sup_{v \in S^{d-1}} \left(\frac{1}{n\lambda_1} \sum_{i=1}^n \psi(\lambda_1 \langle v, X_i \rangle^s) - \mathbf{E} \langle v, X \rangle^s \right). \end{aligned}$$

The analysis of the second term will follow from Lemma 2.11. For $\eta = 4s\kappa$, defining $\lambda_1 = \sqrt{\frac{2\mathbf{r}(\Sigma)}{n(4s\kappa)^{2s}\|\Sigma\|^s}}$, we have, with probability at least $1 - 2\exp(-\mathbf{r}(\Sigma))$,

$$\sup_{v \in S^{d-1}} \left(\frac{1}{n\lambda_1} \sum_{i=1}^n \psi(\lambda_1 \langle v, X_i \rangle^2) - \mathbf{E} \langle v, X \rangle^2 \right) \leq C_s (4s\kappa)^s \|\Sigma\|^{s/2} \sqrt{\frac{\mathbf{r}(\Sigma)+t}{n}},$$

where C_s depends only on s . Observe that $\lambda_1 |\langle v, X_i \rangle|^s > 1$ is equivalent to $\langle v, X_i \rangle^2 > \lambda_1^{-2/s}$. Denoting $\lambda^{-1} = \lambda_1^{-2/s}$, we repeat the arguments of the proof of Theorem 1.8 for the term

$$\sup_{v \in S^{d-1}} \frac{1}{n} \sum_{i=1}^n \langle v, X_i \rangle^2 \mathbf{1}(\lambda \langle v, X_i \rangle^2 > 1)$$

with just two modifications: We use that, with probability at least $1 - n \exp(-\sqrt{\mathbf{r}(\Sigma)})$,

$$\max_i \|X_i\| \leq \epsilon C(1+2\kappa)\sqrt{\text{Tr}(\Sigma)},$$

and choose $t = \sqrt{\mathbf{r}(\Sigma)}$ when applying Lemma 2.12 in the form (2.22). The same lines imply that if $\lambda_1^{-2/s} \geq 4 \left(c_2 \kappa \sqrt{\|\Sigma\|} \log\left(\frac{en}{m}\right) \right)^2$ for $m \geq \frac{\mathbf{r}(\Sigma)}{\log^2\left(\frac{e^2 n}{\mathbf{r}(\Sigma)}\right)}$, which is

$$\left(\frac{n(4s\kappa)^{2s}\|\Sigma\|^s}{2\mathbf{r}(\Sigma)} \right)^{\frac{1}{s}} \geq 4 \left(c_2 \kappa \sqrt{\|\Sigma\|} \log\left(\frac{en}{m}\right) \right)^2, \tag{2.27}$$

then, with probability at least $1 - (n+c_1)\exp(-\sqrt{\mathbf{r}(\Sigma)})$,

$$\begin{aligned} \sup_{v \in S^{d-1}} \left(\sum_{i=1}^n |\langle v, X_i \rangle|^s \mathbf{1}(|\langle v, X_i \rangle|^s > \lambda_1^{-1}) \right)^{\frac{1}{s}} & \leq \sup_{v \in S^{d-1}} \left(\sum_{i=1}^n \langle v, X_i \rangle^2 \mathbf{1}(\langle v, X_i \rangle^2 > \lambda^{-1}) \right)^{\frac{1}{2}} \\ & \leq c_8 \sqrt{\|\Sigma\| \mathbf{r}(\Sigma)}, \end{aligned}$$

where c_8 is an absolute constant. This inequality will contribute the term $\left(c_8 \sqrt{\|\Sigma\| \mathbf{r}(\Sigma)} \right)^s$ to the final bound. For this term to be consistent with our final bound, we need

$$\frac{(c_8 \sqrt{\|\Sigma\| \mathbf{r}(\Sigma)})^s}{n} \leq \|\Sigma\|^{s/2} \sqrt{\frac{\mathbf{r}(\Sigma)}{n}}.$$

It is easy to verify that both the last inequality and (2.27) are satisfied when (1.8) is satisfied. Indeed, in (2.27) we are essentially comparing $\left(\frac{n}{\mathbf{r}(\Sigma)}\right)^{\frac{1}{s}}$ with $\log^2\left(\frac{en}{\mathbf{r}(\Sigma)}\right)$, which implies the need for an s -dependent factor c_s in (1.8).

The analysis of the lower tail can be done analogously. Indeed, we may repeat the analysis for the following decomposition:

$$\begin{aligned} \sup_{v \in S^{d-1}} \left(\mathbf{E} \langle v, X \rangle^s - \frac{1}{n} \sum_{i=1}^n \langle v, X_i \rangle^s \right) &\leq \sup_{v \in S^{d-1}} \left(\frac{1}{n} \sum_{i=1}^n |\langle v, X_i \rangle|^s \mathbf{1}(\lambda_1 |\langle v, X_i \rangle|^s > 1) \right) \\ &\quad + \sup_{v \in S^{d-1}} \left(\mathbf{E} \langle v, X \rangle^s - \frac{1}{n\lambda_1} \sum_{i=1}^n \psi(\lambda_1 \langle v, X_i \rangle^s) \right). \end{aligned}$$

The second part of the claim follows.

