



Generic error bounds for the generalized Lasso with sub-exponential data

Martin Genzel¹  · Christian Kipp²

Received: 17 September 2020 / Accepted: 15 July 2022 / Published online: 12 August 2022
© The Author(s) 2022

Abstract

This work performs a non-asymptotic analysis of the generalized Lasso under the assumption of sub-exponential data. Our main results continue recent research on the benchmark case of (sub-)Gaussian sample distributions and thereby explore what conclusions are still valid when going beyond. While many statistical features remain unaffected (e.g., consistency and error decay rates), the key difference becomes manifested in how the complexity of the hypothesis set is measured. It turns out that the estimation error can be controlled by means of two complexity parameters that arise naturally from a generic-chaining-based proof strategy. The output model can be non-realizable, while the only requirement for the input vector is a generic concentration inequality of Bernstein-type, which can be implemented for a variety of sub-exponential distributions. This abstract approach allows us to reproduce, unify, and extend previously known guarantees for the generalized Lasso. In particular, we present applications to semi-parametric output models and phase retrieval via the lifted Lasso. Moreover, our findings are discussed in the context of sparse recovery and high-dimensional estimation problems.

Keywords Generalized Lasso · High-dimensional signal estimation · Statistical learning · Sub-exponential data · Generic chaining

Mathematics Subject Classification 60D05 · 62F30 · 62F35 · 90C25

Communicated by Felix Kraemer.

Martin Genzel: work done while at Technische Universität Berlin and Utrecht University.

✉ Martin Genzel
martin.genzel@helmholtz-berlin.de

Christian Kipp
kipp@math.tu-berlin.de

¹ Helmholtz-Zentrum Berlin für Materialien und Energie, Berlin, Germany

² Department of Mathematics, Technische Universität Berlin, Berlin, Germany

1 Introduction

This paper is concerned with the following common inference problem in statistical learning: Let $(x_1, y_1), \dots, (x_n, y_n) \in \mathbb{R}^p \times \mathbb{R}$ be samples of a random input-output pair $(x, y) \in \mathbb{R}^p \times \mathbb{R}$, whose joint probability distribution is unknown. What information about the relationship between x and y can we retrieve only based on the knowledge of $(x_1, y_1), \dots, (x_n, y_n)$?

A classical instance of this problem is *linear regression*, where y depends linearly on x , say $y = \langle x, \beta_0 \rangle + v$ for an unknown parameter vector $\beta_0 \in \mathbb{R}^p$ and independent additive noise v . While the resulting task of estimating β_0 is nowadays fairly well understood in the low-dimensional regime $n \geq p$, it is still subject of ongoing research in the high-dimensional regime $n \ll p$. In the latter scenario, it is indispensable to impose additional conditions on the input-output model. A typical assumption is that β_0 belongs to a known, convex *hypothesis set* $K \subset \mathbb{R}^p$ that is of low complexity in a certain sense. In such a model setup, a natural estimation procedure is based on solving the *generalized Lasso*:¹

$$\min_{\beta \in K} \frac{1}{n} \sum_{i=1}^n (y_i - \langle x_i, \beta \rangle)^2. \quad (\text{LS}_K)$$

The popularity of Lasso-type estimators is due to several desirable properties. Perhaps most importantly, many efficient algorithmic implementations are available for (LS_K) due to the convexity of K (e.g., see [12, 56, 66]), accompanied by the suitability for a statistical analysis due to its simple variational formulation (e.g., see the textbooks [6, 13, 23]). A more astonishing feature of the generalized Lasso (LS_K) is its ability to deal with *non-linear* relations between x and y . In fact, inspired by a classical result of Brillinger [5], a recent work of Plan and Vershynin [42] shows that for Gaussian input vectors, (LS_K) yields a consistent estimator for *single-index models*, i.e., $y = f(\langle x, \beta_0 \rangle)$ with an unknown, non-linear distortion function $f : \mathbb{R} \rightarrow \mathbb{R}$. This finding has triggered a lot of related and follow-up research, e.g., see [14, 16, 19, 38, 46, 52–54]. We note that these works form only a small fraction of a whole research area on non-linear observation models, lying at the interface of statistics, learning theory, signal processing, and compressed sensing. A comprehensive review of the literature goes beyond the scope of this paper, and we refer the reader to [15, Sec. 4.2] and the references therein for more details in that regard.

The present work is inspired by the general framework developed in [15], which enables a theoretical analysis of (LS_K) for a large class of *semi-parametric observation models* (see also the technical report [17]). More specifically, we will not assume an explicit functional relationship between x and y , such as for single-index models; note that a similar viewpoint is taken by Sattar and Oymak [46], who refrain from

¹ We adopt the common term ‘generalized Lasso’ from the literature (e.g., see [42]), as a tribute to the original *Lasso* estimator introduced by Tibshirani [55], where the hypothesis set corresponds to a scaled ℓ^1 -ball, serving as a convex relaxation of *sparse* parameter vectors. Taking the viewpoint of statistical learning, (LS_K) is a specific instance of *(constraint) empirical risk minimization*, but this terminology appears somewhat too general for the purpose of this paper.

a *realizable model* connecting the input and output. Adopting this abstract setup, we intend to address the following parameter estimation problem:

Problem 1.1 *Under what conditions is the generalized Lasso (LS_K) capable of estimating a certain target vector $\beta^\natural \in K$ that carries information about the relationship between x and y ? What is the impact of the sample size n and the complexity of the (convex) hypothesis set $K \subset \mathbb{R}^p$?*

For the moment, it is convenient to assume that the vector which provides the desired ‘information’ is an *expected risk minimizer*, i.e., we set $\beta^\natural := \beta^*$ where $\beta^* \in K$ is a solution to the expected risk minimization problem (on K):

$$\min_{\beta \in K} \mathbb{E}[(y - \langle x, \beta \rangle)^2]. \quad (1.1)$$

Indeed, this simplification reduces the above problem to a well-known challenge in statistical learning theory, namely finding the best possible (linear) predictor of y by empirical risk minimization.² However, we wish to emphasize that the absolute magnitude of the prediction error is only of minor importance to our approach, since the predictive capacity of the Lasso is likely to be poor unless y follows a linear model. On the other hand, one can still hope for a satisfactory outcome in the sense of Problem 1.1, which explains why all guarantees presented in this paper concern the (parameter) estimation error. More details on relevant scenarios, where β^\natural is not necessarily equal to the expected risk minimizer, are discussed later in the context of semi-parametric (non-linear) output models; see Sect. 3.1.

The first author’s dissertation [15] gives a far-reaching answer to Problem 1.1 for (*sub*-)Gaussian input vectors—an assumption that is made in almost all of the above-mentioned works on the generalized Lasso and related estimators. However, the setup of [15] does not adequately address output models with even non-linearities, most prominently, the *phase retrieval problem* where $y = |\langle x, \beta_0 \rangle|$.³ The initial motivation of the present article was to address this fundamental shortcoming by applying the “phase lift trick” [7, 8] to the Lasso. A key feature of the lifted Lasso is that the input vector $x \in \mathbb{R}^p$ is replaced by the tensor product xx^T , see Sect. 3.5 for details. If x is an isotropic sub-Gaussian random vector, then xx^T is a (generally anisotropic) *sub-exponential* random matrix satisfying a mixed-tail inequality of Bernstein type (whose precise form depends on the exact assumptions on x). The strategy we follow here is tailored to this problem, but abstracts away from all irrelevant details. Therefore, going clearly beyond non-asymptotic error bounds for phase-retrieval-like models, we obtain a general solution to Problem 1.1 if the sample data have sub-exponential tails. Indeed, research on this subject is still in its infancy, and we intend to make progress by presenting a unified approach. This particularly includes the derivation of new and extension of known results on statistical estimation with sub-exponential

² To simplify the presentation even further, one might simply assume a noisy linear model $y = \langle x, \beta_0 \rangle + v$. In this case, the key messages of our main results remain valid; in particular, our conclusions on heavier tailed input data are also of interest to high-dimensional linear regression.

³ More specifically, if x is standard Gaussian, then (LS_K) becomes a useless estimator in this case, since the expected risk minimizer β^* in (1.1) is just the zero vector.

data (see Sect. 3). Before outlining our general approach in Sect. 1.2, we would like to familiarize the reader with the setup of this paper by presenting a prototypical estimation guarantee for sub-exponential input vectors.

1.1 A simple error bound for sub-exponential input vectors

Let us begin with the formal definition of sub-Gaussian and sub-exponential random variables:

Definition 1.2 (*Sub-Gaussian/sub-exponential random variables*) For $\alpha \in \{1, 2\}$, we define the *exponential Orlicz norm* of a random variable $Z : \Omega \rightarrow \mathbb{R}$ by⁴

$$\|Z\|_{\psi_\alpha} := \inf \left\{ t > 0 \mid \mathbb{E} \left[\exp \left(\frac{|Z|^\alpha}{t^\alpha} \right) \right] \leq 2 \right\}.$$

The *exponential Orlicz space* L_{ψ_α} is then denoted by

$$L_{\psi_\alpha} := \{Z : \Omega \rightarrow \mathbb{R} \mid \|Z\|_{\psi_\alpha} < \infty\}.$$

The elements of the exponential Orlicz spaces L_{ψ_1} and L_{ψ_2} are called *sub-exponential* and *sub-Gaussian random variables*, respectively.

The notions of sub-exponentiality and sub-Gaussianity impose restrictions on the tails of a random variable, which must not be “too heavy”. This intuition gives rise to several equivalent versions of Definition 1.2, which are summarized in Proposition A.1 in Appendix A; for a more detailed introduction, we refer to [58, Chap. 2 & 3].

The definition of sub-Gaussian and sub-exponential random vectors is characterized by their one-dimensional marginals (i.e., projections onto one-dimensional subspaces):

Definition 1.3 (*Sub-Gaussian/sub-exponential random vectors*) For a random vector $x \in \mathbb{R}^p$ and $\alpha \in \{1, 2\}$, we set

$$\|x\|_{\psi_\alpha} := \sup_{v \in \mathbb{S}^{p-1}} \|\langle x, v \rangle\|_{\psi_\alpha}.$$

If $\|x\|_{\psi_2} < \infty$, we say that x is (*uniformly*) *sub-Gaussian*, and if $\|x\|_{\psi_1} < \infty$, we say that x is (*uniformly*) *sub-exponential*.

The following result states a non-asymptotic error bound for the generalized Lasso (LS_K) with sub-exponential input vectors. Its proof is provided in Sect. 5.6, being a “by-product” of one of our main results, Corollary 2.15 in Sect. 2.3. For the sake of simplicity, we restrict ourselves to a polytopal hypothesis set K here, as this allows for explicit bounds on the complexity parameters. Moreover, it is worth emphasizing that for linear models, i.e., if $y = \langle x, \beta_0 \rangle$, we simply obtain an estimation guarantee for $\beta^* = \beta_0$.

⁴ We do not explicitly mention the underlying probability space here. In fact, our analysis does not require any treatment of measure theoretic issues and we simply assume that the probability space is rich enough to model all random quantities and processes that we are interested in.

Proposition 1.4 *Let $(x, y) \in \mathbb{R}^p \times \mathbb{R}$ be a joint random pair such that $y \in \mathbb{R}$ is sub-exponential and $x \in \mathbb{R}^p$ is isotropic and sub-exponential with $\|x\|_{\psi_1} \leq \kappa$ for some $\kappa > 0$. Let $K \subset \mathbb{R}^p$ be a convex polytope with D vertices and Euclidean diameter $\Delta_2(K)$, and let $\beta^* \in K$ be the expected risk minimizer on K , i.e., a solution to (1.1). Finally, let the observed sample pairs $(x_1, y_1), \dots, (x_n, y_n) \in \mathbb{R}^p \times \mathbb{R}$ be independent copies of (x, y) . Then there exists a universal constant $C > 0$ such that for every $u \geq 8$, the following holds true with probability at least $1 - 5 \exp(-C \cdot u^2) - 2 \exp(-C \cdot \sqrt{n})$: If the sample size obeys*

$$n \gtrsim \left(\kappa^{10} \cdot \Delta_2(K) \cdot \left(\frac{\log(D)}{\sqrt{n}} + \sqrt{\log(D)} \right) + \kappa^6 \cdot u \right)^2,$$

then every minimizer $\hat{\beta}$ of (LS_K) satisfies

$$\|\hat{\beta} - \beta^*\|_2 \lesssim \kappa^{18} \cdot \max\{1, u^2 \cdot \sigma(\beta^*)\} \cdot \frac{\sqrt{\kappa \cdot \Delta_2(K) \cdot \log(D)}}{n^{1/4}},$$

where $\sigma(\beta^*) := \|y - \langle x, \beta^* \rangle\|_{\psi_1}$.

Informally speaking, Proposition 1.4 shows that estimation of the expected risk minimizer succeeds with overwhelmingly high probability as long as $n \gg \Delta_2(K)^2 \cdot \log(D)^2$. Such a statement is particularly appealing to high-dimensional problems such as sparse recovery. Another remarkable conclusion is that the estimator (LS_K) essentially performs as well as if the sample data were sub-Gaussian (cf. [15, Thm. 4.3]). Our main results in Sect. 2 confirm this observation in much greater generality, but they will also reveal several important differences to the sub-Gaussian case; first and foremost, we will be concerned with defining appropriate complexity measures for K , which do not explicitly appear in the polytopal setting of Proposition 1.4. In this respect, it is important to note that there are relevant special cases of sub-exponential vectors (e.g., those with independent coordinates) for which the above estimate is too pessimistic and can be improved. The elaboration of this aspect is a key concern of this article and motivates the introduction of a generic tail condition that takes the underlying “geometry” of the problem into account. Apart from this, let us also emphasize that the simplifications of Proposition 1.4 come along with a suboptimal behavior regarding, (a), the error decay rate $O(n^{-1/4})$, (b), the sub-exponential parameter κ , and (c), the model deviation parameter $\sigma(\beta^*)$.

To the best of our knowledge, Proposition 1.4 is a new result, but it bears resemblance with a recent finding of Sattar and Oymak [46, Thm. 3.4], who consider a similar model setup with sub-exponential input vectors. Their analysis focuses on the projected gradient descent method, as an algorithmic implementation of (LS_K) , and is therefore related to our estimation guarantees; see Sect. 3.4 for a more detailed comparison.

1.2 Contributions and overview

The main purpose of this work is to shed more light on the estimation capacity of the generalized Lasso (LS_K) when the sample data are not sub-Gaussian. While Proposition 1.4 already gives a first glimpse into the prototypical situation of sub-exponential input vectors, we intend to address this problem in a more systematic and abstract way (cf. Problem 1.1). At the heart of our statistical analysis stands the so-called *generic Bernstein concentration*, which is introduced in Sect. 2.1 (see Definition 2.2). This concept is the outcome of a somewhat uncommon proof strategy: Instead of assuming a specific (sub-exponential) distribution for x , we study the associated excess risk of (LS_K) in an abstract sense, relying on an advanced *generic chaining* argument due to Mendelson [32]. Consequently, the key step of our approach is to understand the increment behavior of the underlying stochastic processes, and in fact, this precisely leads to generic Bernstein concentration as a natural condition for x . In that way, we are able to explore (LS_K) for a whole class of input distributions and thereby to refine the assumption of uniform sub-exponentiality in Proposition 1.4. Another important outcome of our analysis are two general complexity parameters for the hypothesis set K (see Definition 2.5 and 2.6), which are compatible with the notion of generic Bernstein concentration.

With these preliminaries at hand, we formulate our main result in Sect. 2.2 (see Theorem 2.10), which provides a novel, non-asymptotic error bound for (LS_K) under generic Bernstein concentration. However, a direct application of this guarantee to specific model situations is not always straightforward, since the aforementioned complexity parameters are of local nature, implicitly depending on the desired precision level. For this reason, we present two more easily accessible corollaries of Theorem 2.10 in Sect. 2.3. These results are based on simplified complexity parameters (see Definition 2.11 and 2.13, respectively), but come with the price of looser error bounds and sample-size conditions.

While the purpose of Sect. 2 is to develop a unified analysis for the generalized Lasso (LS_K), Sect. 3 is devoted to various applications and examples of our findings. We begin with a brief discussion on semi-parametric modeling in Sect. 3.1, demonstrating how our general results may be applied to specific parameter estimation problems. This is followed by several relevant examples of generic Bernstein concentration (see Sects. 3.2–3.4), leading to off-the-shelf guarantees for (LS_K) with sub-exponential and sub-Gaussian sample data; these parts also provide a comparison to related approaches in the literature. In Sect. 3.5, we then revisit our motivating case study on the *lifted Lasso* for phase-retrieval-like problems—a scenario where sub-exponential distributions arise naturally. Finally, Sect. 3.6 contains a more detailed discussion of the complexity parameters from Sect. 2. In this context, it will become clearer that measuring complexity beyond sub-Gaussianity is a delicate issue and comes along with unexplored difficulties. Nevertheless, we are able to establish simple bounds in the prototypical case of ℓ^1 -balls, making our error bounds applicable to high-dimensional estimation and sparse recovery. Some concluding remarks are made in Sect. 4.

1.3 Differentiation from previous works

Apart from enabling heavier-tailed data, a crucial feature of generic Bernstein concentration is that it does *not require any type of isotropy*. Instead, the “geometric” behavior of the input vectors is captured by selecting two appropriate semi-norms (see Definition 2.2). This relaxation is key to the applicability of our results, as it allows us to handle structured, but anisotropic input vectors, such as arising in the phase lift approach (see Sect. 3.5). It is not clear to us how this important challenge could be addressed with other techniques, especially those suggested in [15] and related articles mentioned above.⁵ The concept of generic Bernstein concentration therefore also presents a novel and systematic solution to this open problem, which is arguably one of the most significant achievements compared to previous works.

We close this part with another clarification: The present article is concerned with the generalized Lasso (LS_K) when the sample data are heavier tailed than sub-Gaussian, in particular, the underlying distribution may be *unbounded*. An alternative strategy is to first truncate the raw data at an appropriate threshold and then to apply (LS_K) or a similar estimator. In fact, the latter approach is quite common in practice, but it also facilitates a theoretical study due to the boundedness of the involved random variables, e.g., see [19, 60, 62–64] for related results on non-linear observation models. However, the (concentration-based) machinery for bounded sample data is certainly not applicable to the model setup of the present paper. Instead, we rather follow the conceptual ideas of Mendelson [31, 33], who points out the downsides of the bounded framework and develops a general theory for heavy-tailed problems. Although our analysis is concerned with more specific model assumptions, it is just general enough to allow for a rigorous understanding of estimation with sub-exponential data, thereby unifying and improving a series of previously known results from the literature (see Sect. 3). Thus, to a certain degree, our work can be seen as a “connecting piece” between such highly customized approaches and the abstract theory of Mendelson [31, 33].

1.4 Notation

The letter C is reserved for constants, whose values could change from time to time, and we say that C is *universal* if its value does not depend on any other involved parameter. If an inequality holds true up to a universal constant $C > 0$, we usually write $A \lesssim B$ instead of $A \leq C \cdot B$; the notation $A \asymp B$ means that both $A \lesssim B$ and $B \lesssim A$ hold true. Furthermore, the *positive part* of a real number $s \in \mathbb{R}$ is denoted by $[s]_+ := \max\{s, 0\}$.

The *cardinality* of a finite set I is denoted by $|I|$. The j -th entry of a vector $v \in \mathbb{R}^p$ is denoted by v_j and the *support* of v is defined as $\text{supp}(v) := \{j \mid v_j \neq 0\}$. The cardinality of $\text{supp}(v)$ is referred to as the *sparsity* of v and we write $\|v\|_0 := |\text{supp}(v)|$. For $1 \leq q \leq \infty$, we denote the ℓ^q -norm on \mathbb{R}^p by $\|\cdot\|_q$ and the associated *unit ball* by B_q^p . The *Euclidean unit sphere* is given by $\mathbb{S}^{p-1} := \{v \in \mathbb{R}^p \mid \|v\|_2 = 1\}$.

⁵ Unfortunately, the idea of [15] to capture anisotropic and correlated structures through *mixing matrices* is no remedy in our situations of interest, e.g., phase-lifted input vectors.

The *Frobenius norm* is denoted by $\|\cdot\|_F$ and the *spectral norm* by $\|\cdot\|_{\text{op}}$. We write $I_p \in \mathbb{R}^{p \times p}$ for the *identity matrix*.

Let $L \subset \mathbb{R}^p$. By $\text{span}(L)$, $\text{cone}(L)$, and $\text{conv}(L)$, we denote the *linear hull*, *conic hull*, and *convex hull*, respectively. The *diameter* of L with respect to a (pseudo-)metric d is defined as $\Delta(L) := \sup_{v_1, v_2 \in L} d(v_1, v_2)$.

Let x be a random vector in \mathbb{R}^p . We say that x is *centered* if $\mathbb{E}[x] = 0$, it is *symmetric* if x has the same distribution as $-x$, and it is *isotropic* if $\mathbb{E}[xx^T] = I_p$. The L^q -norm of a real-valued random variable Z is $\|Z\|_{L^q} := (\mathbb{E}[|Z|^q])^{1/q}$ for $1 \leq q < \infty$. Moreover, we write $g \sim \mathcal{N}(0, I_p)$ if g is a *standard Gaussian random vector* in \mathbb{R}^p .

For $v \in \mathbb{R}^p$, we use the notation v^* for the linear functional $\langle \cdot, v \rangle$, i.e., v^* is the image of v under the Riesz isomorphism; analogously, we write A^* for the image of a subset $A \subset \mathbb{R}^p$ under the Riesz isomorphism. Furthermore, if \mathbb{R}^p is equipped with a probability measure μ , we can interpret v^* as a random variable, i.e., $v^* = \langle x, v \rangle$, where x is distributed according to μ . In particular, we have that $\|v^*\|_{L^q}^q = \mathbb{E}[|\langle v, x \rangle|^q]$ for $1 \leq q < \infty$.

2 Main results

This section presents the main results of this work. We begin with several technical preliminaries in Sect. 2.1, including the central concept of generic Bernstein concentration (see Definition 2.2) as well as the related complexity parameters (see Definition 2.5 and 2.6). The most general estimation guarantee is then formulated and discussed in Sect. 2.2 (see Theorem 2.10). This is followed by two corollaries in Sect. 2.3, employing simplified variants of our complexity parameters. Note that all proofs for this section are postponed to Sect. 5.

2.1 Preliminaries and generic Bernstein concentration

An error bound for the generalized Lasso (LS_K) is a statement about the minimizer of the following function:

Definition 2.1 (*Empirical risk, excess risk*) The objective function minimized in (LS_K), i.e.,

$$\bar{\mathcal{L}}(\beta) := \frac{1}{n} \sum_{i=1}^n (y_i - \langle x_i, \beta \rangle)^2,$$

is called the *empirical risk* of $\beta \in K$. Given $\beta, \beta^\natural \in K$, we call

$$\mathcal{E}(\beta, \beta^\natural) := \bar{\mathcal{L}}(\beta) - \bar{\mathcal{L}}(\beta^\natural)$$

the *excess risk* of β over β^\natural .

Since the map $\beta \mapsto \bar{\mathcal{L}}(\beta)$ depends on the random pairs (x_i, y_i) , it can be seen as a stochastic process on the hypothesis set K . If the excess risk is strictly positive on a subset of K , the minimizer must be outside of this subset. In other words, we can

localize the empirical risk minimizer in a certain set $L \subset K$ if we have a positive lower bound for the excess risk on $K \setminus L$ (see Fact 2.4 below). A powerful technique for proving such lower bounds is *generic chaining* for stochastic processes (see [50, 51]).

