# The Effect of the Intrinsic Dimension on the Generalization of Quadratic Classifiers 

Fabian Latorre, Leello Dadi, Paul Rolland and Volkan Cevher<br>Laboratory for Information and Inference Systems, EPFL, Switzerland.<br>\{first\}.\{last\}@epfl.ch


#### Abstract

It has been recently observed that neural networks, unlike kernel methods, enjoy a reduced sample complexity when the distribution is isotropic (i.e., when the covariance matrix is the identity). We find that this sensitivity to the data distribution is not exclusive to neural networks, and the same phenomenon can be observed on the class of quadratic classifiers (i.e., the sign of a quadratic polynomial) with a nuclear-norm constraint. We demonstrate this by deriving an upper bound on the Rademacher Complexity that depends on two key quantities: (i) the intrinsic dimension, which is a measure of isotropy, and (ii) the largest eigenvalue of the second moment (covariance) matrix of the distribution. Our result improves the dependence on the dimension over the best previously known bound and precisely quantifies the relation between the sample complexity and the level of isotropy of the distribution.


## 1 Introduction

We revisit the problem of supervised classification using quadratic features of the data. We do so to highlight the influence of properties of data distrbution on the generalization error. Most of the existing results on this error only use a bound on the support of the distribution. By leveraging results from matrix concentration, we show an improved bound that uses more refined properties of the data distribution, like the second moment matrix.

The use of the second moment matrix in the error bound shows that the intrinsic dimension of the data distribution plays an important role. This is of particular interest because it is widely believed that real-world data distributions have nice properties that allow classifiers, namely neural networks, to avoid the worst-case sample complexities predicted by generalization bounds Jiang* et al. (2020).

Indeed, assumptions like the manifold hypothesis, which state that the data lies on lower dimensional embedded manifold, are often made to explain the practical success of some generative methods. A recent paper by Pope et al. (2021) computes estimates of this true dimensionality of common machine learning datasets and shows that they are much lower than the ambient dimension of the pixel space $[0,1]^{d}$. It is therefore important that properties of the data distribution, going beyond simple bounds on the support, intervene in the study of generalization.

This influence of intrinsic dimension on generalization has been recently observed in the context of differentiating neural networks and from their kernel approximations, like the neural tangent kernel Jacot et al. (2018) or random feature models Yehudai and Shamir (2019). In particular, Ghorbani et al. (2020) observe that neural networks seem to require fewer samples than kernel methods to learn when the data distribution is isotropic.

We show that a similar phenomenon occurs in the simpler setting of quadratic classifiers, which leads to a better understanding of the causes. An improvement in sample complexity on isotropic data
distributions can be proved when comparing nuclear-norm constrained quadratic classifiers and the corresponding kernel method (Frobenius norm constrained classifiers).

The study of quadratic classifiers can serve as an important first step in understanding how neural networks take advantage of the intrinsic dimension to learn with fewer samples Du and Lee (2018); Bai and Lee (2020). The nuclear-norm constraint is a natural one to study in this context. Indeed, when applying weight-decay (or $\ell_{2}$-regularization) on a single-hidden layer neural network with quadratic activations, the regularization is in effect encouraging a low nuclear norm of the coefficient matrix of the quadratic polynomial.

A better understanding of quadratics is also a worthwhile goal in its own right: complex architectures like those in Jayakumar et al. (2020) use quadratics as building blocks, attention layers Vaswani et al. (2017), which have seen great success in language processing tasks, are multiplicative interactions.

For these reasons, we present theoretical and practical developments of nuclear-norm regularization for quadratic classification. We summarize our contributions as follows

Rademacher complexity bounds. We present a new bound on the Rademacher complexity of quadratic classifiers with a nuclear norm constraint c.f. Theorem 1. It improves upon the previously known bound, implied by the results by Kakade et al. (2012), by up to a square-root factor of the dimension, depending on the distribution of the data c.f. Lemma 4

As a consequence of our bound, we draw attention to a clear difference between the complexity of nuclear-norm constrained and Frobenius-norm constrained quadratic classifiers. When the input data distribution is nearly-isotropic, the former enjoys a reduced dependency on the dimension. In contrast, the complexity of Frobenius-norm constrained classifiers has the same dependency on the dimension, independently of how isotropic the input data distribution is Corollary 2 .

This observation motivates the use of data whitening pre-processing steps, which are commonly used in practice: such transformation might bring the second-order moment (covariance) matrix of the distribution close to the identity matrix and thus to nearly-isotropicity.