The sub-Gaussian case We mainly follow the previous proof with several modifications. First, when applying Lemma 2.11 we can use $\eta = 3\kappa\sqrt{2s}$ as discussed above. Second, we use (see [50, Lemma 2.7.7])

$$\sqrt{\log 2} \|\langle y, X \rangle\|_{\psi_1} \leq \|\langle y, X \rangle\|_{\psi_2} \leq \kappa \sqrt{\|\Sigma\|},$$

and repeat the same lines with (2.23) replaced by (2.16) to achieve our bound.

Our final observation is that if n is large, then, with high probability, $\lambda \max_i |\langle v, X_i \rangle|^s \leq \lambda \max_i \|X_i\|^s \leq 1$, where λ is chosen as in Lemma 2.11. On the corresponding event, we have

$$\begin{aligned} \sup_{v \in S^{d-1}} \left| \frac{1}{n} \sum_{i=1}^n \langle v, X_i \rangle^s - \mathbf{E} \langle v, X \rangle^s \right| &= \sup_{v \in S^{d-1}} \left| \frac{1}{n\lambda} \sum_{i=1}^n \psi(\lambda \langle v, X_i \rangle^s) - \mathbf{E} \langle v, X \rangle^s \right| \\ &\leq C_s (3\kappa\sqrt{2s})^s \|\Sigma\|^{s/2} \sqrt{\frac{2\mathbf{r}(\Sigma)}{n}}, \end{aligned}$$

where C_s depends only on s . By (2.16), with probability at least $1 - n \exp(-\mathbf{r}(\Sigma))$, we have $\lambda \max_i \|X_i\|^s \leq 1$ when

$$\sqrt{\frac{2\mathbf{r}(\Sigma)}{n(3\kappa\sqrt{2s})^{2s} \|\Sigma\|^s}} (108\kappa^2 \text{Tr}(\Sigma))^{\frac{s}{2}} \leq 1,$$

Solving this with respect to n , we see that it is enough to take $n \geq c_9(\mathbf{r}(\Sigma))^{s+1}$, where $c_9 > 0$ is an absolute constant. Applying the union bound, we complete the proof of the statement. \square

Remark 2.13. To improve the tail estimate in the Gaussian case for smaller values of n , one needs to prove a ψ_2 -version of Lemma 2.12. In our analysis we simply used that $\sqrt{\log 2} \|\langle y, X \rangle\|_{\psi_1} \leq \|\langle y, X \rangle\|_{\psi_2}$.

3 Additional results

We provide two results extending the ideas used in Theorem 1.2 and Theorem 1.7.

3.1 Deviations of the norm of sub-exponential random vectors

First, assume that X is a zero mean random vector with covariance Σ satisfying

$$\|\langle X, y \rangle\|_{\psi_1} \leq \kappa \sqrt{y^T \Sigma y}, \tag{3.1}$$

for all $y \in \mathbb{R}^d$. For example, the aforementioned result shows that the assumption (3.1) is implied by the log-concavity assumption. We prove the following result, which improves the bound of Alesker [4] on the concentration of the Euclidean norm of a point sampled uniformly at random from a convex body in the isotropic position. Our proof is shorter and does not involve any explicit calculations in polar coordinates. This bound can also be seen as a special instance of the recent result of Latała and Nayar [28] relating the weak and strong moments of random vectors. Although their result is not limited to the sub-exponential assumption (3.1), our bound has the advantage of being dimension-free.

Proposition 3.1. *Assume that X is a zero mean sub-exponential random vector (3.1). Then, for any $t \geq 1$, with probability at least $1 - \exp(-t)$,*

$$\|X\| \leq 8\kappa(\sqrt{t \operatorname{Tr}(\Sigma)} + t\sqrt{\|\Sigma\|}).$$

Proof. We follow the proof of Theorem 1.2, but with a different choice of β . To control the norm, it is enough to upper bound $\langle X, v \rangle$ uniformly over $v \in S^{d-1}$. We assume without loss of generality that Σ is invertible. Set $\Theta = \mathbb{R}^d$. Let the prior distribution μ be the multivariate Gaussian distribution with mean zero and covariance $\beta^{-1}\Sigma$. For $u \in \Sigma^{1/2}S^{d-1}$ let the distribution ρ_u be defined by the density function (2.4). Choosing $r = \sqrt{2\beta^{-1} \operatorname{Tr}(\Sigma)}$, similarly to (2.6) we have