The following definition introduces a generic concentration inequality for linear functions on the parameter space, which leads to an increment condition for the involved stochastic processes. Based on this condition, we will use chaining arguments to derive a generic error bound for (LS_K) (see Theorem 2.10 and its proof in Sect. 5.2). Estimation guarantees for specific classes of input vectors can be then obtained by considering concrete instances of this condition (see Sect. 3).

Definition 2.2 (*Generic Bernstein concentration*) Let $x \in \mathbb{R}^p$ be a random vector and let $\|\cdot\|_g$ and $\|\cdot\|_e$ be two semi-norms on \mathbb{R}^p . We say that x exhibits *generic Bernstein concentration* with respect to $(\|\cdot\|_g, \|\cdot\|_e)$ if for every $v \in \mathbb{R}^p$ and every $t \geq 0$, we have that

$$\mathbb{P}(|\langle x, v \rangle| \geq t) \leq 2 \exp\left(-\min\left\{\frac{t^2}{\|v\|_g^2}, \frac{t}{\|v\|_e}\right\}\right), \tag{2.1}$$

where $\exp(-\infty) := 0$ and

$$\frac{t}{0} := \begin{cases} \infty & \text{for } t > 0, \\ 0 & \text{for } t = 0. \end{cases}$$

The prototypical instance of generic Bernstein concentration is a centered random vector $x = (x_1, \dots, x_p) \in \mathbb{R}^p$ with independent, sub-exponential coordinates: indeed, such an x exhibits generic Bernstein concentration with respect to $(\frac{R}{\sqrt{C_B}}\|\cdot\|_2, \frac{R}{C_B}\|\cdot\|_\infty)$ where $R := \max_{1 \leq j \leq p} \|x_j\|_{\psi_1}$ and a universal constant $C_B > 0$. In this case, (2.1) simply corresponds to the classical Bernstein’s inequality (see Theorem A.3), justifying the terminology of Definition 2.2. More generally, (2.1) can be seen as an example of *mixed-tail conditions*, which are quite common in the generic chaining literature, e.g., see [11, Thm. 3.5] or [51, Thm. 2.2.23]. To be more specific, the semi-norm $\|\cdot\|_g$ governs the Gaussian-like (‘g’) tail, while $\|\cdot\|_e$ governs the exponential-like (‘e’) tail.

The central idea of generic chaining is that the expected infimum (or supremum) of a stochastic process depends on the “size” of the underlying index set, which is equipped with a (pseudo-)metric that reflects the increment behavior of the stochastic process. For certain classes of canonical processes, the appropriate way of measuring the size is given by the well-known γ -functional:

Definition 2.3 (γ -functional; [51, Def. 2.2.19]) Let L be a set equipped with a pseudo-metric d . We call a sequence $(\mathcal{A}_s)_{s \in \mathbb{N}}$ of partitions⁶ of L an *admissible partition sequence* if $|\mathcal{A}_0| = 1$ and $|\mathcal{A}_s| \leq 2^{2^s}$ for $s \geq 1$ and if the sequence is increasing, i.e., for every $A \in \mathcal{A}_{s+1}$ there is some $B \in \mathcal{A}_s$ with $A \subset B$. For $\alpha \in \{1, 2\}$, we set

$$\gamma_\alpha(L, d) := \inf_{v \in L} \sup_{s \in \mathbb{N}} \sum_{s \in \mathbb{N}} 2^{s/\alpha} \Delta(\mathcal{A}_s(v)),$$

⁶ As usual, by a partition of L , we mean a family of pairwise disjoint, non-empty subsets of L whose union is L .

where $A_s(v)$ is the unique set in \mathcal{A}_s containing v and the infimum is taken over all admissible partition sequences. In this work, we will only deal with pseudo-metrics induced by semi-norms. Hence, we may write $\gamma_\alpha(L, \|\cdot\|) := \gamma_\alpha(L, d_{\|\cdot\|})$ where $d_{\|\cdot\|}$ is the pseudo-metric induced by a semi-norm $\|\cdot\|$.

Returning to the issue of finding an error bound for (\mathbf{LS}_K) , let us now fix some precision level $t > 0$ and an arbitrary *target vector* $\beta^\natural \in K$ (see also Problem 1.1). At the present level of abstraction, it is beneficial to leave the notion of the ‘estimation error’ as general as possible. For the sake of mental convenience, β^\natural can be seen as a desirable outcome of an estimation procedure (e.g., the expected risk minimizer on K), but this interpretation is mathematically irrelevant. The error measure that will concern us in this section is the Euclidean distance $\|\hat{\beta} - \beta^\natural\|_2$, where $\hat{\beta} \in K$ is the estimate of the generalized Lasso, i.e., a minimizer of (\mathbf{LS}_K) . Since $\mathcal{E}(\cdot, \beta^\natural)$ is a convex function (on K), one can make use of the following basic, yet important fact:

Fact 2.4 *Let $K \subset \mathbb{R}^p$ be a convex set. For $\beta^\natural \in K$ and $t > 0$, we set*

$$K_{\beta^\natural, t} := \{\beta \in K \mid \|\beta - \beta^\natural\|_2 = t\} = K \cap (t\mathbb{S}^{p-1} + \beta^\natural).$$

If $\mathcal{E}(\beta, \beta^\natural) > 0$ for all $\beta \in K_{\beta^\natural, t}$, then every minimizer $\hat{\beta}$ of (\mathbf{LS}_K) satisfies the error bound $\|\hat{\beta} - \beta^\natural\|_2 < t$.

Consequently, it suffices to control $\mathcal{E}(\cdot, \beta^\natural)$ on the spherical subset $K_{\beta^\natural, t}$ of radius t around β^\natural . To this end, we loosely follow the approach of Mendelson [31] and decompose the excess risk as follows:

$$\begin{aligned} \mathcal{E}(\beta, \beta^\natural) &= \frac{1}{n} \sum_{i=1}^n (y_i - \langle x_i, \beta \rangle)^2 - \frac{1}{n} \sum_{i=1}^n (y_i - \langle x_i, \beta^\natural \rangle)^2 \\ &= \underbrace{\frac{1}{n} \sum_{i=1}^n \langle x_i, \beta - \beta^\natural \rangle^2}_{=: \mathcal{Q}(\beta - \beta^\natural)} + \underbrace{\frac{2}{n} \sum_{i=1}^n (\langle x_i, \beta^\natural \rangle - y_i) \langle x_i, \beta - \beta^\natural \rangle}_{=: \mathcal{M}(\beta, \beta^\natural)}. \end{aligned} \tag{2.2}$$

In this decomposition, the excess risk is expressed as a sum of two empirical processes $\mathcal{Q}(\beta - \beta^\natural)$ and $\mathcal{M}(\beta, \beta^\natural)$, both indexed by $\beta \in K_{\beta^\natural, t}$, which we call the *quadratic process* and the *multiplier process*, respectively. Note that this corresponds to a second-order Taylor expansion of $\mathcal{E}(\cdot, \beta^\natural) = \tilde{\mathcal{L}}(\cdot) - \tilde{\mathcal{L}}(\beta^\natural)$: the quadratic process is the second-order term⁷

$$\mathcal{Q}(\beta - \beta^\natural) = \frac{1}{2}(\beta - \beta^\natural)^T H_{\tilde{\mathcal{L}}}(\beta^\natural)(\beta - \beta^\natural)$$

⁷ Since the Hessian matrix $H_{\tilde{\mathcal{L}}}(\beta^\natural) \in \mathbb{R}^{p \times p}$ is actually independent of β^\natural for the squared loss, the quadratic process is translation-invariant in the sense that it only depends on $\beta - \beta^\natural$. Hence, we write $\mathcal{Q}(\beta - \beta^\natural)$ rather than $\mathcal{Q}(\beta, \beta^\natural)$.

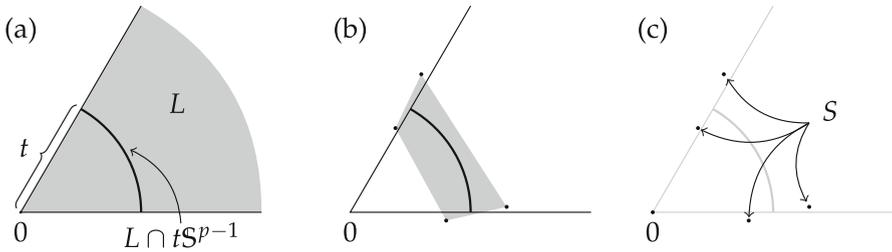


Fig. 1 An illustration of the local q -complexity $q_{t,n}^{(g,e)}(L)$ from Definition 2.5: **(a)** in order to measure the complexity of L locally at scale t , we consider the set $L \cap t\mathbb{S}^{p-1}$. **(b)** $L \cap t\mathbb{S}^{p-1}$ is contained in the convex hull of the four points indicated by small black dots. **(c)** Defining S as these four points, the quantity $\frac{1}{t} \cdot (\gamma_1(S, \|\cdot\|_e) / \sqrt{n} + \gamma_2(S, \|\cdot\|_g + \|\cdot\|_e))$ is an upper bound for the local q -complexity of L at scale t , which is defined as the infimum over all such upper bounds

and the multiplier process is the first-order term

$$\mathcal{M}(\beta, \beta^\natural) = \langle \nabla \bar{\mathcal{L}}(\beta^\natural), \beta - \beta^\natural \rangle.$$

With this notation at hand, the desired uniform lower bound $\mathcal{E}(\beta, \beta^\natural) > 0$ amounts to the event that $\mathcal{Q}(\beta - \beta^\natural)$ dominates $-\mathcal{M}(\beta, \beta^\natural)$ on the whole index set $K_{\beta^\natural,t}$.

Based on the γ -functional, we now define two general complexity parameters, which are adapted to the analysis of the quadratic process and the multiplier process, respectively. Both parameters are tailored to the above notion of generic Bernstein concentration and have in common that they measure the complexity of a set *locally*, i.e., at a certain scale $t > 0$. This reflects the fact that we are only interested in the behavior of the empirical processes on $K_{\beta^\natural,t}$ and not on the full hypothesis set K .

Definition 2.5 (Local q -complexity) Let $L \subset \mathbb{R}^p$ and let $\|\cdot\|_g$ and $\|\cdot\|_e$ be seminorms on \mathbb{R}^p . For $t > 0$, we define the *local q -complexity* of L at scale t and sample size n with respect to $(\|\cdot\|_g, \|\cdot\|_e)$ by

$$q_{t,n}^{(g,e)}(L) := \frac{1}{t} \inf \left\{ \frac{\gamma_1(S, \|\cdot\|_e)}{\sqrt{n}} + \gamma_2(S, \|\cdot\|_g + \|\cdot\|_e) \mid S \subset \mathbb{R}^p, \text{conv}(S) \supset L \cap t\mathbb{S}^{p-1} \right\}.$$

Remarkably, we do not simply measure the size of the set $L \cap t\mathbb{S}^{p-1}$ in Definition 2.5, but optimize over all “skeletons” S of this set; see Fig. 1 for an illustration and [36, Appx. A] for a related approach in the literature.

Definition 2.6 (Local m -complexity) Let $L \subset \mathbb{R}^p$ and let $\|\cdot\|_g$ and $\|\cdot\|_e$ be seminorms on \mathbb{R}^p . For $t > 0$, we define the *local m -complexity* of L at scale t with respect to $(\|\cdot\|_g, \|\cdot\|_e)$ by

$$m_t^{(g,e)}(L) := \frac{1}{t} \inf \left\{ \gamma_1(S, \|\cdot\|_e) + \gamma_2(S, \|\cdot\|_g) \mid S \subset \mathbb{R}^p, \text{conv}(S) \supset (L \cap t\mathbb{S}^{p-1}) \cup \{0\} \right\}.$$

It is worth noting that in the well-understood case of sub-Gaussian sample data, the q -complexity and m -complexity can be both identified with the notion of *local*

Gaussian width; see also Sect. 3.2 for more details. In general, however, this simple geometric interpretation is no longer valid and the behavior of both parameters is highly non-trivial. We will return to this important issue later in Sect. 3.6.

In order to control the quadratic process (in Sect. 5.2.1), we will apply the *small-ball method*, which is a powerful tool to establish uniform lower bounds for non-negative empirical processes (see [26, 31, 33]). For this purpose, the notion of a small-ball function is required:

Definition 2.7 (*Small-ball function*; [26, p. 12995]) Let $L \subset \mathbb{R}^p$ and let x be a random vector in \mathbb{R}^p . For $\theta \geq 0$, we define the *small-ball function*

$$Q_\theta(L, x) := \inf_{v \in L} \mathbb{P}(|\langle x, v \rangle| \geq \theta).$$

Since we are aiming at an error bound relative to an arbitrary target vector $\beta^\natural \in K$, it is natural that this error bound depends on how well the associated linear hypothesis $\langle x, \beta^\natural \rangle$ predicts the actual output variable y (which may depend on x in a non-linear way). In other words, the estimation performance of (LS_K) is also affected by the behavior of the *model mismatch* $y - \langle x, \beta^\natural \rangle$, measuring how much y deviates from the linear model $\langle x, \beta^\natural \rangle$. The following parameters allow us to make this precise:

Definition 2.8 (*Mismatch parameters*) Given $\beta^\natural \in \mathbb{R}^p$ and a random pair $(x, y) \in \mathbb{R}^p \times \mathbb{R}$, the *mismatch deviation* of β^\natural is defined by

$$\sigma(\beta^\natural) := \|y - \langle x, \beta^\natural \rangle\|_{\psi_1}$$

and the (*global*) *mismatch covariance* of β^\natural by

$$\rho(\beta^\natural) := \|\mathbb{E}[(y - \langle x, \beta^\natural \rangle)x]\|_2.$$

Moreover, for $t \geq 0$ and $K \subset \mathbb{R}^p$, we define the (*local*) *mismatch covariance* of β^\natural at scale t by

$$\rho_t(\beta^\natural) := \sup_{v \in K^t} \langle \mathbb{E}[(y - \langle x, \beta^\natural \rangle)x], v \rangle,$$

where $K^t := \frac{1}{t}(K - \beta^\natural) \cap \mathbb{S}^{p-1}$ for $t > 0$ and $K^0 := \text{cone}(K - \beta^\natural) \cap \mathbb{S}^{p-1}$.

As the name suggests, the mismatch covariance captures the covariance between the input vector x and the model mismatch $y - \langle x, \beta^\natural \rangle$. Inspired by linear regression problems, it is useful to think of the model mismatch as “noise” that perturbs the linear model $\langle x, \beta^\natural \rangle$. In particular, if $\mathbb{E}[(y - \langle x, \beta^\natural \rangle)x] = 0$, this noise is uncorrelated with all input variables (but not necessarily independent), implying that $\rho(\beta^\natural) = \rho_t(\beta^\natural) = 0$. In contrast, the mismatch deviation measures the sub-exponential tail behavior of the model mismatch. Note that in the noisy linear case, i.e., $y = \langle x, \beta^\natural \rangle + v$, we simply have that $\sigma(\beta^\natural) = \|v\|_{\psi_1}$. The interested reader is referred to Appendix B.1 for further remarks on the above notions of the mismatch covariance.

2.2 A local error bound for (LS_K)

Before stating the error bound, let us formally summarize our assumptions about the sampling process:

Assumption 2.9 (*Model setup*) Let $(x, y) \in \mathbb{R}^p \times \mathbb{R}$ be a joint random pair where $x \in \mathbb{R}^p$ satisfies generic Bernstein concentration with respect to $(\|\cdot\|_g, \|\cdot\|_e)$ and $y \in \mathbb{R}$ is sub-exponential. Moreover, let $K \subset \mathbb{R}^p$ be a convex hypothesis set. We define the set $K^\Delta := \text{span}(K - K) \cap \mathbb{S}^{p-1}$ and assume that x satisfies the small-ball condition

$$Q_{2\tau}(K^\Delta, x) > 0 \tag{2.3}$$

for some $\tau > 0$. Finally, we assume that the observed sample pairs $(x_1, y_1), \dots, (x_n, y_n)$ are independent copies of (x, y) .

Although it can be helpful to imagine a semi-parametric relationship between x and y (see Sect. 3.1), such an assumption is not required at the current level of abstraction. Indeed, our main result, which is presented next, provides a generic error bound for the generalized Lasso (LS_K) without any specific observation model.

Theorem 2.10 (General error bound for (LS_K) , local version) *Let Assumption 2.9 be satisfied and fix a vector $\beta^\natural \in K$. Then there exists a universal constant $C > 0$ such that for every $u \geq 8$ and $t \geq 0$, the following holds true with probability at least $1 - 5 \exp(-C \cdot u^2) - 2 \exp(-C \cdot \sqrt{n})$: If the sample size obeys⁸*

$$n \gtrsim \left(\frac{q_{t,n}^{(g,e)}(K - \beta^\natural) + \tau \cdot u}{\tau \cdot Q_{2\tau}(K^\Delta, x)} \right)^2 \tag{2.4}$$

and we have that

$$t \gtrsim \frac{1}{(\tau \cdot Q_{2\tau}(K^\Delta, x))^2} \cdot \left[\rho_t(\beta^\natural) + u^2 \cdot \sigma(\beta^\natural) \cdot \frac{m_t^{(g,e)}(K - \beta^\natural)}{\sqrt{n}} \right]_+, \tag{2.5}$$

then every minimizer $\hat{\beta}$ of (LS_K) satisfies $\|\hat{\beta} - \beta^\natural\|_2 \leq t$.

The interpretation of the error bound established in Theorem 2.10 is not straightforward, since the right-hand side of (2.5) depends on the precision level t and the right-hand side of (2.4) depends on both t and n . But regardless of these implicit dependencies, the above statement has almost the same syntactic form as in the case of sub-Gaussian sample data, e.g., see [15, Thm. 3.6], and we can rely on the interpretation suggested there. The following way of reading Theorem 2.10 is quoted from [15, p. 41], except that the mathematical terms and the equation numbers have been altered accordingly:

⁸ For the case of exact recovery, i.e., $t = 0$, the corresponding complexity parameters $q_{0,n}^{(g,e)}(K - \beta^\natural)$ and $m_0^{(g,e)}(K - \beta^\natural)$ are introduced further below in Definition 2.11 in Sect. 2.3.

A convenient way to read the above statement is as follows: First, fix an estimation accuracy t that can be tolerated. Then adjust the sample size n and $\beta^\natural \in K$ such that (2.4) and (2.5) are both fulfilled (if possible at all). In particular, if n is chosen such that (2.5) just holds with equality (up to a constant), we obtain an error bound of the form

$$\|\hat{\beta} - \beta^\natural\|_2 \lesssim \frac{1}{(\tau \cdot Q_{2\tau}(K^\Delta, x))^2} \cdot \left[\rho_t(\beta^\natural) + u^2 \cdot \sigma(\beta^\natural) \cdot \frac{m_t^{(g,e)}(K - \beta^\natural)}{\sqrt{n}} \right]_+ \tag{2.6}$$

With that in mind, one might be tempted to think that not much changes when going beyond sub-Gaussian—but this is far from being true. The key difference becomes manifested in our generalized complexity parameters $q_{t,n}^{(g,e)}(K - \beta^\natural)$ and $m_t^{(g,e)}(K - \beta^\natural)$. In fact, their behavior can be significantly more complicated than in the sub-Gaussian case. We defer a more detailed discussion to Sect. 3.6, but also the applications in Sects. 3.2–3.5 can be helpful for a better understanding of this issue. Finally, several additional remarks on Theorem 2.10 and possible refinements can be found in Appendix B.2.

2.3 Global and conic error bounds for (LS_K)

The local complexity parameters in Theorem 2.10 lead to a fairly strong, but implicit error bound for (LS_K) . In this subsection, we state two corollaries of Theorem 2.10 which achieve a better interpretability at the price of suboptimality. The first one replaces the local complexity terms by their (more pessimistic) conic versions:

Definition 2.11 (*Conic q - and m -complexity*) Let $L \subset \mathbb{R}^p$ and let $\|\cdot\|_g$ and $\|\cdot\|_e$ be semi-norms on \mathbb{R}^p . We define the *conic q -complexity* of L at sample size n with respect to $(\|\cdot\|_g, \|\cdot\|_e)$ by

$$q_{0,n}^{(g,e)}(L) := \inf \left\{ \frac{\gamma_1(S, \|\cdot\|_e)}{\sqrt{n}} + \gamma_2(S, \|\cdot\|_g + \|\cdot\|_e) \mid S \subset \mathbb{R}^p, \text{conv}(S) \supset \text{cone}(L) \cap \mathbb{S}^{p-1} \right\}.$$

Similarly, we define the *conic m -complexity* of L with respect to $(\|\cdot\|_g, \|\cdot\|_e)$ by

$$m_0^{(g,e)}(L) := \inf \left\{ \gamma_1(S, \|\cdot\|_e) + \gamma_2(S, \|\cdot\|_g) \mid S \subset \mathbb{R}^p, \text{conv}(S) \supset (\text{cone}(L) \cap \mathbb{S}^{p-1}) \cup \{0\} \right\}.$$

The subscript ‘0’ in $q_{0,n}^{(g,e)}(L)$ and $m_0^{(g,e)}(L)$ indicates that one can imagine the conic q - and m -complexity as the limit case $t = 0$ of their local counterparts from Definition 2.5 and 2.6.

The conic complexity parameters allow us to remove the dependence of the right-hand sides of (2.4) and (2.5) on t in Theorem 2.10:

Corollary 2.12 (General error bound for (LS_K) , conic version) *The assertion of Theorem 2.10 remains valid if $q_{t,n}^{(g,e)}(K - \beta^\natural)$ is replaced by $q_{0,n}^{(g,e)}(K - \beta^\natural)$ in (2.4),*

while $m_t^{(g,e)}(K - \beta^\natural)$ and $\rho_t(\beta^\natural)$ are replaced by $m_0^{(g,e)}(K - \beta^\natural)$ and $\rho_0(\beta^\natural)$ in (2.5), respectively.

While leading to an explicit error bound for the generalized Lasso (cf. (2.6)), Corollary 2.12 has the following drawback: If β^\natural is an interior point of K , then $\text{cone}(K - \beta^\natural) = \mathbb{R}^p$, and the complexity terms $q_{0,n}^{(g,e)}(K - \beta^\natural)$ and $m_0^{(g,e)}(K - \beta^\natural)$ are equal to $q_{0,n}^{(g,e)}(\mathbb{R}^p)$ and $m_0^{(g,e)}(\mathbb{R}^p)$, respectively, i.e., they no longer reflect any complexity reduction due to the restricted hypothesis set K . Hence, unless the hypothesis set K is perfectly tuned such that β^\natural is located on the boundary of K , Corollary 2.12 fails to provide a useful estimation guarantee in the high-dimensional regime $p \gg n$. Evidently, this tuning problem affects the local error bound in Theorem 2.10 as well, but the situation is much less severe there, at least when β^\natural is close to the boundary of K (more precisely, if $\inf_{\beta \in \mathbb{R}^p \setminus K} \|\beta^\natural - \beta\|_2 < t$). This fact particularly explains why (LS_K) is a *stable* estimator (cf. [15, Prop. 2.6 and Cor. 3.15]).

Our second approach to simplify Theorem 2.10 is to measure the complexity of the hypothesis set “globally”, rather than in a local neighborhood of β^\natural .