Computable generalization bounds. The refined Rademacher complexity bound that we obtain depends on the often unknown second-order moment of the distribution, rather than simple bounds on the diameter of the support as in (Kakade et al., 2012). Even though useful in theory, it is desirable in practice to obtain bounds that can be computed from a sample. We overcome this difficulty in Theorem 3, where we provide high-probability computable generalization error bounds for nuclear-norm constrained quadratic classifiers.

Experiments. We illustrate our theoretical results on synthetic data. We show how the isotropy of the input distribution plays a major role in the generalization properties of quadratic classifiers. As the dimension increases and the sample size remains proportional to it, we observe a constant generalization gap for the nuclear-norm constrained classifier. In contrast, for SVMs, the gap grows at a predicted $\sqrt{d}$ rate. In the case of anisotropic distributions, we observe similar performance for both regularization schemes.

### 1.1 Related work

Kakade et al. (2012) provides generalization error bounds for the more general problem of learning a linear classifier over matrices. An upper bound for quadratic classifiers with a nuclear norm constraint can be derived as a consequence of their results c.f. Corollary 1. To the best of our knowledge, it is the only known bound for the hypothesis class we study, and thus the one we compare to. Our analysis improves the dependency on the dimension. See subsection 2.1 for a technical discussion.

Because of the generality of the results in Kakade et al. (2012), it is only natural that the implied bound in some particular case is not the tightest. We precisely give a step towards tight complexity estimates for classification with quadratic polynomials. We look on our results as relevant, given the simplicity and widespread use of linear learning over features.

Wimalawarne et al. (2016) study linear classifiers over higher-order tensor spaces, using constraints on generalized notions of the nuclear norm. The problem we study is thus a particular case. Generalization error bounds via Rademacher complexity are provided, but they apply only under a highly restrictive assumption: the entries of the tensor are independent standard normal random variables. In contrast we only require a boundedness condition.

Srebro et al. (2005); Srebro and Shraibman (2005) develop the theory of nuclear-norm regularization for matrix completion. Bounds on the Rademacher complexity of the class of matrices with bounded nuclear-norm are obtained in Srebro (2005). However, in this setting matrices are understood as mappings from an index pair to a value, and the generalization error measures how well the missing entries of the target matrix can be predicted. In contrast, in our setting the matrix corresponds to the coefficients of a quadratic polynomial, so the bounds are not comparable.

Pontil and Maurer (2013) study nuclear-norm regularization in the context of multi-task learning (Caruana, 1997). Rademacher complexity bounds are obtained for nuclear-norm constrained multitask classifiers. In this setting, each row of the matrix corresponds to a linear classifier for a different task, and each task corresponds to a different distribution over data-label pairs.
Matrices with bounded trace norm also have been studied by Amit et al. (2007) and Yu et al. (2014) in the related multi-class classification and multi-label learning setting, respectively. Yu et al. (2014) remark essential differences between Rademacher complexity bounds of nuclear- and Frobenius-norm constrained linear classifiers, similar to our conclusion in Corollary 2. In all such problems, however, the setting is not comparable to ours: the matrix acts as a linear mapping of the sample, rather than as a quadratic. Thus, the analysis is not analogous and requires a different set of tools in our case.
The analysis in Yu et al. (2014) is closest in spirit to ours, as their bound also depends on the intrinsic dimension of the distribution, and the largest eigenvalue of the second moment (covariance) matrix. However, we go the extra mile and achieve bounds that can be computed from the sample at hand (Theorem 3), as the true second moment (covariance) matrix is usually unknown.

The papers Du and Lee (2018); Bai and Lee (2020) establish a similar result as ours for a different norm for a non-convex parametrization of quadratic polynomials. The reparametrization consits of writing the coefficient matrix as a sum of $m$ rank 1 matrices which facilitates analogies to neural networks. It is straighforward to see that their studied norm $\|\cdot\|_{2,4}$ is, in essence, a way of upper bounding the nuclear-norm by using the reparametrization.

## 2 Rademacher complexity bounds

Notation. Throughout this section $\mathbf{x} \in \mathbb{R}^{d}$ is a random variable with distribution $\mu$, and $\mathbf{X}_{n}=$ $\left(\mathbf{x}_{1}, \ldots \mathbf{x}_{n}\right)$ is a sample of i.i.d. random variables drawn from $\mu$. The second moment matrix of $\mu$ is denoted by $\Sigma:=\mathbb{E}\left[\mathbf{x x}^{T}\right]$. Note that, for centered (mean zero) random variables, this notion coincides with the covariance matrix. For a square symmetric matrix $\mathbf{A}$ we denote with $\|\mathbf{A}\|_{F},\|\mathbf{A}\|_{2}$ and $\|\mathbf{A}\|_{\text {tr }}$ its Frobenius-norm, spectral-norm and nuclear-norm, respectively. The notation $\lesssim$ stands for less than or equal, but hides constants independent of the dimension or number of samples. The notation $x \approx y$ means that there exist constants $c, C>0$ such that $c y \leq x \leq C y$.