$$\mathcal{KL}(\rho_u, \mu) = \log 2 + \beta/2.$$

For $\lambda > 0$ we want to plug $\lambda \langle X, \Sigma^{-1/2}\theta \rangle$ into Lemma 2.1, where θ is distributed according to ρ_u . Since $\|\theta\| \leq \sqrt{\|\Sigma\| + \sqrt{2\beta^{-1} \operatorname{Tr}(\Sigma)}}$ almost surely, we have, conditionally on θ ,

$$\|\langle X, \Sigma^{-1/2}\theta \rangle\|_{\psi_1} \leq \kappa\|\theta\| \leq \kappa(\sqrt{\|\Sigma\|} + \sqrt{2\beta^{-1} \operatorname{Tr}(\Sigma)}).$$

By Lemma 2.5, conditionally on θ ,

$$\log \mathbf{E}_X \exp(\lambda \langle X, \Sigma^{-1/2}\theta \rangle) \leq 4\lambda^2 \kappa^2 \|\theta\|^2 \leq 4\lambda^2 \kappa^2 (\sqrt{\|\Sigma\|} + \sqrt{2\beta^{-1} \operatorname{Tr}(\Sigma)})^2,$$

whenever $\lambda \leq \frac{1}{2\kappa(\sqrt{\|\Sigma\|} + \sqrt{2\beta^{-1} \operatorname{Tr}(\Sigma)})}$. We choose $\lambda = \frac{1}{2\kappa(\sqrt{\|\Sigma\|} + \sqrt{2\beta^{-1} \operatorname{Tr}(\Sigma)})}$. By Lemma 2.1 we have for all $v \in S^{d-1}$, with probability at least $1 - \exp(-t)$,

$$\langle X, v \rangle \leq 4\lambda \kappa^2 (\sqrt{\|\Sigma\|} + \sqrt{2\beta^{-1} \operatorname{Tr}(\Sigma)})^2 + \frac{\log 2 + \beta/2 + t}{\lambda}.$$

Let $\beta = 2t$. On the same event we have, using $t \geq 1$,

$$\|X\| \leq (2 + 2 \log 2)\kappa(\sqrt{\|\Sigma\|} + \sqrt{\operatorname{Tr}(\Sigma)}) + 4\kappa(t\sqrt{\|\Sigma\|} + \sqrt{t \operatorname{Tr}(\Sigma)}).$$

The claim follows. □

The result of Proposition 3.1 implies the bound of Alesker [4] in the isotropic case (in this case, $\Sigma = I_d$). Importantly, the bound in [4] scales as \sqrt{td} for all values of t . An inspection of their proof shows that this happens because for $t \geq d$ one can use the boundedness of a unit volume convex body in the isotropic position.

Let us shortly explain why the result of Proposition 3.1 cannot be improved in general. Consider the isotropic case $\Sigma = I_d$. Using the standard tools, one can rewrite the bound of Proposition 3.1 as follows: for $p \geq 1$, and some $c > 0$ that depends only on κ ,

$$\mathbf{E}(\|X\|^p)^{1/p} \leq c(\sqrt{pd} + p). \tag{3.2}$$

In [27] the following example is considered: Let $X = \sqrt{d}ZU$, where Z has the standard Gaussian distribution and U is uniformly distributed on S^{d-1} , independent of Z . It is easy to show that (3.1) is satisfied for X with κ being an absolute constant. In this case $\mathbf{E}(\|X\|^p)^{1/p}$ scales as \sqrt{pd} up to multiplicative constant factors, which proves the necessity of the first term in (3.2). The necessity of the second term is implied by the case where $d = 1$, and X is a standard exponentially distributed random variable.

3.2 A lower tail version of Theorem 1.2

It is known that the least singular value of a random matrix can be controlled with high probability guarantees (in fact, it has sub-Gaussian tails) under significantly milder assumptions. This can be seen as a high-dimensional extension of the following fact. For a non-negative random variable Z and for any $t \geq 0$,

$$\Pr(\mathbf{E}Z - Z \geq t) \leq \exp(-t^2/(2\mathbf{E}Z^2)). \tag{3.3}$$

That is, the one-sided bound allows for sub-Gaussian concentration even if the random variable Z has infinite moments after the first two. There is a list of dimension-dependent lower tail bounds for the least singular value under minimal assumptions [44, 40, 26, 46, 52, 11, 38]. Our next result improves the dimension-dependent bound of Oliveira [40] and complements the result of Theorem 1.2. Although one can modify the original proof in [40]² to get a dimension-free bound, we present a short and self-contained proof.