Definition 2.13 (*Global q - and m -complexity*) Let $L \subset \mathbb{R}^p$ and let $\|\cdot\|_g$ and $\|\cdot\|_e$ be semi-norms on \mathbb{R}^p . We define the *global q -complexity* of L at sample size n with respect to $(\|\cdot\|_g, \|\cdot\|_e)$ by

$$q_n^{(g,e)}(L) := \inf \left\{ \frac{\gamma_1(S, \|\cdot\|_e)}{\sqrt{n}} + \gamma_2(S, \|\cdot\|_g + \|\cdot\|_e) \mid S \subset \mathbb{R}^p, \text{conv}(S) \supset L \right\}.$$

Similarly, we define the *global m -complexity* of L with respect to $(\|\cdot\|_g, \|\cdot\|_e)$ by

$$m^{(g,e)}(L) := \inf \left\{ \gamma_1(S, \|\cdot\|_e) + \gamma_2(S, \|\cdot\|_g) \mid S \subset \mathbb{R}^p, \text{conv}(S) \supset L \right\}.$$

The following lemma provides some basic facts about the global complexity parameters and relates them to their local counterparts.

Lemma 2.14 *Let $L \subset \mathbb{R}^p$, $v \in \mathbb{R}^p$, and $t > 0$. Then we have the following:*

- (i) $q_{t,n}^{(g,e)}(L) \leq \frac{1}{t} q_n^{(g,e)}(L)$ and $m_t^{(g,e)}(L) \leq \frac{1}{t} m^{(g,e)}(L \cup \{0\})$,
- (ii) $q_n^{(g,e)}(L) = q_n^{(g,e)}(L + v)$ and $m^{(g,e)}(L) = m^{(g,e)}(L + v)$,
- (iii) $q_n^{(g,e)}(L) \lesssim m^{(g,e)}(L)$,
- (iv) $q_{0,n}^{(g,e)}(L) = q_n^{(g,e)}(\text{cone}(L) \cap \mathbb{S}^{p-1})$ and $m_0^{(g,e)}(L) = m^{(g,e)}((\text{cone}(L) \cap \mathbb{S}^{p-1}) \cup \{0\})$,
- (v) $q_{t,n}^{(g,e)}(L) = q_n^{(g,e)}(\frac{1}{t}L \cap \mathbb{S}^{p-1})$ and $m_t^{(g,e)}(L) = m^{(g,e)}((\frac{1}{t}L \cap \mathbb{S}^{p-1}) \cup \{0\})$.

The second claim of Lemma 2.14 states that the global complexity parameters are translation-invariant. This allows us to decouple the complexity terms in Theorem 2.10 from β^\natural , leading to the following error bound:

Corollary 2.15 (General error bound for (LS_K) , global version) *Let Assumption 2.9 be satisfied and fix a vector $\beta^\natural \in K$. Then there exists a universal constant $C > 0$ such*

that for every $u \geq 8$, the following holds true with probability at least $1 - 5 \exp(-C \cdot u^2) - 2 \exp(-C \cdot \sqrt{n})$: If the sample size obeys

$$n \gtrsim \left(\frac{q_n^{(g,e)}(K) + \tau \cdot u}{\tau \cdot Q_{2\tau}(K^\Delta, x)} \right)^2, \tag{2.7}$$

then every minimizer $\hat{\beta}$ of (LS_K) satisfies

$$\|\hat{\beta} - \beta^\natural\|_2 \lesssim \max \left\{ 1, (\tau \cdot Q_{2\tau}(K^\Delta, x))^{-2} \right\} \cdot \left[\rho_0(\beta^\natural) + \max\{1, u^2 \cdot \sigma(\beta^\natural)\} \cdot \frac{\sqrt{m^{(g,e)}(K)}}{n^{1/4}} \right]_+. \tag{2.8}$$

If the complexity terms $q_n^{(g,e)}(K)$ and $m^{(g,e)}(K)$ are sufficiently small, then (2.8) provides a useful error bound in the high-dimensional regime $p \gg n$, independently of the location of β^\natural in K . Note that a prototypical application of Corollary 2.15 was already presented in Proposition 1.4 where K is a convex polytope (see also Proposition 3.13 in Sect. 3.6). However, the simplification of Corollary 2.15 has its price: the second summand in the error bound (2.15) exhibits a decay rate $O(n^{-1/4})$, which is substantially worse than the rate of $O(n^{-1/2})$ achieved in Theorem 2.10 and Corollary 2.12. Moreover, the dependence on $\sigma(\beta^\natural)$ is suboptimal in the “low-noise” regime, i.e., when $\sigma(\beta^\natural) \ll 1$.

3 Applications and examples

This section is devoted to specific applications of the generic error bounds presented in Sect. 2. We begin with a discussion of semi-parametric estimation problems in Sect. 3.1, in particular, how the generalized Lasso (LS_K) performs with non-linear output models. In Sects. 3.2–3.5, we then demonstrate that generic Bernstein concentration covers a whole “spectrum” of relevant distributions, where (uniformly) sub-exponential and sub-Gaussian input vectors appear just as marginal cases. Finally, we continue our discussion on the notions of q- and m-complexity in Sect. 3.6, thereby focusing on the prototypical situation of sparse recovery via ℓ^1 -constraints.

3.1 Semi-parametric estimation problems and the mismatch principle

We intentionally did not make a concrete choice of the target vector β^\natural in Sect. 2. This strategy has led to very flexible (generic) error bounds for (LS_K) , but it does not address any specific estimation problem. As already pointed out subsequently to the initial Problem 1.1, a valid choice of β^\natural is the expected risk minimizer. Indeed, assuming that x is isotropic and $\beta^\natural := \mathbb{E}[yx] \in K$, then β^\natural is the expected risk minimizer (on both K and \mathbb{R}^p) and we have that $\rho(\beta^\natural) = \rho_r(\beta^\natural) = 0$ (see Appendix B.1 and Fig. 2 there). Hence, according to Theorem 2.10 (or its corollaries), (LS_K) yields a consistent estimator of β^\natural .

While such a statement is common in statistical learning, a much less obvious phenomenon is the capability of (LS_K) to solve *semi-parametric estimation problems*. In the context of this article, we may express a semi-parametric observation model as follows:

$$y = F(x, \beta_0),$$

where $\beta_0 \in \mathbb{R}^p$ is an unknown *parameter vector* and $F : (\mathbb{R}^p \times \mathbb{R}^p) \rightarrow \mathbb{R}$ a scalar *output function* which can be non-linear, random, and unknown. Agreeing on this model setup, the ultimate hope is now that (LS_K) is a (consistent) estimator of β_0 .⁹ It turns out that this is often possible at least to a certain extent, even though fitting a linear model to non-linear observations might appear counterintuitive at first sight. A typical example is the simple classification rule $y = \text{sign}(\langle x, \beta_0 \rangle)$, where there is still hope to recover the direction of β_0 , but not its magnitude. This limitation gives rise to a relaxed estimation problem:

Problem 3.1 *Is the generalized Lasso (LS_K) capable of estimating any element from a certain target set $T_{\beta_0} \subset \mathbb{R}^p$, which contains all those parameter vectors that allow us to extract the information of interest?*

Similarly to the more general formulation of Problem 1.1, the term ‘information’ is left unspecified here and depends on what a user considers as a desirable outcome of an estimation procedure. In the above example of binary classification, a natural choice of target set would be $T_{\beta_0} := \text{span}(\{\beta_0\})$, if one is interested in the recovery of any scalar multiple of β_0 .

Our guarantees from Sect. 2 allow us to tackle Problem 3.1 in a very systematic way:

Select $\beta^\natural \in T_{\beta_0} \cap K$ such that the mismatch covariance $\rho(\beta^\natural)$ becomes as small as possible.^{10,11} Then apply Theorem 2.10 (or one of its corollaries) to obtain an error bound for the estimation error $\|\hat{\beta} - \beta^\natural\|_2$.

This strategy ensures that the resulting target vector β^\natural encodes the desired information, while the (asymptotic) bias of (LS_K) is brought under control. In particular, if $\rho(\beta^\natural) = 0$, we achieve a consistent estimator of β^\natural ; note that the corresponding mismatch deviation $\sigma(\beta^\natural)$ can still be large, but its size only affects the variance of the error $\|\hat{\beta} - \beta^\natural\|_2$. The approach just described was developed by Genzel [15, Chap. 4], where it is referred to as the *mismatch principle* (see also the technical report [17]). It is worth pointing out that there is an important conceptual difference to the “naive” idea of first explicitly computing the expected risk minimizer (on K) and then finding the closest point on the target set T_{β_0} : indeed, we measure the complexity of K locally at β^\natural , which enables us to exploit beneficial geometric features directly on T_{β_0} .

⁹ If F is non-linear, this can be very different from asking for the expected risk minimizer, which would simply yield the best linear predictor of y .

¹⁰ For the sake of clarity, we only consider the global mismatch covariance here, which is easier to interpret and forms an upper bound for $\rho_t(\beta^\natural)$ according to Appendix B.1; but refinements are certainly possible when analyzing $\rho_t(\beta^\natural)$ instead of $\rho(\beta^\natural)$.

¹¹ If x is isotropic, this selection procedure has a nice geometric interpretation due to Appendix B.1: The mismatch covariance $\rho(\beta^\natural)$ is minimized on $T_{\beta_0} \cap K$ if and only if β^\natural is a Euclidean projection of the (global) expected risk minimizer $\mathbb{E}[y \cdot x]$ onto $T_{\beta_0} \cap K$.

We refer the reader to [15, Chap. 4] for a more extensive discussion of the mismatch principle and various applications to semi-parametric estimation problems. In the present work, we confine ourselves to an illustration in the prototypical situation of single-index models.

Proposition 3.2 ([15, Prop. 4.6]) *Let $x \in \mathbb{R}^p$ be a centered, isotropic random vector. We assume that y obeys a single-index model of the form*

$$y = f(\langle x, \beta_0 \rangle) + v, \quad (3.1)$$

where $\beta_0 \in \mathbb{R}^p \setminus \{0\}$ is an unknown parameter vector, $f : \mathbb{R} \rightarrow \mathbb{R}$ is a scalar output function, and v is independent noise with $\mathbb{E}[v] = 0$. Moreover, we choose $T_{\beta_0} := \text{span}(\{\beta_0\})$ as target set. Then $\beta^\natural = \mu\beta_0$ with

$$\mu := \frac{1}{\|\beta_0\|_2^2} \cdot \mathbb{E}[f(\langle x, \beta_0 \rangle) \cdot \langle x, \beta_0 \rangle]$$

minimizes the (global) mismatch covariance over T_{β_0} and we have that

$$\rho(\beta^\natural) = \|\mathbb{E}[f(\langle x, \beta_0 \rangle) P_{\beta_0}^\perp x]\|_2,$$

where $P_{\beta_0}^\perp \in \mathbb{R}^{p \times p}$ is the projection onto the orthogonal complement of $\text{span}(\{\beta_0\})$. In particular, if x is a standard Gaussian random vector, we have that $\rho(\beta^\natural) = 0$.

In the special case of a Gaussian input vector, Proposition 3.2 reproduces the original finding of Plan and Vershynin [42]: despite an unknown, non-linear distortion, the generalized Lasso still allows for consistent estimation of the parameter vector, or at least a scalar multiple of it. When combining Proposition 3.2 with the results from Sect. 2 (for an appropriately tuned hypothesis set K), we observe that their conclusion remains essentially valid for non-Gaussian inputs as long as the mismatch covariance $\rho(\beta^\natural)$ vanishes or gets sufficiently small. On the other hand, if $\rho(\beta^\natural)$ is too large, it can be useful to employ an adaptive estimator instead (e.g., see [61–63]), but there also exist worst-case scenarios where an asymptotic bias is inevitable, regardless of the considered estimator (see [2]). For an overview of the extensive literature on single-index models as well as historical references, we refer the reader to [43, Sec. 6] and [63, Subsec. 1.2]. Moreover, see [15, Subsec. 4.2.2] and the references therein for related works in 1-bit compressed sensing.

Finally, it is worth pointing out that we did not make any (explicit) assumptions on the tail behavior of the distribution of x in this subsection. Therefore, one can easily combine the above described approach with the findings of the forthcoming subsections, which investigate specific instances of generic Bernstein concentration.

3.2 Sub-Gaussian input vectors

The current and subsequent subsections are devoted to several examples of generic Bernstein concentration (see Definition 2.2). Let us begin with the situation of (uniformly) sub-Gaussian input vectors. A characteristic property of sub-Gaussian random

variables is that their tails are essentially not heavier than those of the normal distribution (see Proposition A.1(i)). Combining this property with Definition 1.3, one can easily verify that a sub-Gaussian random vector $x \in \mathbb{R}^p$ with sub-Gaussian norm $\|x\|_{\psi_2}$ exhibits generic Bernstein concentration with respect to $(C\|x\|_{\psi_2} \|\cdot\|_2, 0)$ for a universal constant $C > 0$. Since the sub-exponential part of the mixed-tail condition is effectively erased here by setting $\|\cdot\|_e := 0$, we observe that sub-Gaussian input vectors form a degenerate limit case at the lighter-tailed end of the “spectrum” of generic Bernstein concentration. Regarding the q - and m -complexities, the identity $\|\cdot\|_e = 0$ implies that the γ_1 -functional effectively vanishes in their respective definitions, so that we end up with a rescaled version of the functional $\gamma_2(\cdot, \|\cdot\|_2)$. The celebrated Majorizing Measure Theorem of Talagrand [51, Thm. 2.4.1] relates this functional to a well-known complexity parameter:

Definition 3.3 (*Gaussian width*) Let $L \subset \mathbb{R}^p$ and let $g \sim \mathcal{N}(0, I_p)$ be a standard Gaussian random vector. The *Gaussian width* of L is defined by

$$w(L) := \mathbb{E} \left[\sup_{v \in L} \langle g, v \rangle \right].$$

The Gaussian width originates from classical results in geometric functional analysis and asymptotic convex geometry, e.g., see [18, 20, 35]. More recently, it has emerged as a useful tool for the analysis of high-dimensional estimation problems, e.g., see [3, 10, 34, 37, 44, 49, 57]. The connection to our analysis, which is provided by Talagrand’s Majorizing Measure Theorem, is the fact that for every subset $L \subset \mathbb{R}^p$, we have

$$\gamma_2(L, \|\cdot\|_2) \asymp w(L). \tag{3.2}$$

Apart from a simple geometric interpretation of the resulting complexity parameters, (3.2) implies that the optimization over the “skeleton” is irrelevant (up to constants) in the sub-Gaussian case, since the Gaussian width is invariant under taking the convex hull. This explains why such an optimization is uncommon in the literature dealing with sub-Gaussian input data.

The following fact summarizes the above considerations and allows us to relate the generic (global) error bound from Corollary 2.15 to the sub-Gaussian setting:

Fact 3.4 Let $x \in \mathbb{R}^p$ be a (uniformly) sub-Gaussian random vector, i.e., $\|x\|_{\psi_2} < \infty$. Then x exhibits generic Bernstein concentration with respect to $(\sqrt{C_2}\|x\|_{\psi_2} \|\cdot\|_2, 0)$, where $C_2 > 0$ is the constant from Proposition A.1(i). The global q - and m -complexities satisfy

$$\begin{aligned} q_n^{(g,e)}(L) &\asymp \|x\|_{\psi_2} \cdot \underbrace{\inf \left\{ \gamma_2(S, \|\cdot\|_2) \mid S \subset \mathbb{R}^p, \text{conv}(S) \supset L \right\}}_{=: q_n^{(2,0)}(L)} \asymp \|x\|_{\psi_2} \cdot w(L), \\ m^{(g,e)}(L) &\asymp \|x\|_{\psi_2} \cdot \underbrace{\inf \left\{ \gamma_2(S, \|\cdot\|_2) \mid S \subset \mathbb{R}^p, \text{conv}(S) \supset L \right\}}_{=: m^{(2,0)}(L)} \asymp \|x\|_{\psi_2} \cdot w(L). \end{aligned} \tag{3.3}$$

The comparison of our results with existing ones is facilitated by introducing the *normalized* complexities $q_n^{(2,0)}$ and $m^{(2,0)}$ in (3.3). Indeed, both parameters are unaffected by a rescaling of the input vector, which is a common feature of complexity measures defined in the literature. In contrast, their “unnormalized” counterparts $q_n^{(g,e)}$ and $m^{(g,e)}$ “absorb” the norm $\|x\|_{\psi_2}$ as a scalar pre-factor for the generic semi-norms $\|\cdot\|_e$ and $\|\cdot\|_g$.

With regard to the local and conic q - and m -complexities, Talagrand’s Majorizing Measure Theorem leads to similar conclusions. The (normalized) local q -complexity corresponds to the notion of *local Gaussian width* (cf. [14, 42]):

$$q_{t,n}^{(2,0)}(L) := \frac{1}{t} \inf \left\{ \gamma_2(S, \|\cdot\|_2) \mid S \subset \mathbb{R}^p, \text{conv}(S) \supset L \cap t\mathbb{S}^{p-1} \right\} \asymp \frac{1}{t} w(L \cap t\mathbb{S}^{p-1}).$$

Since the definition of the (normalized) local m -complexity requires that $\text{conv}(S)$ also contains the origin, it is not strictly equivalent to the local Gaussian width, but incorporates an additional constant term (cf. Appendix B.25.6):

$$m_t^{(2,0)}(L) := \frac{1}{t} \inf \left\{ \gamma_2(S, \|\cdot\|_2) \mid S \subset \mathbb{R}^p, \text{conv}(S) \supset (L \cap t\mathbb{S}^{p-1}) \cup \{0\} \right\} \asymp \frac{1}{t} w(L \cap t\mathbb{S}^{p-1}) + 1.$$

Analogously, the (normalized) conic q -complexity and m -complexity correspond to the notion of *conic Gaussian width* (cf. [3, 10]). A combination of these identifications with the corresponding error bounds from Sect. 2 (Theorem 2.10, Corollary 2.12, and Corollary 2.15) allows us to reproduce known estimation guarantees for (sub-)Gaussian sample data, e.g., see [15, 42].

3.3 Input vectors with independent sub-exponential features

Although it is reassuring that the generic error bounds from Sect. 2 are consistent with existing results for the sub-Gaussian case, this setup does not constitute a proper example of the mixed-tail condition in Definition 2.2. A more natural example is given by input vectors with centered, independent, sub-exponential coordinates. In this case, generic Bernstein concentration can be simply implemented by the classical Bernstein’s inequality (see Theorem A.3). The following fact is a direct consequence of Theorem A.3.

Fact 3.5 *Let $x \in \mathbb{R}^p$ be a random vector with centered, independent, sub-exponential coordinates. Set $R := \max_{1 \leq j \leq d} \|x_j\|_{\psi_1}$. Then x exhibits generic Bernstein concentration with respect to $(\frac{R}{\sqrt{C_B}} \|\cdot\|_2, \frac{R}{C_B} \|\cdot\|_\infty)$, where $C_B > 0$ is the constant from Theorem A.3. The global q - and m -complexities satisfy*

$$q_n^{(g,e)}(L) \asymp R \cdot \inf \left\{ \frac{\gamma_1(S, \|\cdot\|_\infty)}{\sqrt{n}} + \gamma_2(S, \|\cdot\|_2) \mid S \subset \mathbb{R}^p, \text{conv}(S) \supset L \right\},$$

$$m^{(g,e)}(L) \asymp R \cdot \underbrace{\inf \left\{ \gamma_1(S, \|\cdot\|_\infty) + \gamma_2(S, \|\cdot\|_2) \mid S \subset \mathbb{R}^p, \text{conv}(S) \supset L \right\}}_{=: m^{(2,\infty)}(L)}.$$

Remarkably, there exists a geometric interpretation of $m^{(2,\infty)}$ that is very similar to (3.2). To this end, let $L \subset \mathbb{R}^p$ and let $Y = (Y_1, \dots, Y_p)$ be a random vector with independent, symmetric coordinates satisfying $\mathbb{P}(|Y_j| \geq t) = \exp(-t)$ for all $j = 1, \dots, p$. By a result of Talagrand [51, Thm. 10.2.8], we then have

$$\gamma_1(L, \|\cdot\|_\infty) + \gamma_2(L, \|\cdot\|_2) \asymp \mathbb{E} \left[\sup_{v \in L} \langle v, Y \rangle \right]. \tag{3.4}$$

Inspired by the notion of ‘Gaussian’ width, the expression on the right-hand side is referred to as the *exponential width* of L . Similarly to the sub-Gaussian case in Sect. 3.2, the relation (3.4) shows that the optimization over the “skeleton” does not make a difference in the present scenario, at least when ignoring universal constants. Therefore, we can conclude that the normalized m-complexity is equivalent to the exponential width, i.e., $m^{(2,\infty)}(L) \asymp \mathbb{E} \left[\sup_{v \in L} \langle v, Y \rangle \right]$.

The idea of using the exponential width as a complexity measure for sub-exponential input vectors was proposed by Sivakumar et al. [48]. In contrast to the Gaussian width, the exponential width is not rotation-invariant: only the sub-Gaussian component of the complexity is tied to the Euclidean structure on \mathbb{R}^p , whereas the sub-exponential component of the complexity is described by the ℓ^∞ -norm. Based on results by Talagrand, it follows from [48, Thm. 1] that for every subset $L \in \mathbb{R}^p$, we have

$$m^{(2,\infty)}(L) \asymp \mathbb{E} \left[\sup_{v \in L} \langle v, Y \rangle \right] \lesssim \sqrt{\log(p)} \cdot w(L). \tag{3.5}$$

Regarding the feasibility of estimation in the high dimensions, this shows that the situation for input vectors with independent, sub-exponential features is not substantially worse than for sub-Gaussian sample data.

According to Lemma 2.14(iii), the normalized global q-complexity $q_n^{(2,\infty)}$ (defined analogously to $m^{(2,\infty)}$) can also be controlled by the exponential width. With this in mind, our lower bound for the quadratic process in Proposition 5.5 resembles a result of Sivakumar et al. [48, Thm. 3], if we assume input vectors with independent, sub-exponential entries. In contrast, [48, Thm. 3] just requires isotropy and uniform sub-exponentiality. Due to a lack of published proofs, we were unable to verify that the exponential width is the correct complexity measure for this more general scenario.

3.4 (Uniformly) sub-exponential input vectors

We were already concerned with the situation of (uniformly) sub-exponential input vectors in Sect. 1.1 (see Proposition 1.4). Taking the more abstract viewpoint from Sect. 2, this setting corresponds to a degenerate limit case at the heavier-tailed end of the “spectrum” of generic Bernstein concentration. Indeed, an application of Proposition A.1(i) leads to the following fact:

Fact 3.6 *Let $x \in \mathbb{R}^p$ be a (uniformly) sub-exponential random vector, that is, $\|x\|_{\psi_1} < \infty$. Then x exhibits generic Bernstein concentration with respect to $(0, C_1 \|x\|_{\psi_1} \|\cdot\|_2)$, where $C_1 > 0$ is the constant from Proposition A.1(i). The global*

q- and *m*-complexities satisfy

$$\begin{aligned}
 q_n^{(g,e)}(L) &\asymp \|x\|_{\psi_1} \cdot \underbrace{\inf \left\{ \frac{\gamma_1(S, \|\cdot\|_2)}{\sqrt{n}} + \gamma_2(S, \|\cdot\|_2) \mid S \subset \mathbb{R}^p, \text{conv}(S) \supset L \right\}}_{=:q_n^{(0,2)}(L)}, \\
 m^{(g,e)}(L) &\asymp \|x\|_{\psi_1} \cdot \underbrace{\inf \left\{ \gamma_1(S, \|\cdot\|_2) \mid S \subset \mathbb{R}^p, \text{conv}(S) \supset L \right\}}_{=:m^{(0,2)}(L)}.
 \end{aligned}$$

According to Talagrand’s Majorizing Measure Theorem, we have that

$$q_n^{(0,2)}(L) \asymp \inf \left\{ \frac{\gamma_1(S, \|\cdot\|_2)}{\sqrt{n}} + w(S) \mid S \subset \mathbb{R}^p, \text{conv}(S) \supset L \right\}. \tag{3.6}$$

The right-hand side of (3.6) agrees with the notion of *perturbed width* that was considered by [36, 46]. Since Sattar and Oymak [46] focus on the projected gradient descent as an algorithmic implementation of the generalized Lasso (LS_K), their error bounds are not directly comparable to ours (in the special case of sub-exponential input vectors), but they bear a resemblance. A remarkable difference is that we achieve an exponentially decaying probability of failure. This is due to the fact that we handle the multiplier process by Mendelson’s chaining approach (see Sect. 5.2.2), which also explains why the notion of *m*-complexity does not appear in the results of [46].