Proposition 3.2. *Assume that M_1, \dots, M_n are independent copies of a positive semi-definite symmetric random matrix M with mean $\mathbf{E}M = \Sigma$. Let M satisfy for some $\kappa \geq 1$,*

$$\sqrt{\mathbf{E}(x^\top Mx)^2} \leq \kappa^2 x^\top \Sigma x, \tag{3.4}$$

for all $x \in \mathbb{R}^d$. Then, for any $t > \log 2$, with probability at least $1 - 2 \exp(-t)$, it holds that

$$\sup_{v \in S^{d-1}} \left(v^\top \Sigma v - \frac{1}{n} \sum_{i=1}^n v^\top M_i v \right) \leq 7\kappa^2 \|\Sigma\| \sqrt{\frac{\mathbf{r}(\Sigma) + t}{n}}.$$

Proof. Set $\Theta = \mathbb{R}^d$. Let the prior distribution μ be a multivariate Gaussian distribution with mean zero and covariance $\beta^{-1}I_d$. For $v \in S^{d-1}$ let ρ_v be a multivariate Gaussian distribution with mean v and covariance $\beta^{-1}I_d$. We have $\mathcal{KL}(\rho_v, \mu) = \beta/2$.

Let θ be distributed according to ρ_v . Given a positive semi-definite matrix C , we have

$$\mathbf{E}_{\rho_v} \theta^\top C \theta = v^\top C v + \beta^{-1} \text{Tr}(C). \tag{3.5}$$

Let Z be a random vector having a multivariate Gaussian distribution with zero mean and covariance $\beta^{-1}\Sigma$. Define $\varphi = \Sigma^{1/2}v$. We have, using $(a + b)^4 \leq 8a^4 + 8b^4$ and the formulas for the moments of the Gaussian distribution,

$$\begin{aligned} \mathbf{E}_{\rho_v} (\theta^\top \Sigma \theta)^2 &= \mathbf{E} (\|Z + \varphi\|)^4 \leq 8 \left(3\beta^{-2} \sum_{i=1}^d \Sigma_{ii}^2 + 2\beta^{-2} \sum_{i < j} (\Sigma_{ii}\Sigma_{jj} + 2\Sigma_{ij}^2) \right) + 8\|\varphi\|^4. \\ &= 8(\beta^{-2}(2 \text{Tr}(\Sigma^2) + \text{Tr}(\Sigma)^2) + \|\Sigma\|^2) \leq 24\beta^{-2} \text{Tr}(\Sigma)^2 + 8\|\Sigma\|^2. \end{aligned} \tag{3.6}$$

Consider the negative part of the truncation function. That is, let $\psi : (-\infty, 0] \rightarrow [-1, 0]$ be defined as follows:

$$\psi(x) = \begin{cases} x, & \text{for } x \in [-1, 0]; \\ -1, & \text{for } x < -1. \end{cases} \tag{3.7}$$

One can easily check that for $x \leq 0$ it holds that $x \leq \psi(x) \leq \log(1+x+x^2/2)$. Fix $\lambda > 0$ and consider the function $\psi(-\lambda\theta^\top M\theta)$. Observe that $-\lambda\theta^\top M\theta \leq 0$ almost surely, so that the function is well defined. We want to plug it into Lemma 2.1. Using $\psi(x) \leq \log(1+x+x^2/2)$ for $x \leq 0$ and $\log(1+y) \leq y$ for $y > -1$, we have

$$\begin{aligned} \mathbf{E}_{\rho_v} \log(\mathbf{E} \exp(\psi(-\lambda\theta^\top M\theta))) &\leq \mathbf{E}_{\rho_v} \log(\mathbf{E}(1 - \lambda\theta^\top M\theta + \lambda^2(\theta^\top M\theta)^2/2)) \\ &\leq -\lambda \mathbf{E}_{\rho_v} \theta^\top M\theta + \lambda^2 \mathbf{E} \mathbf{E}_{\rho_v} (\theta^\top M\theta)^2/2 \\ &\leq -\lambda v^\top \Sigma v - \lambda\beta^{-1} \text{Tr}(\Sigma) + \lambda^2 \kappa^4 \mathbf{E}_{\rho_v} (\theta^\top \Sigma \theta)^2/2 \\ &\leq -\lambda v^\top \Sigma v - \lambda\beta^{-1} \text{Tr}(\Sigma) + \lambda^2 \kappa^4 (12\beta^{-2} \text{Tr}(\Sigma)^2 + 4\|\Sigma\|^2), \end{aligned}$$

²This work is first to apply the variational principle to the analysis of the least singular value.

where we used the moment equivalence assumption together with (3.5) and (3.6). By Lemma 2.1 we have, with probability at least $1 - \exp(-t)$, simultaneously for all $v \in S^{d-1}$,

$$\frac{1}{n} \mathbf{E}_{\rho_v} \sum_{i=1}^n \psi(-\lambda \theta^\top M_i \theta) \leq -\lambda v^\top \Sigma v - \lambda \beta^{-1} \text{Tr}(\Sigma) + \lambda^2 \kappa^4 (12\beta^{-2} \text{Tr}(\Sigma)^2 + 4\|\Sigma\|^2) + \frac{\beta/2 + t}{n}. \tag{3.8}$$