Unlike in the settings of Sects. 3.2 and 3.3, a simple geometric interpretation of the complexity parameters is not available in the case of uniformly sub-exponential input vectors. In particular, there is no reason to believe that the optimization over the “skeleton” is unnecessary in general. Compared to the γ_2 -functional, which can be controlled by means of the Gaussian width, the γ_1 -functional seems more mysterious and intangible. However, it is at least possible to derive informative upper bounds in the important special case where K is a scaled ℓ^1 -ball (see Sect. 3.6).

3.5 The lifted Lasso and phase retrieval

Another relevant example of generic Bernstein concentration occurs when applying the so-called lifted Lasso to sub-Gaussian input vectors. The lifted Lasso introduced below can be seen as a variant of the phase lift approach (see [7, 8]), which is tailored to the phase retrieval problem (e.g., see [47] for an overview). Our statistical analysis is not limited to this specific model setup and covers the more general scenario considered by Thrampoulidis and Rawat [52], namely single-index models with *even* output functions. In fact, Proposition 3.2 indicates that this is a highly non-trivial task: if y obeys (3.1) with Gaussian inputs and an even function f , then we would simply have $\mu = 0$, so that the ordinary Lasso (LS_K) fails to recover the direction of the parameter vector β_0 .

Phase lifting follows a different approach that allows us to reduce the non-linear phase retrieval problem to a more accessible linear problem. It is based on the simple,

yet crucial observation that

$$\langle x, \beta \rangle^2 = \text{tr}(xx^T \beta \beta^T) = \langle xx^T, \beta \beta^T \rangle_F \quad \text{for all } x, \beta \in \mathbb{R}^p,$$

where $\text{tr}(\cdot)$ denotes the trace and $\langle \cdot, \cdot \rangle_F$ the Hilbert-Schmidt inner product. The *lifted Lasso* then corresponds to the following convex optimization problem:

$$\min_{B \in H} \frac{1}{n} \sum_{i=1}^n \left(y_i - \langle x_i x_i^T, B \rangle_F - \mathbb{E}[xx^T], B \right)_F^2, \tag{LLS_H}$$

where $H \subset \mathbb{R}^{p \times p}$ is a convex subset of the positive semidefinite cone in $\mathbb{R}^{p \times p}$ which contains all “lifted” hypotheses, i.e., $H \supset \{\beta \beta^T \mid \beta \in K\}$ for some $K \subset \mathbb{R}^p$. Note that the centering term $\mathbb{E}[xx^T]$ in (LLS_H) is important to achieve consistency (see also Proposition 3.11 below).

If the input vector $x \in \mathbb{R}^p$ is sub-Gaussian, then $xx^T \in \mathbb{R}^{p \times p}$ is sub-exponential.¹² In this sense, the lifted Lasso is a typical application where sub-exponential vectors occur naturally. The connection between the lifted setting and the notion of generic Bernstein concentration is given by the Hanson-Wright inequality (cf. [58, Thm. 6.2.1]):

Fact 3.7 *Let $x \in \mathbb{R}^p$ be a random vector with centered, independent, sub-Gaussian coordinates. Moreover, let $B \in \mathbb{R}^{p \times p}$ and set $R := \max_{1 \leq j \leq p} \|x_j\|_{\psi_2}$. Then $xx^T - \mathbb{E}[xx^T] \in \mathbb{R}^{p \times p}$ exhibits generic Bernstein concentration with respect to $(\frac{R^2}{\sqrt{C}} \|\cdot\|_F, \frac{R^2}{C} \|\cdot\|_{\text{op}})$, where $C > 0$ is a universal constant. In particular, the global q - and m -complexities satisfy¹³*

$$q_n^{(g,e)}(L) \asymp R^2 \cdot \underbrace{\inf \left\{ \frac{\gamma_1(S, \|\cdot\|_{\text{op}})}{\sqrt{n}} + \gamma_2(S, \|\cdot\|_F) \mid S \subset \mathbb{R}^{p \times p}, \text{conv}(S) \supset L \right\}}_{=: q_n^{(F,\text{op})}(L)},$$

$$m^{(g,e)}(L) \asymp R^2 \cdot \underbrace{\inf \left\{ \gamma_1(S, \|\cdot\|_{\text{op}}) + \gamma_2(S, \|\cdot\|_F) \mid S \subset \mathbb{R}^{p \times p}, \text{conv}(S) \supset L \right\}}_{=: m^{(F,\text{op})}(L)}.$$

It is noteworthy that a similar result can be achieved under different assumptions on x . For instance, Adamczak [1, Thm. 2.5] proves a Hanson-Wright inequality under the assumption of a convex concentration property, while Jeong et al. [24, Thm. 1.5] provide a somewhat different version under the additional assumption of unit variance. Alternatively, if the input vector x is just uniformly sub-Gaussian, we can make use of a result by Zajkowski [65, Prop. 2.7, Rem. 2.8] to obtain the following:

¹² Here and in the following, the matrix space $\mathbb{R}^{p \times p}$ is canonically identified with \mathbb{R}^{p^2} . In particular, we interpret xx^T as a random vector in \mathbb{R}^{p^2} , rather than a random matrix.

¹³ Note that the operator norm is absorbed by the Frobenius norm in the γ_2 -part of the q -complexity, which is possible due to $\|\cdot\|_F + \|\cdot\|_{\text{op}} \lesssim \|\cdot\|_F$.

Fact 3.8 Let $x \in \mathbb{R}^p$ be a centered, (uniformly) sub-Gaussian random vector and let $B \in \mathbb{R}^{p \times p}$. Then $xx^T - \mathbb{E}[xx^T]$ exhibits generic Bernstein concentration with respect to $(0, C \|x\|_{\psi_2}^2 \cdot \|B\|_F)$, where $C > 0$ is a universal constant.

All aforementioned results can be integrated into our framework. The following error bound is a direct application of Corollary 2.15 to the setup of Fact 3.7:

Corollary 3.9 (Error bound for (LLS_H) , global version) Let $x \in \mathbb{R}^p$ be as in Fact 3.7 and let y be sub-exponential. Let the sample pairs $(x_1, y_1), \dots, (x_n, y_n)$ be independent copies of (x, y) . Moreover, let $H \subset \mathbb{R}^{p \times p}$ be convex and fix $B^\natural \in H$. We also assume that

$$Q := Q_{2\tau}(\text{span}(H - H) \cap \mathbb{S}^{p^2-1}, xx^T - \mathbb{E}[xx^T]) > 0.$$

Then there exists a universal constant $C > 0$ such that for every $u \geq 8$, the following holds true with probability at least $1 - 5 \exp(-C \cdot u^2) - 2 \exp(-C \cdot \sqrt{n})$: If the sample size obeys

$$n \gtrsim \left(\frac{R^2 \cdot q_n^{(F, \text{op})}(H) + \tau \cdot u}{\tau \cdot Q} \right)^2,$$

then every minimizer \hat{B} of (LLS_H) satisfies

$$\|\hat{B} - B^\natural\|_F \lesssim \max \left\{ 1, (\tau \cdot Q)^{-2} \right\} \cdot \left[\rho_0(B^\natural) + \max \{ 1, u^2 \cdot \sigma(B^\natural) \} \cdot R \cdot \frac{\sqrt{m^{(F, \text{op})}(H)}}{n^{1/4}} \right]_+, \tag{3.7}$$

where the mismatch parameters $\rho_0(B^\natural)$ and $\sigma(B^\natural)$ are defined with respect to the “lifted” random pair $(xx^T - \mathbb{E}[xx^T], y)$.

Obviously, one can derive analogous estimation guarantees for the local and conic complexity parameters based on Theorem 2.10 and Corollary 2.12, respectively.

Remark 3.10 From a practical viewpoint, the error bound for the lifted Lasso in Corollary 3.9 is only of indirect interest, since the actual goal is to construct an estimator for an appropriate target vector $\beta^\natural \in \mathbb{R}^p$ with $B^\natural = \beta^\natural(\beta^\natural)^T$. For this purpose, one may simply extract the rank-one component from a solution \hat{B} to (LLS_H) . Indeed, let $\hat{\beta} \in \mathbb{S}^{p-1}$ be a unit-norm eigenvector of \hat{B} corresponding to the largest eigenvalue $\hat{\lambda}_1$ of \hat{B} (recall that \hat{B} is positive semidefinite). Then, on that same event as in Corollary 3.9, the following error bound holds true:

$$\min \left\{ \|\hat{\lambda}_1 \hat{\beta} - \beta^\natural\|_2, \|\hat{\lambda}_1 \hat{\beta} + \beta^\natural\|_2 \right\} \lesssim \min \left\{ \|\beta^\natural\|_2, \frac{\text{Err}}{\|\beta^\natural\|_2} \right\},$$

where Err is the error term on the right-hand side of (3.7); see [8, Sec. 6] and [52, Subsec. 2.1] for more details. In other words, $\hat{\lambda}_1 \hat{\beta}$ is an estimator of either β^\natural or $-\beta^\natural$, due to sign ambiguities.

Since the error bound (3.7) is affected by the (constant) additive term $\rho_0(B^\natural)$, it is natural to study situations where it vanishes. The following proposition concerns two

model setups where this is the case. In conjunction with Corollary 3.9, this shows that the lifted Lasso (LLS_H) can provide a consistent estimator for phase-retrieval-like problems.

Proposition 3.11 1. *Let $x \in \mathbb{R}^p$ be a random vector with finite second moments (so that the centering term $\mathbb{E}[xx^T]$ is well-defined). We assume that y obeys a quadratic observation model of the form*

$$y = \langle x, \beta_0 \rangle + v)^2,$$

where $\beta_0 \in \mathbb{R}^p$ is an unknown parameter vector and v is independent noise with $\mathbb{E}[v] = 0$. Then we have that $\rho_0(B^\natural) = 0$ for $B^\natural := \beta_0\beta_0^T$.

2. *Let $x \sim \mathcal{N}(0, I_p)$. We assume that y obeys a single-index model the form*

$$y = f(\langle x, \beta_0 \rangle) + v,$$

where $\beta_0 \in \mathbb{S}^{p-1}$ is an unknown parameter vector, $f : \mathbb{R} \rightarrow \mathbb{R}$ is a scalar output function, and v is independent noise with $\mathbb{E}[v] = 0$. Set $\mu := \frac{1}{2}\mathbb{E}[f(Z)(Z^2 - 1)]$ with $Z \sim \mathcal{N}(0, 1)$. Then, we have that $\rho_0(B^\natural) = 0$ for $B^\natural := \mu\beta_0\beta_0^T$.

For the sake of brevity, we omit a proof of Proposition 3.11; it is straightforward but especially the second part requires some lengthy calculations, see [52, Appx. B.2] for more details. While the first statement of Proposition 3.11 addresses the classical phase retrieval problem under no additional assumptions on the input vector, the second one indicates that the lifted Lasso can handle much more general non-linearities, at least in the Gaussian case. In fact, one can achieve a consistent estimator of $\pm\beta_0$ as long as $\mu \neq 0$, which includes a large subclass of even output functions. This observation allows us to reproduce a main result of Thrampoulidis and Rawat [52], thereby integrating it into a more general statistical framework for the lifted Lasso. They also investigate the important special case of sparse recovery, where β_0 is sparse and H is a subset of a scaled ℓ^1 -ball in \mathbb{R}^{p^2} . A detailed analysis of this situation goes beyond the scope of this paper, but we emphasize that the complexity bounds presented in the next subsection could be used to derive results in that regard.¹⁴ Finally, we refer to [52, Subsec. 1.3] for further reading on recent approaches to phase retrieval and related problems.

3.6 Sparse recovery and the complexity of polytopes

In this subsection, we discuss our complexity parameters in the context of high-dimensional estimation problems where $n \ll p$, with a particular emphasis on sparse recovery; for a comprehensive introduction to high-dimensional statistics, we refer to the textbooks [6, 13, 23, 58, 59]. The common ground of sparse recovery problems is the assumption that the underlying parameter vector is sparse in a certain sense. In

¹⁴ However, this would produce a further example of the observation that “the same k^2 -barrier appears in most of the algorithms that have been proposed for sparse recovery from quadratic measurements”; quote from [52, Subsec. 2.1]; see also [39].

this part, we focus on the specific case where the target vector $\beta^\natural \in \mathbb{R}^p$ is k -sparse, i.e., at most k of its coordinates are non-zero. Since the set of k -sparse vectors in \mathbb{R}^p is non-convex for $k < p$, it cannot be used as hypothesis set for the generalized Lasso (LS_K), and one has to come up with an appropriate convex relaxation. Probably the most natural choice is a scaled ℓ^1 -ball, which precisely leads to the standard Lasso studied by Tibshirani [55].

Let us begin with the situation where the hypothesis set is perfectly tuned in the sense that the target vector lies exactly on its boundary. The following result for the ℓ^1 -ball provides bounds for the local and conic m - and q -complexities in settings of Sects. 3.2–3.4. A proof is given in Appendix 5.6.

Proposition 3.12 *Let $t \geq 0$ and assume that $k \lesssim p$. Let $\beta^\natural \in \mathbb{R}^p$ be a k -sparse vector and let $K := \|\beta^\natural\|_1 B_1^p$ (i.e., β^\natural lies on the boundary of K). Then the following holds true:¹⁵*

(i) *The local and conic m - and q -complexities with respect to $(\|\cdot\|_2, 0)$ satisfy*

$$q_{t,n}^{(2,0)}(K - \beta^\natural) \lesssim m_t^{(2,0)}(K - \beta^\natural) \lesssim \sqrt{k \log\left(\frac{p}{k}\right)}.$$

(ii) *The local and conic m - and q -complexities with respect to $(\|\cdot\|_2, \|\cdot\|_\infty)$ satisfy*

$$q_{t,n}^{(2,\infty)}(K - \beta^\natural) \lesssim m_t^{(2,\infty)}(K - \beta^\natural) \lesssim \sqrt{k \log\left(\frac{p}{k}\right) \log(p)}.$$

(iii) *The local and conic m - and q -complexities with respect to $(0, \|\cdot\|_2)$ satisfy*

$$q_{t,n}^{(0,2)}(K - \beta^\natural) \lesssim m_t^{(0,2)}(K - \beta^\natural) \lesssim \sqrt{k} \cdot \log(2p).$$

Since all upper bounds in Proposition 3.12 scale only logarithmically with the ambient dimension p , we can conclude that sparse recovery is feasible in the settings of Sects. 3.2–3.4. Moreover, the square-root-dependence on the sparsity k is optimal in each of the above cases. In particular, the sample-size condition (2.4) of Theorem 2.10 takes the familiar form

$$n \gtrsim k \cdot \text{Polylog}(p, k),$$

where we have ignored other model-dependent parameters for the sake of clarity.

In situations where the hypothesis set is not perfectly tuned, it can be more appropriate to apply the global error bound from Corollary 2.15 instead of Theorem 2.10 (or Corollary 2.12). In this context, the “skeleton” optimization in the m - and q -complexities proves very useful. To this end, let us assume that $K \subset \mathbb{R}^p$ is a convex polytope, i.e., $K = \text{conv}(F)$ for a finite set of vertices $F \subset \mathbb{R}^p$. For $\alpha \in \{1, 2\}$, we then have

$$\gamma_\alpha(F, d) \lesssim \Delta(F) \cdot (\log(|F|))^\frac{1}{\alpha}. \tag{3.8}$$

¹⁵ The notations for the m - and q -complexities are adopted from the respective settings of Sects. 3.2–3.4; see also Lemma 2.14(iv) and (v) for the relation to their global counterparts.

This rather crude bound is proved straightforwardly by constructing an admissible partition sequence whose partitions contain all elements of F as singletons “as soon as admissible”, and by bounding every diameter trivially by $\Delta(F)$. With (3.8) at hand, we immediately obtain the following bounds for the global complexity parameters in the case of polytopal hypothesis sets:

Proposition 3.13 *Let $K \subset \mathbb{R}^p$ be a convex polytope with D vertices. Then we have*

$$q_n^{(g,e)}(K) \lesssim \frac{\Delta_e(K) \cdot \log(D)}{\sqrt{n}} + (\Delta_g(K) + \Delta_e(K)) \cdot \sqrt{\log(D)},$$

$$m^{(g,e)}(K) \lesssim \Delta_e(K) \cdot \log(D) + \Delta_g(K) \cdot \sqrt{\log(D)},$$

where $\Delta_e(K)$ and $\Delta_g(K)$ are the diameters of K with respect to the semi-norms $\|\cdot\|_e$ and $\|\cdot\|_g$, respectively.

Since the ℓ^1 -ball in \mathbb{R}^p has only $D = 2p$ vertices, Proposition 3.13 implies that sparse recovery is possible in the high-dimensional regime $n \ll p$ for all variants of generic Bernstein concentration discussed in this section, even if the hypothesis set K is not perfectly tuned. For example, if β^5 is k -sparse and has unit norm, $K = \sqrt{k}B_1^p$ would be a valid choice for Corollary 2.15. Bypassing perfect tuning is in fact a desirable feature in statistics, but we point out that the error bound (2.8) in Corollary 2.15 exhibits a suboptimal decay rate of $O(n^{-1/4})$.

The above findings indicate a noteworthy phenomenon of our complexity parameters, and generic chaining in general. The argument behind Proposition 3.13 is especially effective for those polytopes with few vertices because we then only have to control the empirical processes over this small subset (cf. (3.8)). For the $\gamma_2(\cdot, \|\cdot\|_2)$ -functional, this simplification is irrelevant, since it is equivalent to the Gaussian width according to (3.2). However, the general geometric mechanisms behind this fact remain largely mysterious, e.g., see [51, Sec. 2.4]. In particular, the situation is much less understood beyond this special case and the involved γ -functionals are not necessarily invariant under taking the convex hull. Consequently, controlling the m- and q-complexities in any specific situation is a highly non-trivial task.

4 Conclusion and outlook

Leaving aside the specific aspects and applications discussed in the previous sections, the overall conclusion of our main results reads as follows: The benchmark case of sub-Gaussian sample data can be seen as a “barrier” behind which the estimation behavior of the generalized Lasso (LS_K) can change significantly. The key difference becomes manifested in what way the complexity of the hypothesis set K is measured. Indeed, the m- and q-complexities do not generally enjoy a simple geometric interpretation similar to the Gaussian width, and except for some specific scenarios, the underlying chaining functionals are difficult to control (see Sect. 3.6). On the other hand, we have observed that several statistical and conceptual features remain valid beyond sub-Gaussianity. In particular, semi-parametric estimation problems can be treated as

before, since the consistency of the generalized (or lifted) Lasso is not affected by the tail behavior of the input vectors (see Sects. 3.1 and 3.5).

On the technical side of this paper stands an application of generic chaining, as a means of controlling the quadratic and the multiplier process according to the underlying geometry of their index sets. In our specific analysis, this paradigm appears explicitly in the notion of generic Bernstein concentration: the correct way to measure complexity is determined by the tail behavior of the input vector, which is captured by two (appropriate) semi-norms.

We close our discussion with a short list of open problems and possible extensions of our approach:

- *Beyond sub-exponentiality.* Are our results extendable to input vectors for which generic Bernstein concentration is simply too restrictive? An obvious relaxation would be that x obeys only a α -sub-exponential distribution, i.e., $\|x\|_{\psi_\alpha} < \infty$ for some $0 < \alpha < 1$ (cf. Theorem 5.10). For instance, such distributions occur naturally when studying higher-order variants of the lifted Lasso, where the input data consist of tensor products. We believe that our basic proof strategy would not break down in such scenarios. In fact, even though $\|\cdot\|_{\psi_\alpha}$ is just a quasi-norm for $0 < \alpha < 1$, concentration inequalities are available, similarly to the case $\alpha = 1$, e.g., see [4, 21, 45]. Hence, a careful adaptation of generic Bernstein concentration and the related chaining argument might lead to similar estimation guarantees as in Sect. 2. It is worth pointing out that lower bounds for the quadratic process under heavier tailed inputs are subject of recent research, e.g., see [25, 27, 30, 48], while the behavior of the multiplier process remains largely unclear.
- *The multiplier process.* The conclusion of the previous point gives rise to another relevant issue: How tight is our bound for the multiplier process (in Proposition 5.15)? Can it be improved in general or at least in specific model setups? Let us be a little more precise about this concern: Our approach to controlling the multiplier process is based on a powerful concentration inequality by Mendelson [32], formulated in Theorem 5.12. In contrast, the multiplier process is handled with more elementary arguments in most related works on the generalized Lasso, e.g., see [42, 46, 52]. These approaches suffer from a more pessimistic probability of success and may lead to different error bounds in some situations. We suspect that there exists a certain trade-off between the probability of success and the size of the related complexity terms. In this regard, a particularly interesting phenomenon is that—in contrast to the sub-Gaussian case—the complexities of the multiplier process and the quadratic process may be measured in a different way.
- *Beyond linearity and convexity.* The results of this paper are limited to *convex* hypothesis sets consisting of *linear* functions. Indeed, convexity is an important ingredient of Fact 2.4, while linearity enables the optimization over the “skeleton” in the proofs of Proposition 5.5 and 5.15. We expect that it is possible to drop the convexity assumption on K by analyzing the projected gradient descent method as an algorithmic implementation of (LS_K) , e.g., see [38, 40, 46]. However, it is not clear to us whether our complexity terms would still adequately capture the non-convex nature of the hypothesis set; for instance, recall that the Gaussian width is invariant under taking the convex hull. In general, the analysis of non-convex

optimization problems is very subtle, due to the possible presence of spurious local optima or saddle points. On the other hand, non-convex methods often perform better in both theory and practice, e.g., see [9]. These benefits have triggered a large amount of research in the last decade, but it is fair to say that many important issues in this field remain widely open.

5 Proofs of the main results

This part is dedicated to the proofs for Sects. 1 (provided in Sects. 5.6) and 2 (provided in Sects. 5.1–5.5).

5.1 Implications of generic Bernstein concentration

We begin with two implications of generic Bernstein concentration (see Definition 2.2), which are required for the proof of Theorem 2.10 in the next subsection, but might be also of independent interest. The proofs of both lemmas can be found in Appendix D. The first one concerns the q -th moment of the marginals of a random vector that satisfies generic Bernstein concentration; recall the notation v^* from Sect. 1.4.

Lemma 5.1 *Let $q \geq 1$. Let x be a random vector in \mathbb{R}^P that exhibits generic Bernstein concentration with respect to $(\|\cdot\|_g, \|\cdot\|_e)$, and we equip \mathbb{R}^P with the pushforward measure $\mathbb{P} \circ x^{-1}$. Then for all $v \in \mathbb{R}^P$, we have that*

$$\|v^*\|_{L^q} \lesssim q \cdot \|v\|_e + \sqrt{q} \cdot \|v\|_g.$$

The second lemma addresses the symmetrized sum of i.i.d. random vectors that satisfy generic Bernstein concentration. The resulting random vector still exhibits generic Bernstein concentration but with respect to different semi-norms.

Lemma 5.2 *Let x be a random vector in \mathbb{R}^P that exhibits generic Bernstein concentration with respect to $(\|\cdot\|_g, \|\cdot\|_e)$ and let x_1, \dots, x_n be independent copies of x . Furthermore, let $\varepsilon_1, \dots, \varepsilon_n$ be independent Rademacher random variables (also independent of the x_i). Then the rescaled symmetrized sum*

$$S_n := \frac{1}{\sqrt{n}} \sum_{i=1}^n \varepsilon_i x_i$$

exhibits generic Bernstein concentration with respect to $(C(\|\cdot\|_g + \|\cdot\|_e), \frac{C}{\sqrt{n}}\|\cdot\|_e)$, where $C > 0$ is a universal constant.

5.2 Proof of Theorem 2.10

Throughout this subsection and unless otherwise stated, we assume that the hypotheses of Theorem 2.10 are satisfied, especially Assumption 2.9. Let us recall the decompo-

sition of the excess risk from (2.2):

$$\mathcal{E}(\beta, \beta^\natural) = \underbrace{\frac{1}{n} \sum_{i=1}^n \langle x_i, \beta - \beta^\natural \rangle^2}_{=\mathcal{Q}(\beta - \beta^\natural)} + \underbrace{\frac{2}{n} \sum_{i=1}^n (\langle x_i, \beta^\natural \rangle - y_i) \langle x_i, \beta - \beta^\natural \rangle}_{=\mathcal{M}(\beta, \beta^\natural)}.$$

According to Fact 2.4, our main goal is to show that $\mathcal{E}(\beta, \beta^\natural) > 0$ for all $\beta \in K_{\beta^\natural, t}$. For this purpose, we will first treat the quadratic and multiplier process separately in Sects. 5.2.1 and 5.2.2 below. The outcome of this analysis are Proposition 5.5 and 5.15, respectively, which eventually allows us to derive the desired error bound in Sect. 5.2.3. We note that some results in Sects. 5.2.1 and 5.2.2 are presented in a slightly more general setting, considering a generic set $L \subset \mathbb{R}^p$ instead of specific subsets of K .

5.2.1 The quadratic process

We now address the issue of finding a lower bound for the quadratic process $\mathcal{Q}(\beta - \beta^\natural)$. Setting $v := \beta - \beta^\natural \in K - K$, the square root of the quadratic process takes the form

$$\sqrt{\mathcal{Q}(\beta - \beta^\natural)} = \frac{1}{\sqrt{n}} \left(\sum_{i=1}^n \langle x_i, v \rangle^2 \right)^{\frac{1}{2}}.$$

Evidently, the quadratic process describes an interaction of the input vectors x_i with the difference of two hypotheses $\beta, \beta^\natural \in K$. In this sense, it is intrinsic to K —it does not depend in any way on y , and in particular not on the model mismatch $y - \langle x, \beta^\natural \rangle$ (cf. Appendix B.2(3)).

Since $\mathcal{Q}(\beta - \beta^\natural)$ is a non-negative empirical process, it is suited for an application of the *small-ball method*. We state a version by Tropp [57, Prop. 5.1] here, but it should be emphasized that the original idea is due to Mendelson (e.g., see [31, Thm. 5.4]); recall the notion of small-ball function from Definition 2.7.

Theorem 5.3 (Small-ball method; [57, Prop 5.1]) *Let $L \subset \mathbb{R}^p$. Let x be a random vector in \mathbb{R}^p and let x_1, \dots, x_n be independent copies of x . Then for every $\theta > 0$ and $u > 0$, we have that*

$$\inf_{v \in L} \left(\sum_{i=1}^n \langle x_i, v \rangle^2 \right)^{\frac{1}{2}} \geq \theta \cdot \sqrt{n} \cdot \mathcal{Q}_{2\theta}(L, x) - 2W_n(L, x) - \theta \cdot u$$

with probability at least $1 - \exp(-u^2/2)$, where

$$W_n(L, x) := \mathbb{E} \left[\sup_{v \in L} \left\langle \frac{1}{\sqrt{n}} \sum_{i=1}^n \varepsilon_i x_i, v \right\rangle \right]$$

is the empirical width of L with independent Rademacher random variables $\varepsilon_1, \dots, \varepsilon_n$.

In fact, Theorem 5.3 is a remarkable result because it holds true without a strong tail assumption on x . However, its significance is closely linked to finding an appropriate (upper) bound for the empirical width $W_n(L, x)$, which is usually not a simple task. In the specific context of this paper, where x exhibits generic Bernstein concentration, the following *generic chaining bound* by Talagrand will prove useful; recall the γ -functional from Definition 2.3.

Theorem 5.4 [51, Thm. 2.2.23] *Let d_1, d_2 be two pseudo-metrics on a set L . Consider a real-valued stochastic process $(X_v)_{v \in L}$ which satisfies the increment condition*

$$\mathbb{P}(|X_{v_1} - X_{v_2}| \geq t) \leq 2 \exp \left(- \min \left\{ \frac{t^2}{d_2(v_1, v_2)^2}, \frac{t}{d_1(v_1, v_2)} \right\} \right)$$

for all $v_1, v_2 \in L$ and all $t > 0$. Then, we have that

$$\mathbb{E} \left[\sup_{v_1, v_2 \in L} |X_{v_1} - X_{v_2}| \right] \lesssim \gamma_1(L, d_1) + \gamma_2(L, d_2).$$

Moreover, if $(X_v)_{v \in L}$ is symmetric, we have that

$$\mathbb{E} \left[\sup_{v \in L} X_v \right] \lesssim \gamma_1(L, d_1) + \gamma_2(L, d_2).$$

An appropriate combination of Theorems 5.3 and 5.4 leads to the following lower bound for the quadratic process:

Proposition 5.5 *Let $x, x_1, \dots, x_n, K \subset \mathbb{R}^p$, and $\tau > 0$ be as in Assumption 2.9. For $t > 0$, let $L \subset (K - K) \cap t\mathbb{S}^{p-1}$. Then for every $u > 0$, we have that*

$$\inf_{v \in L} \left(\sum_{i=1}^n \langle x_i, v \rangle^2 \right)^{\frac{1}{2}} \geq t \cdot \left(\sqrt{n} \cdot \tau \cdot Q_{2\tau}(K^\Delta, x) - C_Q \cdot q_{t,n}^{(g,e)}(L) - \tau \cdot u \right)$$

with probability at least $1 - \exp(-u^2/2)$, where $C_Q > 0$ is a universal constant.

Proof Let $\theta := t \cdot \tau$. Then, Theorem 5.3 states that

$$\inf_{v \in L} \left(\sum_{i=1}^n \langle x_i, v \rangle^2 \right)^{\frac{1}{2}} \geq t \cdot \left(\sqrt{n} \cdot \tau \cdot Q_{2\theta}(L, x) - \frac{2}{t} \cdot W_n(L, x) - \tau \cdot u \right)$$

holds true with probability at least $1 - \exp(-u^2/2)$. The claim of Proposition 5.5 follows from the bounds on $Q_{2\theta}(L, x)$ and $W_n(L, x)$ that we establish in the following.

Lower bound for $Q_{2\theta}$: We have that

$$\begin{aligned} Q_{2\theta}(L, x) &= \inf_{v \in L} \mathbb{P}(|\langle x, v \rangle| \geq 2\theta) \\ &\geq \inf_{v \in (K-K) \cap (t\mathbb{S}^{p-1})} \mathbb{P}(|\langle x, v \rangle| \geq 2\theta) \end{aligned}$$

$$\begin{aligned} &\geq \inf_{\tilde{v} \in \text{span}(K-K) \cap \mathbb{S}^{p-1}} \mathbb{P}(|\langle x, t\tilde{v} \rangle| \geq 2t \cdot \tau) \\ &= \inf_{\tilde{v} \in K^\Delta} \mathbb{P}(|\langle x, \tilde{v} \rangle| \geq 2\tau) = Q_{2\tau}(K^\Delta, x). \end{aligned}$$

Upper bound for W_n : According to the definition of the local q-complexity, there exists a set $\tilde{S} \subset \mathbb{R}^p$ with

$$\text{conv}(\tilde{S}) \supset L \cap t\mathbb{S}^{p-1} (= L)$$

that satisfies

$$\frac{\gamma_1(\tilde{S}, \|\cdot\|_e)}{\sqrt{n}} + \gamma_2(\tilde{S}, \|\cdot\|_g + \|\cdot\|_e) \leq 2t \cdot q_{t,n}^{(g,e)}(L). \tag{5.1}$$

Now let $v \in L$. Since $L \subset \text{conv}(\tilde{S})$, the point v can be expressed as a convex combination in \tilde{S} :

$$v = \sum_{j=1}^M \lambda_j s_j, \quad \text{where } \lambda_j \geq 0, \quad \sum_{j=1}^M \lambda_j = 1, \quad s_j \in \tilde{S}. \tag{5.2}$$

Conditioning on the random variables ε_i, x_i , the function

$$h : \mathbb{R}^p \rightarrow \mathbb{R}, \quad w \mapsto \left\langle \frac{1}{\sqrt{n}} \sum_{i=1}^n \varepsilon_i x_i, w \right\rangle$$

is linear. Hence, we have that

$$h(v) = \sum_{j=1}^M \lambda_j h(s_j) \in \text{conv}(\{h(s_1), \dots, h(s_M)\}).$$

In particular, at least one of the $h(s_i)$ is not smaller than $h(v)$. This implies

$$\sup_{v \in L} \left\langle \frac{1}{\sqrt{n}} \sum_{i=1}^n \varepsilon_i x_i, v \right\rangle \leq \sup_{v \in \tilde{S}} \left\langle \frac{1}{\sqrt{n}} \sum_{i=1}^n \varepsilon_i x_i, v \right\rangle$$

and therefore $W_n(L, x) \leq W_n(\tilde{S}, x)$. To obtain an upper bound for $W_n(\tilde{S}, x)$, we consider the associated stochastic process

$$(X_v)_{v \in \tilde{S}} := \left(\left\langle \frac{1}{\sqrt{n}} \sum_{i=1}^n \varepsilon_i x_i, v \right\rangle \right)_{v \in \tilde{S}}$$

and intend to apply Theorem 5.4: Since the x_i and the ε_i are independent, the distribution of $(X_v)_{v \in \tilde{S}}$ only depends on the individual distributions of the x_i and the ε_i . Observing that

$$(-X_v)_{v \in \tilde{S}} = \left(\left(\frac{1}{\sqrt{n}} \sum_{i=1}^n (-\varepsilon_i)x_i, v \right) \right)_{v \in \tilde{S}}$$

and that $-\varepsilon_i$ has the same distribution as ε_i for each i , we conclude that $(X_v)_{v \in \tilde{S}}$ is indeed symmetric. Regarding the increment condition, let $v_1, v_2 \in \tilde{S}$ and set $v := v_1 - v_2$. Then, Lemma 5.2 implies that

$$\mathbb{P}(|X_{v_1} - X_{v_2}| \geq t) \leq 2 \exp \left(- \min \left\{ \frac{t^2}{C^2(\|v\|_g + \|v\|_e)^2}, \frac{\sqrt{nt}}{C\|v\|_e} \right\} \right).$$

Finally, Theorem 5.4 yields

$$\begin{aligned} W_n(\tilde{S}, x) &\lesssim \gamma_1(\tilde{S}, \frac{1}{\sqrt{n}}\|\cdot\|_e) + \gamma_2(\tilde{S}, \|\cdot\|_g + \|\cdot\|_e) \\ &= \frac{\gamma_1(\tilde{S}, \|\cdot\|_e)}{\sqrt{n}} + \gamma_2(\tilde{S}, \|\cdot\|_g + \|\cdot\|_e) \stackrel{(5.1)}{\leq} 2t \cdot q_{t,n}^{(g,e)}(L), \end{aligned}$$

and therefore $W_n(L, x) \leq W_n(\tilde{S}, x) \lesssim t \cdot q_{t,n}^{(g,e)}(L)$. □

5.2.2 The multiplier process

We now turn our attention to the multiplier process. Setting $v := \beta - \beta^\natural$ and $\xi_i := \langle x_i, \beta^\natural \rangle - y_i$, the process takes the form

$$\mathcal{M}(\beta, \beta^\natural) = \frac{2}{n} \sum_{i=1}^n \xi_i \cdot \langle x_i, v \rangle.$$

Unlike the quadratic process, the multiplier process is not intrinsic to the hypothesis set K , but (empirically) describes an interaction of the difference of two hypotheses $\beta, \beta^\natural \in K$ with the model mismatch $\xi := \langle x, \beta^\natural \rangle - y$ (cf. Appendix B.2(3)).

In order to control the multiplier process, we adapt another result by Mendelson [32], which is based on a refined chaining approach: Instead of applying the traditional generic chaining to the function class $\{\xi \cdot \langle \cdot, \beta \rangle \mid \beta \in K\}$, Mendelson isolates the effect of the multiplier term ξ , which leads to a bound in terms of $\|\xi\|_{L^q}$ and geometric properties of the class $\{\langle \cdot, \beta \rangle \mid \beta \in K\}$. In fact, his result holds true for more general (non-linear) function classes, but in view of the objectives of this article, we only recite the special case of linear functions. In order to state this result, several definitions are required.

Definition 5.6 [32, Def. 1.6] For a real-valued random variable Z and $q \geq 1$, we define the (q)-norm by

$$\|Z\|_{(q)} := \sup_{1 \leq r \leq q} \frac{\|Z\|_{L^r}}{\sqrt{r}}.$$

It is worth comparing the above definition to the moment characterization of sub-Gaussian variables (see Proposition A.1(ii)). Mendelson [32, p. 3658] remarks that the (q) -norm “*measure[s] the subgaussian behaviour of the functions involved, but only up to a fixed level, rather than at every level*”.

Definition 5.7 Let L be a set. We call a sequence $(L_s)_{s \in \mathbb{N}} \subset 2^L$ of subsets of L an *admissible approximation sequence* if $|L_0| = 1$ and $|L_s| \leq 2^{2^s}$ for $s \geq 1$.

The following definition introduces a relative of Talagrand’s γ -functional (see Definition 2.3). For this purpose, also recall the notation of dual vectors from Sect. 1.4; more precisely, we equip \mathbb{R}^p with the pushforward measure $\mathbb{P} \circ x^{-1}$ of a (generic) random vector $x \in \mathbb{R}^p$, so that for every $v \in \mathbb{R}^p$, we have $\|v^*\|_{L^q}^q = \mathbb{E}[|\langle x, v \rangle|^q]$.

Definition 5.8 [32, Def. 1.7] For $L \subset \mathbb{R}^p$ and $u \geq 1$, we define

$$\tilde{\Lambda}_u(L, x) := \inf \left\{ \sup_{v \in L} \|(\pi_0 v)^*\|_{(u^2)} + \sup_{v \in L} \sum_{s \geq 0} 2^{s/2} \cdot \|v^* - (\pi_s v)^*\|_{(u^2 2^s)} \right\},$$

where the infimum is taken over all admissible approximation sequences $(L_s)_{s \in \mathbb{N}}$ and $(\pi_s v)^*$ is a nearest point to v^* in $(L_s)^*$ with respect to the $(u^2 2^s)$ -norm.

With these definitions at hand, we can now state Mendelson’s result, which provides a powerful concentration inequality for multiplier processes. We emphasize that the feature vector x and the multiplier ξ are *not* necessarily independent here, which is crucial for our analysis and an important difference to related results in the literature, e.g., see [22].

Theorem 5.9 [32, Thm. 1.9] *Let $L \subset \mathbb{R}^p$ and let $(x, \xi) \in \mathbb{R}^p \times \mathbb{R}$ be a random pair such that $\|\xi\|_{L^q} < \infty$ for some $q > 2$. We assume that $(x_1, \xi_1), \dots, (x_n, \xi_n)$ are independent copies of (x, ξ) . Then there exist constants $C_0, C_1, \dots, C_4 > 0$ (only depending on q) such that for every $w, u > C_0$, the following holds true with probability at least $1 - C_1 \cdot w^{-q} \cdot n^{-(q/2)+1} \cdot \log_q(n) - 4 \exp(-C_2 \cdot u^2)$:*

$$\sup_{v \in L} \left| \sum_{i=1}^n (\xi_i \cdot \langle x_i, v \rangle - \mathbb{E}[\xi \cdot v^*]) \right| \leq C_3 \cdot w \cdot u \cdot \sqrt{n} \cdot \|\xi\|_{L^q} \cdot \tilde{\Lambda}_{C_4 u}(L).$$

The term $C_1 \cdot w^{-q} \cdot n^{-(q/2)+1} \cdot \log_q(n)$ in the probability of success arises from a concentration inequality for the random vector $(\xi_i)_{i=1}^n \in \mathbb{R}^n$, for which we only assume that the q -th moment of its components exists for some $q > 2$. In fact, better rates can be achieved by more restrictive assumptions on the tails of ξ . For example, Mendelson proves a sub-Gaussian variant of Theorem 5.9 using Bernstein’s inequality (see [32, Thm. 4.4]). If we assume that ξ is just sub-exponential (as in Assumption 2.9), Bernstein’s inequality cannot be applied to the squared coordinates appearing in the Euclidean norm of $(\xi_i)_{i=1}^n$. However, the following recent result of Götze et al. [21] allows us to derive a concentration inequality in the sub-exponential case:

Theorem 5.10 [21, Prop. 1.1] *Let X_1, \dots, X_n be independent, centered random variables with $\sigma_i^2 := \mathbb{E}[X_i^2] < \infty$ and $\|X_i\|_{\psi_\alpha} \leq R$ for some $\alpha \in (0, 1] \cup \{2\}$.¹⁶ Let $B = [b_{ij}] \in \mathbb{R}^{n \times n}$ be a symmetric matrix. Then there exists a universal constant $C > 0$ such that for every $t > 0$, we have that*

$$\mathbb{P}\left(\left|\sum_{i,j=1}^n b_{ij} X_i X_j - \sum_{i=1}^n \sigma_i^2 b_{ii}\right| \geq t\right) \leq 2 \exp\left(-C \cdot \min\left\{\frac{t^2}{R^4 \|B\|_F^2}, \left(\frac{t}{R^2 \|B\|_{\text{op}}}\right)^{\frac{\alpha}{2}}\right\}\right).$$

Corollary 5.11 *Let ξ_1, \dots, ξ_n be i.i.d. sub-exponential random variables. Then there exists a universal constant $C > 0$ such that with probability at least $1 - 2 \exp(-C \cdot \sqrt{n})$, we have that*

$$\|(\xi_i)_{i=1}^n\|_2 = \left(\sum_{i=1}^n \xi_i^2\right)^{\frac{1}{2}} \lesssim \sqrt{n} \cdot \|\xi_1\|_{\psi_1}.$$

Proof We apply Theorem 5.10 for $X_i := \xi_i - \mathbb{E}[\xi_i]$, $B := I_n$, $t := R^2 \cdot n$, and obtain

$$\begin{aligned} \mathbb{P}(E) &:= \mathbb{P}\left(\left|\sum_{i=1}^n (\xi_i - \mathbb{E}[\xi_i])^2 - n \cdot \sigma_1^2\right| \geq n \cdot \|\xi_1 - \mathbb{E}[\xi_1]\|_{\psi_1}^2\right) \\ &\leq 2 \exp(-C \cdot \min\{n, \sqrt{n}\}) = 2 \exp(-C \cdot \sqrt{n}). \end{aligned}$$

Let us assume that the complement of the event E has occurred. Then, using Proposition A.1(ii), it follows that

$$\sum_{i=1}^n (\xi_i - \mathbb{E}[\xi_i])^2 < n \cdot (\|\xi_1 - \mathbb{E}[\xi_1]\|_{\psi_1}^2 + \sigma_1^2) \lesssim n \cdot \|\xi_1 - \mathbb{E}[\xi_1]\|_{\psi_1}^2 \lesssim n \cdot \|\xi_1\|_{\psi_1}^2,$$

where the last step is due to

$$\|X - \mathbb{E}[X]\|_{\psi_\alpha} \lesssim \|X\|_{\psi_\alpha}$$

for $\alpha \in \{1, 2\}$; see [58, Lem. 2.6.8] for a proof, which also works for $\alpha = 1$. Consequently, we have that

$$\|(\xi_i)_{i=1}^n\|_2 \leq \|(\xi_i - \mathbb{E}[\xi_i])_{i=1}^n\|_2 + \|(\mathbb{E}[\xi_i])_{i=1}^n\|_2 \lesssim \sqrt{n} \cdot \|\xi_1\|_{\psi_1}.$$

□

The bound of Corollary 5.11 leads to the following sub-exponential version of Theorem 5.9:

¹⁶ As discussed by Götze et al. [21], our Definition 1.2 can be extended to the more general case $\alpha > 0$, leading to the notion of α -sub-exponential random variables. But note that if $0 < \alpha < 1$, the exponential Orlicz “norm” $\|\cdot\|_{\psi_\alpha}$ violates the triangle inequality.