Observe that ψ is convex. Thus, by Jensen's inequality and (3.5) we have

$$\mathbf{E}_{\rho_v} \psi(-\lambda \theta^\top M_i \theta) \geq \psi(-\lambda \mathbf{E}_{\rho_v} \theta^\top M_i \theta) = \psi(-\lambda v^\top M_i v - \lambda \beta^{-1} \text{Tr}(M_i)).$$

Since $v^\top M_i v \geq 0$ by the definition of ψ and using $\psi(x) \geq x$, we have

$$\begin{aligned} \psi(-\lambda v^\top M_i v - \lambda \beta^{-1} \text{Tr}(M_i)) &= \psi(-\lambda v^\top M_i v - \min\{\lambda \beta^{-1} \text{Tr}(M_i), 1\}) \\ &\geq -\lambda v^\top M_i v - \min\{\lambda \beta^{-1} \text{Tr}(M_i), 1\}. \end{aligned}$$

We plug this in (3.8), divide both sides by $\lambda > 0$ and obtain on the corresponding event

$$\begin{aligned} \frac{1}{n} \sum_{i=1}^n (v^\top \Sigma v - v^\top M_i v) &\leq \frac{1}{n} \sum_{i=1}^n (\min\{\lambda^{-1}, \beta^{-1} \text{Tr}(M_i)\} - \beta^{-1} \text{Tr}(\Sigma)) \\ &\quad + \lambda \kappa^4 (12\beta^{-2} \text{Tr}(\Sigma)^2 + 4\|\Sigma\|^2) + \frac{\beta/2 + t}{\lambda n}. \end{aligned}$$

It is left to analyze the sum involving $\text{Tr}(M_i)$. Observe that

$$\mathbf{E} \min\{\lambda^{-1}, \beta^{-1} \text{Tr}(M_i)\} \leq \beta^{-1} \text{Tr}(\Sigma), \text{ and } \mathbf{E} (\min\{\lambda^{-1}, \beta^{-1} \text{Tr}(M_i)\})^2 \leq \beta^{-2} \kappa^4 (\text{Tr}(\Sigma))^2,$$

where we used, denoting the standard basis in \mathbb{R}^d by e_1, \dots, e_d ,

$$\sqrt{\mathbf{E}(\text{Tr}(M_i))^2} = \sqrt{\mathbf{E} \left(\sum_{j=1}^d e_j^\top M_i e_j \right)^2} \leq \sum_{j=1}^d \sqrt{\mathbf{E} (e_j^\top M_i e_j)^2} \leq \kappa^2 \text{Tr}(\Sigma).$$

Applying the Bernstein inequality, we have, with probability at least $1 - \exp(-t)$,

$$\frac{1}{n} \sum_{i=1}^n (\min\{\lambda^{-1}, \beta^{-1} \text{Tr}(M_i)\} - \beta^{-1} \text{Tr}(\Sigma)) \leq \beta^{-1} \kappa^2 \text{Tr}(\Sigma) \sqrt{\frac{2t}{n}} + \frac{2t}{3n\lambda}.$$

By the union bound, with probability at least $1 - 2 \exp(-t)$, for any $v \in S^{d-1}$,

$$\frac{1}{n} \sum_{i=1}^n (v^\top \Sigma v - v^\top M_i v) \leq \beta^{-1} \kappa^2 \text{Tr}(\Sigma) \sqrt{\frac{2t}{n}} + \lambda \kappa^4 (12\beta^{-2} \text{Tr}(\Sigma)^2 + 4\|\Sigma\|^2) + \frac{\beta}{2\lambda n} + \frac{5t}{3n\lambda}.$$

Choose $\beta = \frac{10}{3} \mathbf{r}(\Sigma)$ and $\lambda = \sqrt{\frac{5}{3(12(10/3)^{-2} + 4)}} \frac{1}{\kappa^2 \|\Sigma\|} \sqrt{\frac{\mathbf{r}(\Sigma) + t}{n}}$. Simplifying, we conclude the proof. \square

4 Concluding remarks

A comparison with previous results by Catoni and co-authors

The application of Lemma 2.1 is not new in the context of estimation of quadratic forms. The earliest such application traces back to the works of Audibert and Catoni on robust linear least squares regression [6, 5]. More recently, it was done in the context of covariance matrix estimation [11, 17, 12] and lower tails for the sample covariance matrix [40, 38]. With a few exceptions discussed below, this technique has

not been previously applied to also control the upper tails of the sample covariance matrix. Applications to general multilinear forms were not previously analyzed.