Theorem 5.12 *Let $L \subset \mathbb{R}^p$ and let $(x, \xi) \in \mathbb{R}^p \times \mathbb{R}$ be a random pair such that $\|\xi\|_{\psi_1} < \infty$. We assume that $(x_1, \xi_1), \dots, (x_n, \xi_n)$ are independent copies of (x, ξ) . Then there exist universal constants $C, C', \tilde{C} > 0$ such that for every $u \geq 8$, the following holds true with probability at least $1 - 2 \exp(-C \cdot \sqrt{n}) - 4 \exp(-C \cdot u^2)$:*

$$\sup_{v \in L} \left| \sum_{i=1}^n (\xi_i \cdot \langle x_i, v \rangle - \mathbb{E}[\xi \cdot v^*]) \right| \leq C' \cdot u \cdot \sqrt{n} \cdot \|\xi\|_{\psi_1} \cdot \tilde{\Lambda}_{\tilde{C}u}(L, x).$$

Proof Analogously to the proof of [32, Thm. 4.4], it is enough to adapt the last step of the proof of [32, Thm. 1.9]. To this end, we set the variables arising in the proof of [32, Thm. 4.4] to the values $q := 6$ and $w := 1$, which entails $r = r' = 2$ and $q_1 = 8$. Then, with probability at least $1 - 2 \exp(-C \cdot \sqrt{n})$, Corollary 5.11 implies that

$$\|(\xi_i)_{i=1}^n\|_2 \lesssim \sqrt{n} \cdot \|\xi_1\|_{\psi_1}.$$

Since $\|\xi\|_{L^6} \lesssim \|\xi\|_{\psi_1}$, we also have that

$$\left(\sum_{i \geq j_0} (\xi_i^*)^{2r} \right)^{\frac{1}{2r}} = \left(\sum_{i \geq j_0} (\xi_i^*)^4 \right)^{\frac{1}{4}} \lesssim \|\xi\|_{L^6} \cdot n^{1/4} \lesssim \|\xi\|_{\psi_1} \cdot n^{1/4},$$

with probability at least $1 - 2 \exp(-C \cdot u^2)$, where ξ^* and j_0 are objects defined in the proof of [32, Thm. 1.9]. The rest of the proof remains unchanged. \square

The following lemma is a centerpiece of our statistical analysis, as it allows us to control the complexity term $\tilde{\Lambda}_u(L, x)$ via generic Bernstein concentration:

Lemma 5.13 *Let $L \subset \mathbb{R}^p$ with $0 \in \text{conv}(L)$ and $u \geq 1$. Let x be a random vector in \mathbb{R}^p that exhibits generic Bernstein concentration with respect to $(\|\cdot\|_g, \|\cdot\|_e)$. Then, we have that*

$$\tilde{\Lambda}_u(L, x) \lesssim u \cdot \gamma_1(L, \|\cdot\|_e) + \gamma_2(L, \|\cdot\|_g).$$

Proof From Lemma 5.1, we obtain

$$\|v^*\|_{(u^2 2^{2^s})} = \sup_{1 \leq q \leq u^2 2^s} \frac{\|v^*\|_{L^q}}{\sqrt{q}} \lesssim \sup_{1 \leq q \leq u^2 2^s} \frac{q \cdot \|v\|_e + \sqrt{q} \cdot \|v\|_g}{\sqrt{q}} = u 2^{s/2} \cdot (\|v\|_e + \|v\|_g) \tag{5.3}$$

for every $v \in \mathbb{R}^p$. Adopting the notation from [32, Def. 1.7], we set

$$\Lambda_u(L, x) := \inf_{v \in L} \sup_{s \geq 0} \sum_{s \geq 0} 2^{s/2} \cdot \|v^* - (\pi_s v)^*\|_{(u^2 2^{2^s})},$$

which implies that

$$\tilde{\Lambda}_u(L, x) \leq \Lambda_u(L, x) + \sup_{v \in L} \|v^*\|_{(u^2)}. \tag{5.4}$$

According to (5.3), the second summand of (5.4) can be bounded as follows:

$$\begin{aligned} \sup_{v \in L} \|v^*\|_{(u^2)} &\lesssim \sup_{v \in L} (u \cdot \|v\|_e + \|v\|_g) \leq \sup_{v \in L} (u \cdot \|v\|_e) + \sup_{v \in L} \|v\|_g \\ &\stackrel{0 \in \text{conv}(L)}{\leq} u \cdot \Delta_e(L) + \Delta_g(L) \leq u \cdot \gamma_1(L, \|\cdot\|_e) + \gamma_2(L, \|\cdot\|_g), \end{aligned}$$

where $\Delta_e(L)$ and $\Delta_g(L)$ are the diameters of L with respect to $\|\cdot\|_e$ and $\|\cdot\|_g$, respectively. To handle the first summand of (5.4), we apply (5.3) once again:

$$\begin{aligned} \Lambda_u(L, x) &= \inf_{v \in L} \sup_{s \geq 0} \sum 2^{s/2} \cdot \|v^* - (\pi_s v)^*\|_{(u^2 2^s)} \\ &\leq \inf_{v \in L} \sup_{s \geq 0} \sum 2^{s/2} \cdot \|v^* - (\pi_s^+ v)^*\|_{(u^2 2^s)} \\ &\lesssim \inf_{v \in L} \sup_{s \geq 0} \sum (u 2^s \cdot \|v - \pi_s^+ v\|_e + 2^{s/2} \cdot \|v - \pi_s^+ v\|_g), \end{aligned}$$

where $\pi_s^+ : L \rightarrow L_s$ is an arbitrary map (depending on the respective admissible approximation sequence $(L_s)_{s \in \mathbb{N}}$ indexed by the infimum); note that we can indeed replace π_s by π_s^+ here, since by definition, $(\pi_s)_{s \in \mathbb{N}}$ is an optimal (functional-minimizing) sequence of projections with respect to the $(u^2 2^s)$ -norms.

We now show that the above expression is upper bounded by

$$5 \cdot \left(u \cdot \gamma_1(L, \|\cdot\|_e) + \gamma_2(L, \|\cdot\|_g) \right),$$

which would imply the claim of Lemma 5.13. For this purpose, let $(\mathcal{E}_s)_{s \in \mathbb{N}}$ and $(\mathcal{G}_s)_{s \in \mathbb{N}}$ be two admissible partition sequences which approximate $\gamma_1(L, \|\cdot\|_e)$ and $\gamma_2(L, \|\cdot\|_g)$ up to a factor of 2, respectively. Furthermore, let $(\mathcal{F}_s)_{s \in \mathbb{N}}$ be given by $\mathcal{F}_0 = \{L\}$ and

$$\mathcal{F}_s := \{E \cap G \mid E \in \mathcal{E}_{s-1}, G \in \mathcal{G}_{s-1}\} \quad \text{for } s \geq 1. \tag{5.5}$$

It is not hard to see that $(\mathcal{F}_s)_{s \in \mathbb{N}}$ is indeed an admissible partition sequence.

Next, we use the sequence $(\mathcal{F}_s)_{s \in \mathbb{N}}$ to construct an admissible approximation sequence $(L_s)_{s \in \mathbb{N}}$ and a corresponding sequence of maps $(\pi_s^\circ)_{s \in \mathbb{N}}$: For each $s \in \mathbb{N}$, the set $L_s \subset L$ is obtained by selecting exactly one (arbitrary) point v_F from every $F \in \mathcal{F}_s$, while π_s° maps every point in F to the respective v_F . This construction ensures that for $s \geq 1$ and $v \in L$, we have

$$\|v - \pi_s^\circ v\|_e \leq \Delta_e(F_s(v)) \leq \Delta_e(E_{s-1}(v))$$

and

$$\|v - \pi_s^\circ v\|_g \leq \Delta_g(F_s(v)) \leq \Delta_g(G_{s-1}(v)).$$

This implies

$$\begin{aligned}
 \Delta_u(L, x) &\lesssim \sup_{v \in L} \sum_{s \geq 0} (u 2^s \cdot \|v - \pi_s^\circ v\|_e + 2^{s/2} \cdot \|v - \pi_s^\circ v\|_g) \\
 &\leq u \cdot \sup_{v \in L} \sum_{s \geq 0} 2^s \cdot \|v - \pi_s^\circ v\|_e + \sup_{v \in L} \sum_{s \geq 0} 2^{s/2} \cdot \|v - \pi_s^\circ v\|_g \\
 &\leq u \cdot \left(\Delta_e(L) + \sup_{v \in L} \sum_{s \geq 1} 2^s \cdot \Delta_e(E_{s-1}(v)) \right) + \Delta_g(L) + \sup_{v \in L} \sum_{s \geq 1} 2^{s/2} \cdot \Delta_g(G_{s-1}(v)) \\
 &= u \cdot \left(\Delta_e(L) + \sup_{v \in L} \sum_{s \geq 0} 2^{s+1} \cdot \Delta_e(E_s(v)) \right) + \Delta_g(L) + \sup_{v \in L} \sum_{s \geq 0} 2^{(s+1)/2} \cdot \Delta_g(G_s(v)) \\
 &\leq u \cdot (1 + 4) \cdot \gamma_1(L, \|\cdot\|_e) + (1 + 2\sqrt{2}) \cdot \gamma_2(L, \|\cdot\|_g),
 \end{aligned}$$

where we have used in the last line that $(\mathcal{E}_s)_{s \in \mathbb{N}}$ and $(\mathcal{G}_s)_{s \in \mathbb{N}}$ approximate $\gamma_1(L, \|\cdot\|_e)$ and $\gamma_2(L, \|\cdot\|_g)$ up to a factor of 2, respectively. \square

Remark 5.14 The proof of Lemma 5.13 is inspired by [32, Subsec. 4.3], where the upper bound $\tilde{\Lambda}_u(L, x) \lesssim u \cdot \gamma_1(L, \|\cdot\|_\infty) + \gamma_2(L, \|\cdot\|_2)$ is derived under the assumption that x obeys an unconditional, isotropic, log-concave distribution. This assumption implies that x is stochastically dominated by a random vector with i.i.d. standard exponential coordinates, which enables a bound for $\|v^*\|_{(q)}$ in terms of $\|v\|_2$ and $\|v\|_\infty$ (cf. Sect. 3.3).

The estimate from Lemma 5.13 leads us to our final result for the multiplier process:

Proposition 5.15 *Let $L \subset t\mathbb{S}^{p-1}$ for some $t > 0$ and let $(x, \xi) \in \mathbb{R}^p \times \mathbb{R}$ be a random pair such that $\|\xi\|_{\psi_1} < \infty$ and x exhibits generic Bernstein concentration with respect to $(\|\cdot\|_g, \|\cdot\|_e)$. We assume that $(x_1, \xi_1), \dots, (x_n, \xi_n)$ are independent copies of (x, ξ) . Then there exist universal constants $C, C' > 0$ such that for every $u \geq 8$, the following holds true with probability at least $1 - 2 \exp(-C \cdot \sqrt{n}) - 4 \exp(-C \cdot u^2)$:*

$$\sup_{v \in L} \left| \sum_{i=1}^n (\xi_i \cdot \langle x_i, v \rangle - \mathbb{E}[\xi \cdot v^*]) \right| \leq C' \cdot t \cdot u^2 \cdot \sqrt{n} \cdot \|\xi\|_{\psi_1} \cdot m_t^{(g,e)}(L).$$

Proof According to the definition of the local m-complexity, there exists a set $\tilde{S} \subset \mathbb{R}^p$ with

$$\text{conv}(\tilde{S}) \supset (L \cap t\mathbb{S}^{p-1}) \cup \{0\} (= L \cup \{0\})$$

that satisfies

$$\gamma_1(\tilde{S}, \|\cdot\|_e) + \gamma_2(\tilde{S}, \|\cdot\|_g) \leq 2t \cdot m_t^{(g,e)}(L). \tag{5.6}$$

By Theorem 5.12, with probability at least $1 - 2 \exp(-C \cdot \sqrt{n}) - 4 \exp(-C \cdot u^2)$, we have that

$$\sup_{v \in \tilde{S}} \left| \sum_{i=1}^n (\xi_i \cdot \langle x_i, v \rangle - \mathbb{E}[\xi \cdot v^*]) \right| \lesssim u \cdot \sqrt{n} \cdot \|\xi\|_{\psi_1} \cdot \tilde{\Lambda}_{\tilde{C}u}(\tilde{S}, x). \tag{5.7}$$

Now let $v \in L$. Since $L \subset \text{conv}(\tilde{S})$, the point v can be expressed as a convex combination of points $s_1, \dots, s_M \in \tilde{S}$ as in (5.2). Conditioning on the random variables x_i and ξ_i , the function

$$h : \mathbb{R}^p \rightarrow \mathbb{R}, \quad w \mapsto \left| \sum_{i=1}^n (\xi_i \cdot \langle x_i, w \rangle - \mathbb{E}[\xi \cdot w^*]) \right| = \left| \left\langle \sum_{i=1}^n (\xi_i x_i - \mathbb{E}[\xi x]), w \right\rangle \right|$$

is a composition of a linear function and the convex function $z \mapsto |z|$. Hence, h is convex and we can apply Jensen’s inequality to obtain

$$h(v) = h \left(\sum_{j=1}^M \lambda_j s_j \right) \leq \sum_{j=1}^M \lambda_j h(s_j) \leq \sum_{j=1}^M \left(\lambda_j \cdot \sup_{w \in \tilde{S}} h(w) \right) = \sup_{w \in \tilde{S}} h(w).$$

Since $v \in L$ was arbitrarily chosen, we can conclude that the following bound holds true if the event from (5.7) has occurred:

$$\sup_{v \in L} \left| \sum_{i=1}^n (\xi_i \cdot \langle x_i, v \rangle - \mathbb{E}[\xi \cdot v^*]) \right| \lesssim u \cdot \sqrt{n} \cdot \|\xi\|_{\psi_1} \cdot \tilde{\Lambda}_{\tilde{C}_u}(\tilde{S}, x).$$

Finally, Lemma 5.13 implies

$$\tilde{\Lambda}_{\tilde{C}_u}(\tilde{S}, x) \lesssim \tilde{C} \cdot u \cdot \left(\gamma_1(\tilde{S}, \|\cdot\|_e) + \gamma_2(\tilde{S}, \|\cdot\|_g) \right) \stackrel{(5.6)}{\lesssim} u \cdot t \cdot m_t^{(g,e)}(L),$$

where we have also used that $u \geq 8 > 1$. □

5.2.3 Controlling the excess risk

With the results of Propositions 5.5 and 5.15 at hand, we are now ready to prove Theorem 2.10. Let us first consider the case $t > 0$. According to Fact 2.4, it suffices to show that $\mathcal{E}(\beta, \beta^\natural) > 0$ for all $\beta \in K_{\beta^\natural, t}$. Remarkably, this argument is actually the only point in our proof where we rely on the convexity of the hypothesis set K . The remainder of the proof is divided into several substeps.

Step 1 (quadratic process): Applying Proposition 5.5 to $L := K_{\beta^\natural, t} - \beta^\natural = (K - \beta^\natural) \cap t\mathbb{S}^{p-1}$, the following holds with probability at least $1 - \exp(-u^2/2)$: For every $\beta \in K_{\beta^\natural, t}$, we have that

$$\begin{aligned} \sqrt{Q(\beta - \beta^\natural)} &= \left(\frac{1}{n} \sum_{i=1}^n \langle x_i, \beta - \beta^\natural \rangle^2 \right)^{\frac{1}{2}} \\ &\geq t \cdot \left(\tau \cdot Q_{2\tau}(K^\Delta, x) - \frac{C_Q \cdot q_{t,n}^{(g,e)}(K - \beta^\natural) + \tau \cdot u}{\sqrt{n}} \right) \\ &\geq \frac{1}{2} \cdot t \cdot \tau \cdot Q_{2\tau}(K^\Delta, x), \end{aligned}$$

where the last step follows from the condition (2.4) for an appropriate hidden constant.

Step 2 (multiplier process): Applying Proposition 5.15 to $L := K_{\beta^\natural, t} - \beta^\natural = (K - \beta^\natural) \cap t\mathbb{S}^{p-1}$ and $\xi := \langle x, \beta^\natural \rangle - y$, the following holds with probability at least $1 - 2 \exp(-C \cdot \sqrt{n}) - 4 \exp(-C \cdot u^2)$: For every $\beta \in K_{\beta^\natural, t}$, we have that

$$\begin{aligned} \frac{1}{2} \cdot \mathcal{M}(\beta, \beta^\natural) &= \frac{1}{n} \sum_{i=1}^n (\langle x_i, \beta^\natural \rangle - y_i) \langle x_i, \beta - \beta^\natural \rangle \\ &\geq \mathbb{E}[(\langle x, \beta^\natural \rangle - y) \langle x, \beta - \beta^\natural \rangle] - C' \cdot u^2 \cdot \|\langle x, \beta^\natural \rangle - y\|_{\psi_1} \cdot \frac{t \cdot m_t^{(g,e)}(K - \beta^\natural)}{\sqrt{n}} \\ &= -t \cdot \left(\mathbb{E}\left[\langle y - \langle x, \beta^\natural \rangle, x \rangle \left\langle x, \frac{\beta - \beta^\natural}{t} \right\rangle\right] + C' \cdot u^2 \cdot \|\langle x, \beta^\natural \rangle - y\|_{\psi_1} \cdot \frac{m_t^{(g,e)}(K - \beta^\natural)}{\sqrt{n}} \right) \\ &\geq -t \cdot \left(\rho_t(\beta^\natural) + C' \cdot u^2 \cdot \sigma(\beta^\natural) \cdot \frac{m_t^{(g,e)}(K - \beta^\natural)}{\sqrt{n}} \right) \\ &\geq -t \cdot \max\{1, C'\} \cdot \left(\rho_t(\beta^\natural) + u^2 \cdot \sigma(\beta^\natural) \cdot \frac{m_t^{(g,e)}(K - \beta^\natural)}{\sqrt{n}} \right), \end{aligned}$$

where the second inequality is due to

$$\begin{aligned} \mathbb{E}\left[\langle y - \langle x, \beta^\natural \rangle, x \rangle \left\langle x, \frac{\beta - \beta^\natural}{t} \right\rangle\right] &= \left\langle \mathbb{E}[(y - \langle x, \beta^\natural \rangle)x], \underbrace{\frac{\beta - \beta^\natural}{t}}_{\in \frac{1}{t}(K - \beta^\natural) \cap \mathbb{S}^{p-1} = K^t} \right\rangle \leq \rho_t(\beta^\natural). \end{aligned}$$

If the aforementioned event has occurred and if t satisfies the condition (2.5) for an appropriate hidden constant, it follows that

$$\mathcal{M}(\beta, \beta^\natural) \geq -\frac{t^2}{8} \cdot (\tau \cdot Q_{2\tau}(K^\Delta, x))^2 \quad \text{for all } \beta \in K_{\beta^\natural, t}.$$

Step 3 (excess risk): Finally, we assume that the events from Step 1 and Step 2 have occurred jointly, which indeed happens with probability at least $1 - 2 \exp(-C \cdot \sqrt{n}) - 5 \exp(-C \cdot u^2)$ for an appropriately chosen constant $C > 0$. Then, we obtain

$$\begin{aligned} \mathcal{E}(\beta, \beta^\natural) &= Q(\beta - \beta^\natural) + \mathcal{M}(\beta, \beta^\natural) \\ &\geq \frac{t^2}{4} \cdot (\tau \cdot Q_{2\tau}(K^\Delta, x))^2 - \frac{t^2}{8} \cdot (\tau \cdot Q_{2\tau}(K^\Delta, x))^2 > 0 \end{aligned}$$

for all $\beta \in K_{\beta^\natural, t}$, which concludes the proof for $t > 0$. It remains to consider the case $t = 0$.

Step 4 ($t = 0$): In this case, $q_{t,n}^{(g,e)}(K - \beta^\natural)$ and $m_t^{(g,e)}(K - \beta^\natural)$ correspond to the conic complexities from Definition 2.11.

Applying Propositions 5.5 and 5.15 simultaneously (as in the preceding steps) to $L := \text{cone}(K - \beta^\natural) \cap \mathbb{S}^{p-1}$ and radius $\tilde{t} := 1$, we have that with probability at least $1 - 2 \exp(-C \cdot \sqrt{n}) - 5 \exp(-C \cdot u^2)$, both

$$\inf_{v \in L} \mathcal{Q}(v) \stackrel{(2.4)}{\geq} \frac{1}{4} \cdot (\tau \cdot \mathcal{Q}_{2\tau}(K^\Delta, x))^2 > 0$$

and

$$\inf_{v \in L} \mathcal{M}(v + \beta^\natural, \beta^\natural) \geq -1 \cdot \max\{1, C'\} \cdot \left(\rho_0(\beta^\natural) + u^2 \cdot \sigma(\beta^\natural) \cdot \frac{m_0^{(g,e)}(K - \beta^\natural)}{\sqrt{n}} \right) \stackrel{(2.5)}{\geq} 0,$$

where we have also used that $q_{1,n}^{(g,e)}(L) = q_{0,n}^{(g,e)}(K - \beta^\natural)$ and $m_1^{(g,e)}(L) = m_0^{(g,e)}(K - \beta^\natural)$. Finally, let us assume that this event has occurred and let $\beta \in K \setminus \{\beta^\natural\}$. Then, we have

$$\mathcal{E}(\beta, \beta^\natural) = \|\beta - \beta^\natural\|_2^2 \cdot \underbrace{\mathcal{Q}\left(\frac{\beta - \beta^\natural}{\|\beta - \beta^\natural\|_2}\right)}_{\in L} + \|\beta - \beta^\natural\|_2 \cdot \mathcal{M}\left(\underbrace{\frac{\beta - \beta^\natural}{\|\beta - \beta^\natural\|_2}}_{\in L} + \beta^\natural, \beta^\natural\right) > 0,$$

which implies that β^\natural is the only solution to (LS_K) . □

5.3 Proof of Corollary 2.12

For $t > 0$ and $L \subset \mathbb{R}^p$, we have that $\frac{1}{t}L \subset \text{cone}(L)$. Due to the homogeneity of the semi-norms $\|\cdot\|_g$ and $\|\cdot\|_e$, we can rewrite the definition of $q_{t,n}^{(g,e)}(L)$ as

$$q_{t,n}^{(g,e)}(L) = \inf \left\{ \frac{\gamma_1(S, \|\cdot\|_e)}{\sqrt{n}} + \gamma_2(S, \|\cdot\|_g + \|\cdot\|_e) \mid S \subset \mathbb{R}^p, \text{conv}(S) \supset \frac{1}{t}L \cap \mathbb{S}^{p-1} \right\},$$

which coincides with the definition of $q_{0,n}^{(g,e)}(L)$ except that the infimum is taken over an inclusion-wise larger domain of sets. Therefore, it holds that $q_{t,n}^{(g,e)}(L) \leq q_{0,n}^{(g,e)}(L)$, and analogously, we have $m_t^{(g,e)}(L) \leq m_0^{(g,e)}(L)$. It follows that the replacement of the local complexities by the conic complexities (and $\rho_t(\beta^\natural)$ by $\rho_0(\beta^\natural)$) leads to stronger conditions in (2.4) and (2.5), which obviously cannot harm the validity of the theorem. □

5.4 Proof of Lemma 2.14

The claim of (i) follows from the fact that the infima in the global complexity parameters are taken over inclusion-wise smaller domains of sets. The claim of (ii) follows from the fact that the affine term v does not affect the pseudo-metrics induced by the semi-norms $\|\cdot\|_g$ and $\|\cdot\|_e$.

For the claim of (iii), we first observe that

$$\gamma_2(L, \|\cdot\|_g + \|\cdot\|_e) \lesssim \gamma_2(L, \|\cdot\|_g) + \gamma_2(L, \|\cdot\|_e), \tag{5.8}$$

which is stated as an exercise by Talagrand [51, Exc. 2.2.24]; its proof is based on the same strategy as the proof of Lemma 5.13: based on two admissible partition sequences that approximate $\gamma_2(L, \|\cdot\|_g)$ and $\gamma_2(L, \|\cdot\|_e)$ up to a factor of 2, a third partition sequence is defined as in (5.5). Making use of (5.8), we obtain

$$\begin{aligned} q_n^{(g,e)}(L) &\lesssim \inf_S \left\{ \gamma_1(S, \|\cdot\|_e) + \gamma_2(S, \|\cdot\|_g) + \gamma_2(S, \|\cdot\|_e) \right\} \\ &\lesssim \inf_S \left\{ \gamma_1(S, \|\cdot\|_e) + \gamma_2(S, \|\cdot\|_g) \right\} \\ &= m^{(g,e)}(L), \end{aligned}$$

where the second inequality is due to the fact that $\gamma_2(L, d) \leq \gamma_1(L, d)$ holds true for all sets L and pseudo-metrics d .

The claims of (iv) and (v) follow directly from the respective definitions. □

5.5 Proof of Corollary 2.15

We apply Theorem 2.10 for the precision level

$$t := \tilde{C} \cdot \max \left\{ 1, (\tau \cdot Q_{2\tau}(K^\Delta, x))^{-2} \right\} \cdot \left[\rho_0(\beta^\natural) + \max\{1, u^2 \cdot \sigma(\beta^\natural)\} \cdot \frac{\sqrt{m^{(g,e)}(K)}}{n^{1/4}} \right]_+,$$

where the universal constant $\tilde{C} \geq 1$ will be specified later.

For this specific choice of t , Lemma 2.14(i)-(iii) imply that

$$\begin{aligned} q_{t,n}^{(g,e)}(K - \beta^\natural) &\leq \frac{1}{t} \cdot q_n^{(g,e)}(K) \\ &\leq \tilde{C} \cdot \min \left\{ 1, (\tau \cdot Q_{2\tau}(K^\Delta, x))^2 \right\} \cdot \frac{n^{1/4}}{\sqrt{m^{(g,e)}(K)}} \cdot q_n^{(g,e)}(K) \\ &\lesssim \sqrt{\tau \cdot Q_{2\tau}(K^\Delta, x)} \cdot \frac{n^{1/4}}{\sqrt{q_n^{(g,e)}(K)}} \cdot q_n^{(g,e)}(K) \\ &= \sqrt{\tau \cdot Q_{2\tau}(K^\Delta, x)} \cdot n^{1/4} \cdot \sqrt{q_n^{(g,e)}(K)}, \end{aligned} \tag{5.9}$$

where we have used that

$$\min \left\{ 1, (\tau \cdot Q_{2\tau}(K^\Delta, x))^2 \right\} \leq \sqrt{\tau \cdot Q_{2\tau}(K^\Delta, x)}.$$

This implies

$$\begin{aligned}
 q_{t,n}^{(g,e)}(K - \beta^\natural) + \tau \cdot u &\lesssim \sqrt{\tau \cdot Q_{2\tau}(K^\Delta, x)} \cdot n^{1/4} \cdot \sqrt{q_n^{(g,e)}(K)} + \tau \cdot u \\
 &\stackrel{(2.7)}{\lesssim} \sqrt{\tau \cdot Q_{2\tau}(K^\Delta, x)} \cdot n^{1/4} \cdot \sqrt{\tau \cdot Q_{2\tau}(K^\Delta, x)} \cdot n^{1/4} + \tau \cdot Q_{2\tau}(K^\Delta, x) \cdot \sqrt{n} \\
 &\leq 2 \cdot \sqrt{n} \cdot \tau \cdot Q_{2\tau}(K^\Delta, x),
 \end{aligned}$$

and if the hidden constant in (2.7) is appropriately chosen, then the condition (2.4) is indeed fulfilled.

Analogously to the estimates in (5.9), we obtain from Lemma 2.14 that (note that $0 \in K - \beta^\natural$)

$$m_t^{(g,e)}(K - \beta^\natural) \leq \sqrt{\tau \cdot Q_{2\tau}(K^\Delta, x)} \cdot n^{1/4} \cdot \sqrt{m^{(g,e)}(K)}.$$

This implies

$$\begin{aligned}
 &\frac{1}{(\tau \cdot Q_{2\tau}(K^\Delta, x))^2} \cdot \left[\rho_t(\beta^\natural) + u^2 \cdot \sigma(\beta^\natural) \cdot \frac{m_t^{(g,e)}(K - \beta^\natural)}{\sqrt{n}} \right]_+ \\
 &\leq \max \left\{ 1, (\tau \cdot Q_{2\tau}(K^\Delta, x))^{-2} \right\} \cdot \left[\rho_0(\beta^\natural) + u^2 \cdot \sigma(\beta^\natural) \cdot \frac{\sqrt{m^{(g,e)}(K)}}{n^{1/4}} \right]_+ \\
 &\leq \tilde{C}^{-1} \cdot t,
 \end{aligned}$$

where we have used the definition of t and the fact that

$$\frac{\sqrt{\tau \cdot Q_{2\tau}(K^\Delta, x)}}{(\tau \cdot Q_{2\tau}(K^\Delta, x))^2} = (\tau \cdot Q_{2\tau}(K^\Delta, x))^{-\frac{3}{2}} \leq \max \left\{ 1, (\tau \cdot Q_{2\tau}(K^\Delta, x))^{-2} \right\}.$$

Finally, if the constant \tilde{C} is chosen sufficiently large, then the condition (2.5) is fulfilled and Theorem 2.10 yields the claim. □

5.6 Proof of Proposition 1.4

Proposition 1.4 is a direct consequence of Corollary 2.15 with the following specifications. Since $\beta^\natural := \beta^*$, we have $\rho_0(\beta^*) \leq 0$ according to Appendix B.1. The global complexity terms can be bounded according to Proposition 3.13. Finally, we make use of the bound (B.3) from Appendix B.2(2). Let α and δ be as in (B.2) and set $\tau := \alpha/4$. Since x is isotropic, we have $\delta = 1$. To see that $\alpha \gtrsim \kappa^{-3}$, observe that

$$1 = \mathbb{E}[\langle x, v \rangle^2]^2 \leq \mathbb{E}[|\langle x, v \rangle|] \cdot \mathbb{E}[|\langle x, v \rangle|^3] \leq \mathbb{E}[|\langle x, v \rangle|] \cdot 3^3 \cdot \kappa^3 \quad \text{for all } v \in \mathbb{S}^{p-1},$$

where we have used the isotropy of x , the Cauchy-Schwarz inequality, and Proposition A.1(ii). Now, (B.3) implies that $\tau \cdot Q_{2\tau}(K^\Delta, x) \gtrsim \kappa^{-9}$. Since $\tau = \alpha/4$, (B.3) is equivalent to $Q_{2\tau}(K^\Delta, x) \geq \alpha^2/(4\delta)$, which implies that $Q_{2\tau}(K^\Delta, x) \gtrsim \kappa^{-6}$.

Finally, note that due to the isotropy of x , we have that $\kappa \gtrsim 1$. This concludes the proof. \square

Appendix A: Basic facts on sub-Gaussian and sub-exponential random variables

The following proposition provides two characterizations of sub-exponential and sub-Gaussian random variables. The first one concerns the exponential decay behavior of their tails, while the second one addresses the growth of their absolute moments. Note that the dependence of the constants on α can be dropped here, since we consider only two values of α .

Proposition A.1 [58, Prop. 2.5.2, Prop. 2.7.1] *For $\alpha \in \{1, 2\}$, let $Z \in L_{\psi_\alpha}$. Then the following holds true:*

(i) *Z satisfies the concentration inequality*

$$\mathbb{P}(|Z| \geq t) \leq 2 \exp\left(-\left(\frac{t}{C_\alpha \cdot \|Z\|_{\psi_\alpha}}\right)^\alpha\right) \text{ for all } t \geq 0,$$

where $C_\alpha > 0$ is a constant depending on α .

(ii) *The moments of Z satisfy*

$$\|Z\|_{L^q} \leq C_\alpha \cdot \|Z\|_{\psi_\alpha} \cdot q^{1/\alpha} \text{ for all } q \geq 1,$$

where $C_\alpha > 0$ is a constant depending on α .

The following two results are well-known inequalities for sub-Gaussian and sub-exponential random vectors, respectively. A comparison of both shows that the (weighted) sum of independent sub-exponential variables exhibits a mixed-tail behavior, as if it “were a mixture of sub-gaussian and sub-exponential distributions”; quote from [58, p. 35].

Theorem A.2 (Hoeffdings’s inequality; [58, Thm. 2.6.3]) *Let $x = (x_1, \dots, x_p) \in \mathbb{R}^p$ be a centered random vector with independent, sub-Gaussian coordinates and $R := \max_{1 \leq j \leq p} \|x_j\|_{\psi_2}$. Then, for every $v \in \mathbb{R}^p$ and $t \geq 0$, we have that*

$$\mathbb{P}(|\langle x, v \rangle| \geq t) \leq 2 \exp\left(-C_H \cdot \frac{t^2}{R^2 \|v\|_2^2}\right),$$

where $C_H > 0$ is a universal constant.

Theorem A.3 (Bernstein’s inequality; [58, Thm. 2.8.2]) *Let $x = (x_1, \dots, x_p) \in \mathbb{R}^p$ be a centered random vector with independent, sub-exponential coordinates and $R := \max_{1 \leq j \leq p} \|x_j\|_{\psi_1}$. Then, for every $v \in \mathbb{R}^p$ and $t \geq 0$, we have that*

$$\mathbb{P}(|\langle x, v \rangle| \geq t) \leq 2 \exp\left(-C_B \cdot \min\left\{\frac{t^2}{R^2 \|v\|_2^2}, \frac{t}{R \|v\|_\infty}\right\}\right),$$

where $C_B > 0$ is a universal constant.

Appendix B: Further details on Sect. 2

B.1: Remarks on the mismatch covariance (Definition 2.8)

In this part, we adopt the notation of Definition 2.8, which has introduced the mismatch parameters. Since $K^t \subset K^0 \subset \mathbb{S}^{p-1}$, we observe that the mismatch covariance satisfies

$$\rho_t(\beta^\natural) \leq \rho_0(\beta^\natural) \leq \rho(\beta^\natural). \tag{B.1}$$

These bounds imply that all estimation guarantees presented in Sect. 2 remain true when replacing the local mismatch covariance by its global variant. However, there exist relevant scenarios where the second inequality in (B.1) becomes strict, so that considering $\rho(\beta^\natural)$ leads to suboptimal results. To this end, it is useful to first relate the mismatch covariance to the expected risk minimization problem (1.1): Let $\mathcal{L}(\beta) := \mathbb{E}[(y - \langle x, \beta \rangle)^2]$ be the objective function in (1.1) and assume that $K \subset \mathbb{R}^p$ is compact and convex. Then, we have that $\nabla \mathcal{L}(\beta^\natural) = 2\mathbb{E}[\langle x, \beta^\natural \rangle - y]x$ and therefore

$$\rho_t(\beta^\natural) = \sup_{v \in K^t} \left\langle -\frac{1}{2} \nabla \mathcal{L}(\beta^\natural), v \right\rangle \quad \text{and} \quad \rho(\beta^\natural) = \left\| \frac{1}{2} \nabla \mathcal{L}(\beta^\natural) \right\|_2.$$

Now, let $\beta^\natural \in K$ be an expected risk minimizer on K , i.e., a solution to (1.1). A well-known optimality condition in convex analysis then implies that $\rho_t(\beta^\natural) \leq 0$. On the other hand, if K does not contain a global expected risk minimizer, i.e., a solution to $\min_{\beta \in \mathbb{R}^p} \mathcal{L}(\beta)$, we have that $\nabla \mathcal{L}(\beta^\natural) \neq 0$ and therefore $\rho(\beta^\natural) > 0$; and vice versa, $\rho(\beta^\natural) = 0$ implies that K contains a global expected risk minimizer. We refer to Fig. 2 for an illustration of this argument when x is isotropic.¹⁷

In view of our main result, Theorem 2.10, the local mismatch covariance measures the asymptotic impact of the model mismatch. When positive, $\rho_t(\beta^\natural)$ can be seen as an asymptotic bias term, while a negative value can have favorable effects on the estimation performance of (LS_K) ; see Appendix B.2(3).

B.2: Remarks on Theorem 2.10

This part compiles several additional remarks on our main result, Theorem 2.10.

1. *Possible extensions.* Theorem 2.10 is amenable to various extensions and generalizations. For instance, replacing the ℓ^2 -norm by an arbitrary semi-norm $\|\cdot\|$ in Fact 2.4 would lead to an error bound in terms of $\|\cdot\|$; note that such a step would also require an adaptation of the spherical intersections in the q- and m-complexities and the small-ball condition (2.3). This extension becomes particularly useful when the covariance matrix of the input vector x is poorly conditioned

¹⁷ In the isotropic case, there also exists a nice functional-analytic interpretation: The mapping $\beta \mapsto \langle x, \beta \rangle$ is an isometric embedding of the Hilbert space \mathbb{R}^p into $L_2(\Omega, \mathbb{P})$. Then the components $x_1, \dots, x_p \in L_2$ of the random vector x constitute an orthonormal basis for the subspace $G := \{\langle x, \beta \rangle \mid \beta \in \mathbb{R}^p\}$. This implies $\rho(\beta^\natural) = \|\langle y - \langle x, \beta^\natural \rangle, x_j \rangle_{L_2}\|_{j=1}^p \|\cdot\|_2 = \|P_G(y - \langle x, \beta^\natural \rangle)\|_{L_2}$, where P_G is the orthogonal projection onto G ; in other words, $\rho(\beta^\natural)$ corresponds to the “linear component” of the expected risk of β^\natural .

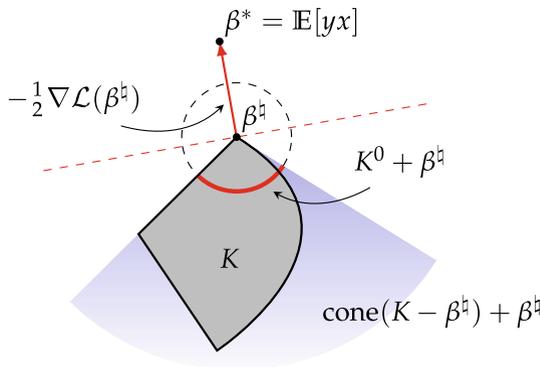


Fig. 2 An illustration of the argument in Appendix B.1 when x is isotropic: in this case, we have $\frac{1}{2}\nabla\mathcal{L}(\beta) = \beta - \mathbb{E}[yx]$, implying that $\beta^* := \mathbb{E}[yx]$ is the (unique) global expected risk minimizer. The above figure shows a situation where $\beta^* \notin K$ and β^\dagger is the expected risk minimizer on K . This implies that the negative gradient at β^\dagger points out of K in the direction of β^* , or more geometrically, the dashed supporting hyperplane separates K and β^* . Hence, we have that $\langle -\nabla\mathcal{L}(\beta), v \rangle \leq 0$ for all $v \in K - \beta^\dagger$, and in particular, $\rho_t(\beta^\dagger) \leq 0$. On the other hand, it holds that $\rho(\beta^\dagger) = \|\frac{1}{2}\nabla\mathcal{L}(\beta^\dagger)\|_2 = \|\beta^* - \beta^\dagger\|_2 > 0$

or even degenerate.¹⁸ In this case, an appropriate linear transform of the ℓ^2 -error can account for the underlying covariance structure; see [15, Chap. 3 and Sec. 4.3] for a detailed discussion of this issue.

Apart from that, it is possible to incorporate different loss functions or adversarial noise into Theorem 2.10; cf. [14] and [15, Chap. 3]. Working out the details goes beyond the scope of this paper, but is expected to be relatively straightforward. Furthermore, one might show similar estimation guarantees for the “basis-pursuit” version or the unconstrained version of the generalized Lasso; cf. [15, Chap. 3] and [28, 29]. Finally, it is worth mentioning that the sub-exponentiality of y in Assumption 2.9 could be replaced by a less restrictive tail condition. This modification would concern the analysis of the multiplier process in Sect. 5.2.2; for example, a finite moment assumption for y would be sufficient when using Theorem 5.9 instead of Theorem 5.12.

- 2. *Small-ball condition.* The small-ball condition (2.3) in Assumption 2.9 is based on Mendelson’s more general condition (see [31, Asm. 3.1]), which reads

$$\exists u > 0 : Q_{\mathcal{F}-\mathcal{F}}(u) := \inf_{f \in \mathcal{F}-\mathcal{F}} \mathbb{P}(|f(x)| \geq u \|f\|_{L^2}) > 0,$$

where \mathcal{F} is a class of (not necessarily linear) hypothesis functions. We emphasize that the small-ball condition (2.3) is stated relative to the hypothesis set K , and it particularly suffices for x to be non-degenerate relative to the subspace $\text{span}(K - K)$. This reflects the fact that the input vectors x_i are only of interest to us insofar as they enable us to discern differences between the hypotheses in K . Furthermore, one can easily replace the small-ball function $Q_{2\tau}(K^\Delta, x)$ in Theorem 2.10 by a more explicit expression. For instance, the Paley-Zygmund

¹⁸ In principle, Assumption 2.9 imposes no explicit conditions on the covariance structure of x , but if it becomes too degenerate, the small-ball condition (2.3) might become unrealizable for $\|\cdot\| = \|\cdot\|_2$.

inequality implies the following lower bound (cf. [57, Subsec. 2.6.5]): Let

$$\alpha := \inf_{v \in K^\Delta} \mathbb{E}[|\langle x, v \rangle|] \quad \text{and} \quad \delta := \sup_{v \in K^\Delta} \mathbb{E}[\langle x, v \rangle^2] \tag{B.2}$$

and set $\tau := \alpha/4$. As long as $\alpha > 0$, we have that

$$\tau \cdot Q_{2\tau}(K^\Delta, x) \geq \frac{\alpha^3}{16\delta}. \tag{B.3}$$

This lower bound is a more convenient expression, since α measures the degeneracy of x relative to K^Δ while δ can be seen as an (an-)isotropy parameter.

3. *Low- and high-noise regime.* The statement of Theorem 2.10 could be further refined by specifying the smallest value of t such that the conditions (2.4) and (2.5) still hold true, while all other model parameters remain fixed. Such an optimization strategy is elaborated in the general learning framework of Mendelson [31]. Although the latter has certainly a wider scope than ours, there are important conceptual overlaps. Indeed, (2.4) is closely related to what Mendelson refers to as the “low-noise” regime: this condition is intrinsic to the hypothesis set K and does not depend on the model mismatch (the “noise”) $y - \langle x, \beta^\natural \rangle$; in particular, it specifies how many samples are required for (LS_K) to recover a linear hypothesis function exactly. In contrast, the condition (2.5) is associated with the “high-noise” regime, as it strongly depends on the model mismatch in terms of $\rho_t(\beta^\natural)$ and $\sigma(\beta^\natural)$. A remarkable conclusion is possible when $\rho_0(\beta^\natural) < 0$ (cf. Appendix B.1): in this case, we may simply set $t = 0$, while (2.5) can be even satisfied if $\sigma(\beta^\natural) > 0$. In other words, exact recovery of β^\natural is feasible in certain scenarios, despite the presence of noise or model misspecifications.
4. *Alternative m -complexity.* An important detail in the definition of the m -complexity $m_t^{(g,e)}$ is that $\text{conv}(S)$ needs to contain the origin as well. Without this minor modification, an extra term would have to be added to $m_t^{(g,e)}$ in the error bound of (2.5), capturing the radii of $(K - \beta^\natural) \cap t\mathbb{S}^{p-1}$ with respect to $\|\cdot\|_g$ and $\|\cdot\|_e$ (see the proof of Lemma 5.13). The appearance of such an additive term is in fact quite common in the literature, e.g., see [11, 32].

Appendix C: Proof of Proposition 3.12

For the claims of (i) and (ii), we first apply Lemma 2.14(iii)–(v) to observe that

$$q_{t,n}^{(g,e)}(K - \beta^\natural) \lesssim m_t^{(g,e)}(K - \beta^\natural) \leq m_0^{(g,e)}(K - \beta^\natural) = m^{(g,e)}(\underbrace{(\text{cone}(K - \beta^\natural) \cap \mathbb{S}^{p-1}) \cup \{0\}}_{=: \mathcal{D}_1(\beta^\natural)}), \tag{C.1}$$

where the second inequality is due to $\frac{1}{t}(K - \beta^\natural) \subset \text{cone}(K - \beta^\natural)$ for all $t > 0$. Moreover, according to (3.3) and (3.5), it holds that

$$m^{(2,0)}((\mathcal{D}_1(\beta^\natural) \cap \mathbb{S}^{p-1}) \cup \{0\}) \asymp w((\mathcal{D}_1(\beta^\natural) \cap \mathbb{S}^{p-1}) \cup \{0\})$$

and

$$m^{(2,\infty)}((\mathcal{D}_1(\beta^\natural) \cap \mathbb{S}^{p-1}) \cup \{0\}) \lesssim \sqrt{\log(p)} \cdot w((\mathcal{D}_1(\beta^\natural) \cap \mathbb{S}^{p-1}) \cup \{0\}).$$

Since $\mathcal{D}_1(\beta^\natural)$ corresponds to the descent cone of $\|\cdot\|_1$ at a k -sparse vector β^\natural , the claims of (i) and (ii) now follow from a standard bound for the conic Gaussian width (e.g., see [57, Eq. 2.11]):

$$w((\mathcal{D}_1(\beta^\natural) \cap \mathbb{S}^{p-1}) \cup \{0\}) \lesssim \sqrt{k \log\left(\frac{p}{k}\right)}.$$

The proof of (iii) is a little more involved. In a first step, we intend to show that

$$\mathcal{D}_1(\beta^\natural) \cap \mathbb{S}^{p-1} \subset \text{conv}(S) \tag{C.2}$$

with $S := \{v \in \mathbb{R}^p \mid \|v\|_0 \leq k, \|v\|_2 \leq 3\}$. In order to verify this inclusion, we adapt the proof of Pilanci and Wainwright [41, Lem. 13]: Writing $A := \mathcal{D}_1(\beta^\natural) \cap B_2^p$, we define the support functions $h_A, h_S : \mathbb{R}^p \rightarrow \mathbb{R}$ by

$$h_A(z) := \sup_{v \in A} \langle z, v \rangle \quad \text{and} \quad h_S(z) := \sup_{v \in S} \langle z, v \rangle = \sup_{v \in \text{conv}(S)} \langle z, v \rangle.$$

If there would exist a point $v \in A \setminus \text{conv}(S)$, then a variant of the hyperplane separation theorem implies that there exists some $z \in \mathbb{R}^p$ with $\langle z, v \rangle > h_S(z)$ and therefore $h_A(z) > h_S(z)$. Consequently, it suffices to show that $h_A(z) \leq h_S(z)$ holds true for all $z \in \mathbb{R}^p$.

To this end, let $z \in \mathbb{R}^p \setminus \{0\}$ be fixed and let $I_1 \subset \{1, \dots, p\}$ contain k indices whose corresponding entries in z have the k largest absolute values. Then, it is not hard to see that $h_S(z) = \langle z, v \rangle$ for

$$v = \frac{3}{\|P_{I_1}(z)\|_2} \cdot P_{I_1}(z) \in S,$$

where $P_{I_1}(z) \in \mathbb{R}^{p \times p}$ denotes the orthogonal projection onto the coordinate space associated with I_1 . Consequently, we obtain

$$h_S(z) = 3 \cdot \sqrt{\sum_{j \in I_1} z_j^2},$$

and for every $j \notin I_1$, we have that

$$|z_j| \leq \frac{1}{k} \cdot \sum_{j' \in I_1} |z_{j'}| \leq \frac{1}{\sqrt{k}} \cdot \sqrt{\sum_{j' \in I_1} z_{j'}^2}.$$

Now let $v \in A$. We denote by $I_2 \subset \{1, \dots, p\}$ the set of indices corresponding to the non-zero entries of β^\natural , which satisfies $|I_2| \leq k$. Since v belongs to descent cone of

$\|\cdot\|_1$ at β^{\natural} and to the Euclidean unit ball, it follows that

$$\sum_{j \notin I_2} |v_j| \leq \sum_{j \in I_2} |v_j| \leq \sqrt{k} \sqrt{\sum_{j \in I_2} v_j^2} \leq \sqrt{k}.$$

Combining these observations, we obtain

$$\begin{aligned} \langle z, v \rangle &= \sum_{j \in I_2} z_j v_j + \sum_{\substack{j \notin I_2 \\ j \in I_1}} z_j v_j + \sum_{\substack{j \notin I_2 \\ j \notin I_1}} z_j v_j \\ &\leq 2 \cdot \|v\|_2 \cdot \sqrt{\sum_{j \in I_1} z_j^2} + \max_{j \notin I_1} |z_j| \cdot \sqrt{k} \leq 3 \cdot \sqrt{\sum_{j \in I_1} z_j^2} = h_S(z), \end{aligned}$$

which concludes the proof of (C.2).