The closest (in terms of proof techniques) to our results are the bounds appearing in [11] and in the follow-up work [17]. In particular, the work of Catoni [11] provides the dimension-dependent analysis of the sample covariance matrix under two assumptions: the equivalence of L_4 and L_2 marginal norms as well as the exponential moment assumption on the distribution of $\|X\|$ (see their Proposition 2.2 and Proposition 2.3). However, their results fall short in providing the rate of Proposition 2.3 (our simplest bound), as in their case the $\sqrt{d/n}$ convergence rate requires at least $n \geq d^5$ due to the additive term denoted there by γ_+ . The sub-optimality comes from the step of the analysis at which $\max_{i=1, \dots, n} \|X_i\|$ is used to control the moment generating function of the quadratic form. The problem is explicitly pointed out by Catoni [11, pages 15 and 16], where the author is asking if it is possible to show that γ_+ scales as d/n at least in some cases. The same limitations are inherited in the dimension-free extension of this analysis appearing in [17, Proposition 5.1 and Proposition 5.2]: In addition to unnecessary logarithmic factors, their bound similarly requires the sample size n to be greater than some integer power of $r(\Sigma)$. To resolve these problems and achieve the optimal guarantees, we introduce a truncated non-isotropic posterior distribution (2.4). This idea also plays a central role in the proof of Proposition 3.1 for which no non-isotropic analog is known.

Our second key idea is the analysis of the truncated powers of linear forms in Lemma 2.11. When $s = 2$, the bound of Lemma 2.11 shares some similarities with the uniform bound of Proposition 4.2 in [13]. However, we work with an explicit truncation applied to powers of linear forms, which allows to combine the bound of Lemma 2.11 with the classical decoupling-chaining argument analyzed in [3]. This plays a central role in the proofs of Theorem 1.7 and Theorem 1.8.

Possible further extensions

In some applications, one is interested in the following generalization of our problem: In the setup of Theorem 1.2 let T be a bounded subset of \mathbb{R}^d . Assume that we want to provide a high probability upper bound on

$$\sup_{v \in T} \left| \frac{1}{n} \sum_{i=1}^n v^\top M_i v - v^\top \Sigma v \right|.$$

Consider the case where T can be approximated by an ellipsoid; that is, assume that there is a symmetric positive definite matrix Γ such that $\frac{1}{2}\Gamma^{1/2}S^{d-1} \subseteq T \subseteq 2\Gamma^{1/2}S^{d-1}$. In this case, by considering the random matrix $\Gamma^{1/2}M\Gamma^{1/2}$ we can verify that the norm equivalence assumption (1.2) holds and the bound of Theorem 1.2 is applicable.

Theorem 1.7 allows various extensions that can be achieved using the same approach. For example, with a slight modification of the proof one can write a similar high probability upper bound on

$$\frac{1}{n} \sup_{v_1, \dots, v_s \in S^{d-1}} \sum_{i=1}^n \left(\prod_{k=1}^s \langle X_i^k, v_k \rangle - \mathbf{E} \prod_{k=1}^s \langle X_i^k, v_k \rangle \right),$$

where for any given index i the random vectors X_i^k are either sub-Gaussian or log-concave but not necessarily the same for $k = 1, \dots, s$. Our results can also be written in the same form as the bounds in [19, 3, 49] allowing isometric approximations at any scale ε . To do so, we need to take a smaller value of λ in the proof. The restriction on n will be weakened accordingly.

Another extension is a multilinear version of the lower tail bound of Proposition 3.2. When s is even, we have $\langle X, v \rangle^s \geq 0$, so that we can prove the following: With probability

at least $1 - \exp(-t)$, it holds that

$$\sup_{v \in S^{d-1}} \left(\mathbf{E} \langle X, v \rangle^s - \frac{1}{n} \sum_{i=1}^n \langle X_i, v \rangle^s \right) \leq c_s \eta^s \|\Sigma\|^{s/2} \sqrt{\frac{\mathbf{r}(\Sigma) + t}{n}},$$

under the only assumption $(\mathbf{E}|\langle y, X \rangle|^{2s})^{1/2s} \leq \eta \sqrt{y^T \Sigma y}$ for all $y \in \mathbb{R}^d$. The proof follows from Lemma 2.11 by the lower tail argument of Theorem 1.8. This complements the dimension-dependent lower tail bound of Mendelson [35, Corollary 1.8] valid for general L_s norms with $s > 2$.

In our proofs, we used that the indexing set is essentially the product of either unit spheres or ellipsoids in \mathbb{R}^d and the functions we are considering are linear with respect to the elements of the indexing set. Recently, Koltchinskii [24] studied the asymptotic properties of smooth functions of the sample covariance operators. It is likely that if the linearity is replaced by a certain smoothness assumption, our techniques are still applicable.