Next, we observe that $S \subset 3\sqrt{k}B_1^p$, which is a consequence of the Cauchy-Schwarz inequality. Now, consider the finite set $F := 3\sqrt{k} \cdot \{\pm u_1, \dots, \pm u_p\}$, where $u_1, \dots, u_p \in \mathbb{R}^p$ denote the Euclidean unit vectors. Then, we have that

$$\mathcal{D}_1(\beta^{\natural}) \cap \mathbb{S}^{p-1} \subset \text{conv}(S) \subset \text{conv}(3\sqrt{k}B_1^p) = 3\sqrt{k}B_1^p = \text{conv}(F). \tag{C.3}$$

Finally, we make use of (3.8) to obtain the following bound:

$$\gamma_1(F, \|\cdot\|_2) \lesssim \Delta_2(F) \cdot \log(|F|) \lesssim \sqrt{k} \cdot \log(2p).$$

The claim of (iii) follows directly from a combination of this bound with (C.1), (C.3), and the definition of the m-complexity. \square

Appendix D: Proofs for Sect. 5.1

Proof of Lemma 5.1 Using the generic Bernstein concentration of x , we observe that

$$\begin{aligned} \|v^*\|_{L^q}^q &= \mathbb{E}[|\langle v, x \rangle|^q] = \int_0^\infty \mathbb{P}(|\langle v, x \rangle|^q \geq s) ds \\ &= \int_0^\infty \mathbb{P}(|\langle v, x \rangle| \geq r) \cdot qr^{q-1} dr \\ &\leq \int_0^\infty 2 \exp\left(-\min\left\{\frac{r^2}{\|v\|_g^2}, \frac{r}{\|v\|_e}\right\}\right) \cdot qr^{q-1} dr \\ &= \int_0^\infty 2 \max\left\{\exp\left(-\frac{r^2}{\|v\|_g^2}\right), \exp\left(-\frac{r}{\|v\|_e}\right)\right\} \cdot qr^{q-1} dr \\ &\leq \int_0^\infty 2 \exp\left(-\frac{r^2}{\|v\|_g^2}\right) \cdot qr^{q-1} dr + \int_0^\infty 2 \exp\left(-\frac{r}{\|v\|_e}\right) \cdot qr^{q-1} dr \\ &= 2q \cdot \|v\|_g^q \cdot \int_0^\infty \exp(-r^2) \cdot r^{q-1} dr + 2q \cdot \|v\|_e^q \cdot \int_0^\infty \exp(-r) \cdot r^{q-1} dr \\ &\stackrel{\text{Definition of } \Gamma\text{-function}}{=} q \cdot \|v\|_g^q \cdot \Gamma(q/2) + 2q \cdot \|v\|_e^q \cdot \Gamma(q) \end{aligned}$$

$$\begin{aligned} &\text{Stirling's approximation} \\ &\leq 2q \cdot \|v\|_g^q \cdot (q/2)^{q/2} + 2q \cdot \|v\|_e^q \cdot q^q. \end{aligned}$$

Since $(a + b)^{1/q} \leq a^{1/q} + b^{1/q}$ for all $a, b \geq 0$ and $q \geq 1$, this implies

$$\|v^*\|_{L^q} \leq \sqrt[q]{2q} \cdot \left(\|v\|_g \cdot \frac{\sqrt{q}}{\sqrt{2}} + \|v\|_e \cdot q \right). \tag{D.1}$$

Observing that $q \mapsto \sqrt[q]{2q}$ is bounded on the interval $[1, \infty)$, the claim follows. \square

Proof of Lemma 5.2 Our basic proof strategy is adapted from the proof of Bernstein’s inequality by Vershynin [58, Thm. 2.8.1]. Let $\tilde{x}_i := \varepsilon_i x_i$ and $S := \sum_{i=1}^n \tilde{x}_i$. Then, the generic Chernoff bound (for arbitrary $v \in \mathbb{R}^p$, $\lambda \in \mathbb{R}$, and $t \geq 0$) reads as follows:

$$\mathbb{P}(\langle v, S \rangle \geq t) = \mathbb{P}\left(\sum_{i=1}^n \langle v, \tilde{x}_i \rangle \geq t\right) \leq \exp(-\lambda t) \cdot \prod_{i=1}^n \mathbb{E}[\exp(\lambda \langle v, \tilde{x}_i \rangle)].$$

To proceed, we need an upper bound for the moment generating function $\mathbb{E}[\exp(\lambda \langle v, \tilde{x}_1 \rangle)]$. Since \tilde{x}_1 is centered, we have that

$$\mathbb{E}[\exp(\lambda \langle v, \tilde{x}_1 \rangle)] = \mathbb{E}\left[1 + \lambda \langle v, \tilde{x}_1 \rangle + \sum_{q=2}^{\infty} \frac{(\lambda \langle v, \tilde{x}_1 \rangle)^q}{q!}\right] = 1 + \sum_{q=2}^{\infty} \frac{\lambda^q \cdot \mathbb{E}[\langle v, \tilde{x}_1 \rangle^q]}{q!}.$$

Due to symmetry, (2.1) also holds true for the symmetrized random vector \tilde{x}_1 and from the proof of Lemma 5.1, we therefore obtain

$$\begin{aligned} \mathbb{E}[\exp(\lambda \langle v, \tilde{x}_1 \rangle)] &\leq 1 + \sum_{q=2}^{\infty} \frac{\lambda^q \cdot \mathbb{E}[|\langle v, \tilde{x}_1 \rangle|^q]}{q!} \\ &\leq 1 + \sum_{q=2}^{\infty} \frac{\lambda^q \cdot (2q \cdot \|v\|_g^q \cdot (q/2)^{q/2} + 2q \cdot \|v\|_e^q \cdot q^q)}{q!} \\ &= 1 + \underbrace{\sum_{q=2}^{\infty} \frac{\lambda^q \cdot 2q \cdot \|v\|_e^q \cdot q^q}{q!}}_{=:A} + \underbrace{\sum_{q=2}^{\infty} \frac{\lambda^q \cdot 2q \cdot \|v\|_g^q \cdot (q/2)^{q/2}}{q!}}_{=:B}. \end{aligned}$$

Upper bound for A: According to Stirling’s approximation, we have $q! \geq (q/e)^q$, which implies

$$A \leq 1 + \sum_{q=2}^{\infty} \frac{2q \cdot (\lambda \|v\|_e)^q \cdot q^q}{(q/e)^q} = 1 + \sum_{q=2}^{\infty} 2q \cdot (e\lambda \|v\|_e)^q.$$

If $\lambda < \frac{1}{2e\|v\|_e}$, the above series is convergent and we have

$$A \leq 1 + 2 \cdot \frac{(e\lambda\|v\|_e)^2 \cdot (2 - e\lambda\|v\|_e)}{(1 - e\lambda\|v\|_e)^2} \leq 1 + 16(e\lambda\|v\|_e)^2 \leq \exp(16(\lambda\|v\|_e)^2).$$

Upper bound for B: Since we do not want to introduce further restrictions on λ , let us distinguish two cases: If $\lambda < 1/(2e\|v\|_g)$, we use that $(q/2)^{q/2} \leq q^q$ (due to $q \geq 2$) and apply the same strategy as for the bound for A , which yields $B \leq \exp(16(\lambda\|v\|_g)^2)$. Now, let $\lambda \geq 1/(2e\|v\|_g)$. In this case, we use the basic inequality $2q(q/2)^{q/2} \leq 3q!/\lfloor q/2 \rfloor!$, which holds true for all $q \geq 2$. This implies

$$B \leq 3 \sum_{q=2}^{\infty} \frac{(\lambda\|v\|_g)^q}{\lfloor \frac{q}{2} \rfloor!}.$$

Splitting this series into even and odd indices then yields

$$B \leq 3 \left(\sum_{j=1}^{\infty} \frac{(\lambda\|v\|_g)^{2j}}{\lfloor \frac{2j}{2} \rfloor!} + \sum_{j=1}^{\infty} \frac{(\lambda\|v\|_g)^{2j+1}}{\lfloor \frac{2j+1}{2} \rfloor!} \right) = 3 \left(\underbrace{\sum_{j=1}^{\infty} \frac{(\lambda\|v\|_g)^{2j}}{j!}}_{=\exp((\lambda\|v\|_g)^2)-1} + \underbrace{\sum_{j=1}^{\infty} \frac{(\lambda\|v\|_g)^{2j+1}}{j!}}_{\leq \lambda\|v\|_g \exp((\lambda\|v\|_g)^2)} \right).$$

By $(1 + \lambda\|v\|_g) \leq \exp((\lambda\|v\|_g)^2)$, it follows that

$$B \leq 3 \exp(2(\lambda\|v\|_g)^2) - 3.$$

Since $\lambda\|v\|_g \geq 1/(2e)$, there exists a universal constant $C' > 0$ with $3 \leq \exp(C' \cdot (\lambda\|v\|_g)^2)$, so that we obtain

$$B \leq \exp((2 + C') \cdot (\lambda\|v\|_g)^2) - 3.$$

Next, we combine the above bounds for A and B : the basic inequality $\exp(a) + \exp(b) - 3 \leq \exp(a + b)$ for all $a, b \geq 0$ implies that

$$\mathbb{E}[\exp(\lambda \langle v, \tilde{x}_1 \rangle)] \leq \exp\left(\max\{2 + C', 16\} \cdot (\lambda\|v\|_g)^2 + 16(\lambda\|v\|_e)^2\right) \leq \exp(C\lambda^2(\|v\|_g + \|v\|_e)^2)$$

for all $\lambda < 1/(2e\|v\|_e)$ and a universal constant $C > 0$. Plugging this into (5.6) and assuming that $\lambda < 1/(2e\|v\|_e)$, we have that

$$\mathbb{P}(\langle v, S \rangle \geq t) \leq \exp(-\lambda t) \cdot \prod_{i=1}^n \mathbb{E}[\exp(\lambda \langle v, \tilde{x}_i \rangle)] = \exp(-\lambda t + nC\lambda^2(\|v\|_g + \|v\|_e)^2).$$

Setting

$$\lambda := \min \left\{ \frac{t}{2nC(\|v\|_g + \|v\|_e)^2}, \frac{1}{4e\|v\|_e} \right\} < \frac{1}{2e\|v\|_e},$$

we finally obtain

$$\mathbb{P}(\langle v, S \rangle \geq t) \leq \exp\left(-\tilde{C} \cdot \min\left\{\frac{t^2}{n(\|v\|_g + \|v\|_e)^2}, \frac{t}{\|v\|_e}\right\}\right)$$

for a universal constant $\tilde{C} > 0$. Now, the claim follows if we replace t by \sqrt{nt} . \square

Funding Open Access funding enabled and organized by Projekt DEAL. M.G. acknowledges support by the German Federal Ministry of Education and Research (BMBF) through the Berlin Center for Machine Learning (BZML), Project AP4. C.K. is supported by the Berlin Mathematics Research Center MATH+.

Availability of data and material Not applicable.

Code availability Not applicable.

Declarations

Conflict of interest Not applicable.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

References

- Adamczak, R.: A note on the Hanson–Wright inequality for random vectors with dependencies. *Electron. Commun. Probab.* **20**(72), 1–13 (2015)
- Ai, A., Lapanowski, A., Plan, Y., Vershynin, R.: One-bit compressed sensing with non-Gaussian measurements. *Linear Algebra Appl.* **441**, 222–239 (2014)
- Amelunxen, D., Lotz, M., McCoy, M.B., Tropp, J.A.: Living on the edge: phase transitions in convex programs with random data. *Inf. Infer.* **3**(3), 224–294 (2014)
- Bakhshizadeh, M., Maleki, A., de la Pena, V.H.: Sharp concentration results for heavy-tailed distributions (2020). Preprint [arXiv:2003.13819](https://arxiv.org/abs/2003.13819)
- Brillinger, D.R.: A generalized linear model with “Gaussian” regressor variables. In: Bickel, P.J., Doksum, K., Hodges, J. (eds.) *A Festschrift For Erich L. Lehmann*, pp. 97–114. Chapman and Hall/CRC, Boca Raton (1982)
- Bühlmann, P., van de Geer, S.: *Statistics for High-Dimensional Data: Methods. Theory and Applications*. Springer Series in Statistics. Springer, Berlin (2011)
- Candès, E.J., Eldar, Y.C., Strohmer, T., Vershynin, V.: Phase retrieval via matrix completion. *SIAM J. Imag. Sci.* **6**(1), 199–225 (2013)
- Candès, E.J., Strohmer, T., Vershynin, V.: Phaselift: exact and stable signal recovery from magnitude measurements via convex programming. *Commun. Pure Appl. Math.* **66**(8), 1241–1274 (2013)
- Celentano, M., Montanari, A.: Fundamental barriers to high-dimensional regression with convex penalties (2019). Preprint [arXiv:1903.10603](https://arxiv.org/abs/1903.10603)
- Chandrasekaran, V., Recht, B., Parrilo, P.A., Willsky, A.S.: The convex geometry of linear inverse problems. *Found. Comput. Math.* **12**(6), 805–849 (2012)
- Dirksen, S.: Tail bounds via generic chaining. *Electron. J. Probab.* **20**(53), 1–29 (2015)

12. Efron, B., Hastie, T., Johnstone, I.M., Tibshirani, R.: Least angle regression. *Ann. Stat.* **32**(2), 407–499 (2004)
13. Foucart, S., Rauhut, H.: *A Mathematical Introduction to Compressive Sensing*. Applied and Numerical Harmonic Analysis. Birkhäuser, Basel (2013)
14. Genzel, M.: High-dimensional estimation of structured signals from non-linear observations with general convex loss functions. *IEEE Trans. Inf. Theory* **63**(3), 1601–1619 (2017)
15. Genzel, M.: The mismatch principle and ℓ^1 -analysis compressed sensing: a unified approach to estimation under large model uncertainties and structural constraints. Phd thesis, available online: <https://doi.org/10.14279/depositonce-8394>, Technische Universität, Berlin (2019)
16. Genzel, M., Jung, P.: Recovering structured data from superimposed non-linear measurements. *IEEE Trans. Inf. Theory* **66**(1), 453–477 (2020)
17. Genzel, M., Kutyniok, G.: The mismatch principle: the generalized Lasso under large model uncertainties. Technical report, not intended for publication (2018). [arXiv:1808.06329](https://arxiv.org/abs/1808.06329)
18. Giannopoulos, A.A., Milman, V.D.: Asymptotic convex geometry short overview. In: Donaldson, S., Eliashberg, Y., Gromov, M. (eds.) *Different Faces of Geometry*, pp. 87–162. Springer, Boston (2004)
19. Goldstein, L., Minsker, S., Wei, X.: Structured signal recovery from non-linear and heavy-tailed measurements. *IEEE Trans. Inf. Theory* **64**(8), 5513–5530 (2018)
20. Gordon, Y.: On Milman’s inequality and random subspaces which escape through a mesh in \mathbb{R}^n . In: Lindenstrauss, J., Milman, V.D. (eds.) *Geometric Aspects of Functional Analysis*. Lecture Notes in Mathematics, vol. 1317, pp. 84–106. Springer, Berlin (1988)
21. Götze, F., Sambale, H., Sinulis, A.: Concentration inequalities for polynomials in α -sub-exponential random variables. *Electron. J. Probab.* **26**, 1–22 (2021)
22. Han, Q., Wellner, J.A.: Convergence rates of least squares regression estimators with heavy-tailed errors. *Ann. Stat.* **47**(4), 2286–2319 (2019)
23. Hastie, T., Tibshirani, R., Wainwright, M.J.: *Statistical Learning with Sparsity: The Lasso and Generalizations*. Chapman and Hall/CRC, Boca Raton (2015)
24. Jeong, H., Li, X., Plan, Y., Yılmaz, Ö.: Sub-gaussian matrices on sets: optimal tail dependence and applications (2020). Preprint [arXiv:2001.10631](https://arxiv.org/abs/2001.10631)
25. Kasiviswanathan, S.P., Rudelson, M.: Restricted eigenvalue from stable rank with applications to sparse linear regression. In: Bubeck, S., Perchet, V., Rigollet, P. (eds) *Proceedings of the 31st Conference On Learning Theory (COLT)*, vol. 75, pp. 1011–1041 (2018)
26. Koltchinskii, V., Mendelson, S.: Bounding the smallest singular value of random matrix without concentration. *Int. Math. Res. Not.* **2015**(23), 12991–13008 (2015)
27. Kuchibhotla, A.K., Chakraborty, A.: Moving beyond sub-Gaussianity in high-dimensional statistics: Applications in covariance estimation and linear regression (2018). Preprint [arXiv:1804.02605](https://arxiv.org/abs/1804.02605)
28. Lecué, G., Mendelson, S.: Regularization and the small-ball method I: sparse recovery. *Ann. Stat.* **46**(2), 611–641 (2016)
29. Lecué, G., Mendelson, S.: Regularization and the small-ball method II: complexity dependent error rates. *J. Mach. Learn. Res.* **18**(146), 1–48 (2017)
30. Lecué, G., Mendelson, S.: Sparse recovery under weak moment assumptions. *J. Eur. Math. Soc.* **19**(3), 881–904 (2017)
31. Mendelson, S.: Learning without concentration. *J. ACM* **62**(3), 21 (2015)
32. Mendelson, S.: Upper bounds on product and multiplier empirical processes. *Stoch. Proc. Appl.* **126**(12), 3652–3680 (2016)
33. Mendelson, S.: Learning without concentration for general loss functions. *Probab. Theory Relat. Fields* **171**(1–2), 459–502 (2018)
34. Mendelson, S., Pajor, A., Tomczak-Jaegermann, N.: Reconstruction and subgaussian operators in asymptotic geometric analysis. *Geom. Funct. Anal.* **17**(4), 1248–1282 (2007)
35. Milman, V.D.: Random subspaces of proportional dimension of finite dimensional normed spaces: approach through the isoperimetric inequality. In: Kalton, N.J., Saab, E. (eds.) *Banach Spaces*. Lecture Notes in Mathematics, vol. 1166, pp. 106–115. Springer, Berlin (1985)
36. Oymak, S.: Learning compact neural networks with regularization (2018). Preprint [arXiv:1802.01223](https://arxiv.org/abs/1802.01223), published in *Proceedings of the 35th International Conference on Machine Learning (ICML)*
37. Oymak, S., Hassibi, B.: Sharp MSE bounds for proximal denoising. *Found. Comput. Math.* **16**(4), 965–1029 (2016)
38. Oymak, S., Soltanolkotabi, M.: Fast and reliable parameter estimation from nonlinear observations. *SIAM J. Optim.* **27**(4), 2276–2300 (2017)

39. Oymak, S., Jalali, A., Fazel, M., Eldar, Y.C., Hassibi, B.: Simultaneously structured models with application to sparse and low-rank matrices. *IEEE Trans. Inf. Theory* **61**(5), 2886–2908 (2015)
40. Oymak, S., Recht, B., Soltanolkotabi, M.: Sharp time-data tradeoffs for linear inverse problems. *IEEE Trans. Inf. Theory* **64**(6), 4129–4158 (2018)
41. Pilanci, M., Wainwright, M.J.: Randomized sketches of convex programs with sharp guarantees. *IEEE Trans. Inf. Theory* **61**(9), 5096–5115 (2015)
42. Plan, Y., Vershynin, R.: The generalized Lasso with non-linear observations. *IEEE Trans. Inf. Theory* **62**(3), 1528–1537 (2016)
43. Plan, Y., Vershynin, R., Yudovina, E.: High-dimensional estimation with geometric constraints. *Inf. Infer.* **6**(1), 1–40 (2016)
44. Rudelson, M., Vershynin, R.: On sparse reconstruction from Fourier and gaussian measurements. *Commun. Pure Appl. Math.* **61**(8), 1025–1045 (2008)
45. Sambale, H.: Some notes on concentration for α -subexponential random variables (2020). Preprint [arXiv:2002.10761](https://arxiv.org/abs/2002.10761)
46. Sattar, Y., Oymak, S.: Quickly finding the best linear model in high dimensions via projected gradient descent. *IEEE Trans. Signal Process.* **68**, 818–829 (2020)
47. Shechtman, Y., Eldar, Y.C., Cohen, O., Chapman, H., Miao, J., Segev, M.: Phase retrieval with application to optical imaging: a contemporary overview. *IEEE Signal Process. Mag.* **32**(3), 87–109 (2015)
48. Sivakumar, V., Banerjee, A., Ravikumar, P.K.: Beyond sub-Gaussian measurements: high-dimensional structured estimation with sub-exponential designs. In: Cortes, C., Lawrence, N.D., Lee, D.D., Sugiyama, M., Garnett, R. (eds.) *Advances in Neural Information Processing Systems*, vol. 28. Curran Associates, pp. 2206–2214 (2015)
49. Stojnic, M.: Various thresholds for ℓ_1 -optimization in compressed sensing (2009). Preprint [arXiv:0907.3666](https://arxiv.org/abs/0907.3666)
50. Talagrand, M.: *The Generic Chaining*. Springer Monographs in Mathematics. Springer, Berlin (2006)
51. Talagrand, M.: *Upper and Lower Bounds for Stochastic Processes. Ergebnisse der Mathematik und ihrer Grenzgebiete*, vol. 3. Springer, Berlin (2014)
52. Thrampoulidis, C., Rawat, A.S.: Lifting high-dimensional non-linear models with Gaussian regressors (2017). Preprint [arXiv:1712.03638](https://arxiv.org/abs/1712.03638), published in Proceedings of the Twenty-Second International Conference on Artificial Intelligence and Statistics (AISTATS)
53. Thrampoulidis, C., Rawat, A.S.: The generalized Lasso for sub-gaussian measurements with dithered quantization. *IEEE Trans. Inf. Theory* **66**(4), 2487–2500 (2020)
54. Thrampoulidis, C., Abbasi, E., Hassibi, B.: LASSO with non-linear measurements is equivalent to one with linear measurements. In: Cortes, C., Lawrence, N.D., Lee, D.D., Sugiyama, M., Garnett, R. (eds.) *Advances in Neural Information Processing Systems*, vol. 28, pp. 3402–3410. Red Hook, Curran Associates (2015)
55. Tibshirani, R.: Regression shrinkage and selection via the Lasso. *J. Roy. Stat. Soc. Ser. B* **58**(1), 267–288 (1996)
56. Tibshirani, R.J., Taylor, J.: The solution path of the generalized lasso. *Ann. Stat.* **39**(3), 1335–1371 (2011)
57. Tropp, J.A.: Convex recovery of a structured signal from independent random linear measurements. In: Pfander, G.E. (ed.) *Sampling Theory, a Renaissance, Applied and Numerical Harmonic Analysis*, pp. 76–101. Cham, Birkhäuser (2015)
58. Vershynin, R.: *High-Dimensional Probability: An Introduction with Applications in Data Science*, volume 47 of Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, Cambridge (2018)
59. Wainwright, M.J.: *High-Dimensional Statistics: A Non-Asymptotic Viewpoint*, volume 48 of Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, Cambridge (2019)
60. Wei, X.: Structured recovery with heavy-tailed measurements: a thresholding procedure and optimal rates (2018). Preprint [arXiv:1804.05959](https://arxiv.org/abs/1804.05959)
61. Yang, Z., Wang, Z., Liu, H., Eldar, Y. C., Zhang, T.: Sparse nonlinear regression: Parameter estimation and asymptotic inference (2015). Preprint [arXiv:1511.04514](https://arxiv.org/abs/1511.04514), published in Proceedings of The 33rd International Conference on Machine Learning (ICML)
62. Yang, Z., Balasubramanian, K., Liu, H.: High-dimensional non-Gaussian single index models via thresholded score function estimation. In: Precup, D., Teh, Y.W. (eds.) *Proceedings of the 34th International Conference on Machine Learning (ICML)*, vol. 70, pp. 3851–3860 (2017a)

63. Yang, Z., Balasubramanian, K., Liu, H.: On Stein's identity and near-optimal estimation in high-dimensional index models (2017b). Preprint [arXiv:1709.08795](https://arxiv.org/abs/1709.08795)
64. Yang, Z., Balasubramanian, K., Wang, Z., Liu, H.: Estimating high-dimensional non-Gaussian multiple index models via Stein's lemma. In: Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R. (eds.) *Advances in Neural Information Processing Systems*, vol. 30, pp. 6097–6106. Red Hook, Curran Associates (2017)
65. Zajkowski, K.: Bounds on tail probabilities for quadratic forms in dependent sub-Gaussian random variables. *Stat. Probab. Lett.* **167**, 108898 (2020)
66. Zou, H., Hastie, T.: Regularization and variable selection via the elastic net. *J. R. Stat. Soc. Ser. B* **67**(2), 301–320 (2005)