References

- [1] R. Adamczak. A note on the Hanson-Wright inequality for random vectors with dependencies. *Electronic Communications in Probability*, 20:1–13, 2015. MR3407216
- [2] R. Adamczak, R. Latała, A. E. Litvak, K. Oleszkiewicz, A. Pajor, and N. Tomczak-Jaegermann. A short proof of Paouris’ inequality. *Canadian Mathematical Bulletin*, 57(1):3–8, 2014. MR3150710
- [3] R. Adamczak, A. Litvak, A. Pajor, and N. Tomczak-Jaegermann. Quantitative estimates of the convergence of the empirical covariance matrix in log-concave ensembles. *Journal of the American Mathematical Society*, 23(2):535–561, 2010. MR2601042
- [4] S. Alesker. ψ_2 -estimate for the Euclidean norm on a convex body in isotropic position. In *Geometric Aspects of Functional Analysis*, pages 1–4. Springer, 1995. MR1353444
- [5] J.-Y. Audibert and O. Catoni. Linear regression through PAC-Bayesian truncation. *arXiv preprint arXiv:1010.0072*, 2010.
- [6] J.-Y. Audibert and O. Catoni. Robust linear least squares regression. *The Annals of Statistics*, 39(5):2766–2794, 2011. MR2906886
- [7] A. S. Bandeira, M. T. Boedihardjo, and R. van Handel. Matrix concentration inequalities and free probability. *arXiv preprint arXiv:2108.06312*, 2021. MR4635836
- [8] S. Boucheron, G. Lugosi, and P. Massart. *Concentration Inequalities: A Nonasymptotic Theory of Independence*. Oxford University Press, 2013. MR3185193
- [9] S. Bubeck, Y. Li, and D. Nagaraj. A law of robustness for two-layers neural networks. *Conference on Learning Theory*, 2021.
- [10] O. Catoni. Challenging the empirical mean and empirical variance: a deviation study. *Annales de l’Institut Henri Poincaré, Probabilités et Statistiques*, 48(4):1148–1185, 2012. MR3052407
- [11] O. Catoni. PAC-Bayesian bounds for the Gram matrix and least squares regression with a random design. *arXiv preprint arXiv:1603.05229*, 2016.
- [12] O. Catoni and I. Giulini. Dimension-free PAC-Bayesian bounds for matrices, vectors, and linear least squares regression. *arXiv preprint arXiv:1712.02747*, 2017.
- [13] O. Catoni and I. Giulini. Dimension-free PAC-Bayesian bounds for the estimation of the mean of a random vector. *arXiv preprint arXiv:1802.04308*, 2018.
- [14] G. Chinot, G. Lecué, and M. Lerasle. Robust statistical learning with Lipschitz and convex loss functions. *Probability Theory and Related Fields*, pages 1–44, 2019. MR4087486
- [15] M. Even and L. Massoulié. Concentration of non-isotropic random tensors with applications to learning and empirical risk minimization. In *Conference on Learning Theory*, 2021.
- [16] A. A. Giannopoulos and V. D. Milman. Concentration property on probability spaces. *Advances in Mathematics*, 156(1):77–106, 2000. MR1800254

- [17] I. Giullini. Robust dimension-free Gram operator estimates. *Bernoulli*, 24(4B):3864–3923, 2018. MR3788191
- [18] F. Götze, H. Sambale, and A. Sinulis. Concentration inequalities for polynomials in α -sub-exponential random variables. *Electronic Journal of Probability*, 26:1–22, 2021. MR4247973
- [19] O. Guédon and M. Rudelson. L_p -moments of random vectors via majorizing measures. *Advances in Mathematics*, 208(2):798–823, 2007. MR2304336
- [20] D. Hsu, S. Kakade, and T. Zhang. A tail inequality for quadratic forms of subgaussian random vectors. *Electronic Communications in Probability*, 17:1–6, 2012. MR2994877
- [21] H. Jeong, X. Li, Y. Plan, and Ö. Yilmaz. Sub-Gaussian matrices on sets: Optimal tail dependence and applications. *Communications on Pure and Applied Mathematics (forthcoming)*, 2021. MR4465901
- [22] Y. Klochkov, A. Kroshnin, and N. Zhivotovskiy. Robust k -means clustering for distributions with two moments. *The Annals of Statistics (forthcoming)*, 2020. MR4319247
- [23] Y. Klochkov and N. Zhivotovskiy. Uniform Hanson-Wright type concentration inequalities for unbounded entries via the entropy method. *Electronic Journal of Probability*, 25, 2020. MR4073683
- [24] V. Koltchinskii. Asymptotically efficient estimation of smooth functionals of covariance operators. *Journal of the European Mathematical Society*, 23(3):765–843, 2020. MR4210724
- [25] V. Koltchinskii and K. Lounici. Concentration inequalities and moment bounds for sample covariance operators. *Bernoulli*, 23(1):110–133, 2017. MR3556768
- [26] V. Koltchinskii and S. Mendelson. Bounding the smallest singular value of a random matrix without concentration. *International Mathematics Research Notices*, 2015(23):12991–13008, 2015. MR3431642
- [27] R. Latała. On some problems concerning log-concave random vectors. In *Convexity and Concentration*, pages 525–539. Springer, 2017. MR3837281
- [28] R. Latała and P. Nayar. Hadamard products and moments of random vectors. *Advances in Mathematics*, 375:107–114, 2020. MR4170230
- [29] M. Ledoux. *The Concentration of Measure Phenomenon*, volume 89 of *Mathematical surveys and Monographs*. American Mathematical Society, 2001. MR1849347
- [30] C. Liaw, A. Mehrabian, Y. Plan, and R. Vershynin. A simple tool for bounding the deviation of random matrices on geometric sets. In *Geometric Aspects of Functional Analysis*, pages 277–299. Springer, 2017. MR3645128
- [31] M. E. Lopes, N. B. Erichson, and M. W. Mahoney. Bootstrapping the operator norm in high dimensions: Error estimation for covariance matrices and sketching. *arXiv preprint arXiv:1909.06120*, 2019. MR4497253
- [32] G. Lugosi and S. Mendelson. Mean estimation and regression under heavy-tailed distributions: A survey. *Foundations of Computational Mathematics*, 19(5):1145–1190, 2019. MR4017683
- [33] G. Lugosi and S. Mendelson. Sub-Gaussian estimators of the mean of a random vector. *The Annals of Statistics*, 47(2):783–794, 2019. MR3909950
- [34] S. Mendelson. On weakly bounded empirical processes. *Mathematische Annalen*, 340(2):293–314, 2008. MR2368981
- [35] S. Mendelson. Approximating L_p unit balls via random sampling. *Advances in Mathematics*, 386:1–20, 2021. MR4270522
- [36] S. Mendelson and N. Zhivotovskiy. Robust covariance estimation under $L_4 - L_2$ norm equivalence. *The Annals of Statistics*, 48(3):1648–1664, 2020. MR4124338
- [37] S. Minsker. Geometric median and robust estimation in Banach spaces. *Bernoulli*, 21(4):2308–2335, 2015. MR3378468
- [38] J. Mourtada. Exact minimax risk for linear least squares, and the lower tail of sample covariance matrices. *Annals of Statistics (forthcoming)*, 2022. MR4474486
- [39] A. Nemirovski. Interior point polynomial time methods in convex programming. *Lecture Notes*, 42(16):3215–3224, 2004.

- [40] R. Oliveira. The lower tail of random quadratic forms with applications to ordinary least squares. *Probability Theory and Related Fields*, 166(3-4):1175–1194, 2016. MR3568047
- [41] D. M. Ostrovskii and A. Rudi. Affine invariant covariance estimation for heavy-tailed distributions. In *Conference on Learning Theory*, pages 2531–2550. PMLR, 2019.
- [42] G. Paouris. Concentration of mass on convex bodies. *Geometric & Functional Analysis*, 16(5):1021–1049, 2006. MR2276533
- [43] H. Sambale. Some notes on concentration for α -subexponential random variables. *arXiv preprint arXiv:2002.10761*, 2020. MR4633263
- [44] N. Srivastava and R. Vershynin. Covariance estimation for distributions with $2 + \varepsilon$ moments. *The Annals of Probability*, 41(5):3081–3111, 2013. MR3127875
- [45] M. Talagrand. *Upper and Lower Bounds for Stochastic Processes: Modern Methods and Classical Problems*, volume 60. Springer Science & Business Media, 2014. MR3184689
- [46] K. E. Tikhomirov. The smallest singular value of random rectangular matrices with no moment assumptions on entries. *Israel Journal of Mathematics*, 212(1):289–314, 2016. MR3504328
- [47] J. A. Tropp. User-friendly tail bounds for sums of random matrices. *Foundations of Computational Mathematics*, 12(4):389–434, 2012. MR2946459
- [48] R. Van Handel. Structured random matrices. In *Convexity and Concentration*, pages 107–156. Springer, 2017. MR3837269
- [49] R. Vershynin. Approximating the moments of marginals of high-dimensional distributions. *The Annals of Probability*, 39(4):1591–1606, 2011. MR2857251
- [50] R. Vershynin. *High-Dimensional Probability: An Introduction with Applications*, volume 47 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, 2016. MR3837109
- [51] R. Vershynin. Concentration inequalities for random tensors. *Bernoulli*, 26(4):3139–3162, 2020. MR4140540
- [52] P. Yaskov. Lower bounds on the smallest eigenvalue of a sample covariance matrix. *Electronic Communications in Probability*, 19, 2014. MR3291620
- [53] P. Youssef. Estimating the covariance of random matrices. *Electronic Journal of Probability*, 18:1–26, 2013. MR3151727

Acknowledgments. The author would like to thank Pedro Abdalla, Radosław Adamczak and Afonso Bandeira for a valuable feedback and Jaouad Mourtada for many insightful discussions on the topic.