We consider a binary classifier obtained from a homogeneous quadratic polynomial. Such a function can be parametrized as $f_{\mathbf{A}}(\mathbf{x}):=\mathbf{x}^{T} \mathbf{A} \mathbf{x}$, where $\mathbf{A}$ is a square symmetric matrix containing the coefficients of the monomials. In order to control the complexity of a quadratic polynomial we choose a matrix norm $\|\cdot\|$ and consider only elements in a constrained set:

$$
\begin{equation*}
\mathcal{Q}_{\|\cdot\|, \lambda}:=\left\{f_{\mathbf{A}}(\mathbf{x})=\mathbf{x}^{T} \mathbf{A} \mathbf{x}:\|\mathbf{A}\| \leq \lambda\right\}, \quad \mathcal{Q}_{\|\cdot\|}:=\mathcal{Q}_{\|\cdot\|, 1} \tag{1}
\end{equation*}
$$

We quantify the complexity of such function classes using the classical notion of Rademacher complexity. It is well known that high probability generalization error bounds can be obtained in terms of this quantity (Koltchinskii and Panchenko, 2002; Bartlett and Mendelson, 2003). For this reason, we focus only on deriving upper bounds on this complexity measure.
Definition 1 (Rademacher complexity). Let $\sigma$ be uniformly distributed over the set $\{-1,1\}^{n}$ and let $\mathbf{X}_{n}=\left[\mathbf{x}_{1}, \ldots, \mathbf{x}_{n}\right] \subseteq \mathbb{R}^{d}$ be an i.i.d. sample drawn according to $\mu$. For a class of functions $\mathcal{F}: \mathbb{R}^{d} \rightarrow \mathbb{R}$ we define the empirical Rademacher complexity and the Rademacher complexity (with respect to $\mu$ ) of $\mathcal{F}$, respectively, as:

$$
\begin{equation*}
\hat{\mathcal{R}}\left(\mathcal{F} ; \mathbf{X}_{n}\right):=\mathbb{E}_{\sigma}\left[\sup _{f \in \mathcal{F}} \frac{1}{n} \sum_{i=1}^{n} \sigma_{i} f\left(\mathbf{x}_{i}\right)\right], \quad \mathcal{R}_{n, \mu}(\mathcal{F})=\mathbb{E}\left[\hat{\mathcal{R}}\left(\mathcal{F} ; \mathbf{X}_{n}\right)\right] \tag{2}
\end{equation*}
$$

Our bounds depend on the distribution through its intrinsic dimension (Tropp, 2015, Section 7), which measures how much the probability density concentrates near low-dimensional subspaces.

Definition 2. The intrinsic dimension of a distribution $\mu$ is the ratio $1 \leq r(\Sigma):=\operatorname{tr}(\Sigma) /\|\Sigma\|_{2} \leq d$.
We are now ready to state our main results about the Rademacher complexity of homogeneous quadratic polynomials with nuclear norm constraint:
Theorem 1. Let $\mathbf{x} \in \mathbb{R}^{d}$ such that $\|\mathbf{x}\|_{2}^{2} \lesssim \mathbb{E}\|\mathbf{x}\|_{2}^{2}$ almost surely and suppose $n \gtrsim r(\Sigma) \log d$. It holds that

$$
\begin{equation*}
\mathcal{R}_{n, \mu}\left(\mathcal{Q}_{\|\cdot\|_{\mathrm{tr}}, \lambda}\right) \lesssim \lambda \sqrt{\frac{r(\Sigma) \log d}{n}}\|\Sigma\|_{2} \tag{3}
\end{equation*}
$$

Now we proceed to prove Theorem 1. First, we only focus on the class $\mathcal{Q}_{\|\cdot\|}$ corresponding to the unit nuclear-norm ball. This is justified by well-known technical result Lemma 1, whose proof is included for completeness in Appendix D
Lemma 1. $\mathcal{R}_{n, \mu}\left(\mathcal{Q}_{\|\cdot\|, \lambda}\right) \leq \lambda \mathcal{R}_{n, \mu}\left(\mathcal{Q}_{\|\cdot\|}\right)$.
The backbone of Theorem 1 is Lemma 3, which relates the Rademacher complexity of a class of functions to concentration of empirical means to expectations with respect to the dual norm. It makes use of the technical Lemma 2, whose proof is included for completeness in Appendix B
Lemma 2. For all even $n \in \mathbb{N}$, it holds that $\sum_{k=0}^{n}|2 k-n|\binom{n}{k}<\sqrt{n} 2^{n}$
Lemma 3. Denote by $\|\cdot\|_{*}$ the dual norm of $\|\cdot\|$. Define $M_{k}:=\mathbb{E}\left\|\Sigma_{k}-\Sigma\right\|_{*}, \Sigma_{k}:=\frac{1}{k} \sum_{i=1}^{k} \mathbf{x}_{i} \mathbf{x}_{i}^{T}$ The Rademacher complexity of the class $\mathcal{Q}_{\|\cdot\|}$ can be upper bounded as follows:

$$
\begin{equation*}
\mathcal{R}_{n, \mu}\left(\mathcal{Q}_{\|\cdot\|}\right) \leq \frac{1}{n 2^{n-1}} \sum_{k=1}^{n} k\binom{n}{k} M_{k}+\frac{\|\Sigma\|_{*}}{\sqrt{n}} \tag{4}
\end{equation*}
$$

Proof. We first compute an upper bound on the empirical Rademacher complexity. The result will follow after taking expectation of the bound over the sample $\mathbf{X}_{n}=\left[\mathbf{x}_{1}, \ldots, \mathbf{x}_{n}\right]$. By definition of the dual norm, using the basic algebraic identity $\mathbf{x}^{T} \mathbf{A} \mathbf{x}=\left\langle\mathbf{A}, \mathbf{x x}^{T}\right\rangle$ we have

$$
\begin{equation*}
\hat{\mathcal{R}}\left(\mathcal{Q}_{\|\cdot\|} ; \mathbf{X}_{n}\right)=\frac{1}{n} \mathbb{E}_{\sigma} \sup _{\|\mathbf{A}\| \leq 1}\left\langle\mathbf{A}, \sum_{i=1}^{n} \sigma_{i} \mathbf{x}_{i} \mathbf{x}_{i}^{T}\right\rangle=\frac{1}{n} \mathbb{E}_{\sigma}\left\|\sum_{i=1}^{n} \sigma_{i} \mathbf{x}_{i} \mathbf{x}_{i}^{T}\right\|_{*} \tag{5}
\end{equation*}
$$

We now compute the expectation in eq. 5 There is a bijection between the possible configurations of the Rademacher variable $\sigma \in\{-1,1\}^{n}$ and the power set of [n], namely $\sigma \mapsto\left\{i \in[n]: \sigma_{i}=1\right\}$. This allows us to write eq. 5 as:

$$
\begin{equation*}
\hat{\mathcal{R}}\left(\mathcal{Q}_{\|\cdot\|} ; \mathbf{X}_{n}\right)=\frac{1}{n} \mathbb{E}_{\sigma}\left\|\sum_{i=1}^{n} \sigma_{i} \mathbf{x}_{i} \mathbf{x}_{i}^{T}\right\|_{*}=\frac{1}{n 2^{n}} \sum_{B \subseteq[n]} \underbrace{\left\|\sum_{i \in B} \mathbf{x}_{i} \mathbf{x}_{i}^{T}-\sum_{i \in B^{c}} \mathbf{x}_{i} \mathbf{x}_{i}^{T}\right\|_{*}}_{:=D_{B}} \tag{6}
\end{equation*}
$$

Let $\Sigma_{B}:=|B|^{-1} \sum_{i \in B} \mathbf{x}_{i} \mathbf{x}_{i}^{T}$. Using the triangle inequality, we can bound $D_{B}$ as:

$$
\begin{equation*}
D_{B} \leq|B|\left\|\Sigma_{B}-\Sigma\right\|_{*}+\left\|B\left|-\left|B^{c}\right|\right|\right\| \Sigma\left\|_{*}+\left|B^{c}\right|\right\| \Sigma_{B^{c}}-\Sigma \|_{*} \tag{7}
\end{equation*}
$$

To obtain a bound on the Rademacher complexity, we need now sum over $B \subseteq[n]$ the terms on the right hand side of eq. 7 , and take expectation with respect to the sample $\mathbf{X}_{n}$. First, we will deal with the sum of the second term in eq. 7, as it is actually a deterministic value. We can sum over $B \subseteq[n]$ by grouping together subsets $B$ of the same cardinality $|B|=k$. We obtain:

$$
\begin{equation*}
\sum_{B \subseteq[n]}\left\|B\left|-\left|B^{c}\right|\right|\right\| \Sigma\left\|_{*}=\sum_{k=0}^{n}\binom{n}{k}|2 k-n|\right\| \Sigma\left\|_{*} \leq \sqrt{n} 2^{n}\right\| \Sigma \|_{*} \tag{8}
\end{equation*}
$$

where the last inequality follows from Lemma 2
Finally, we compute the expectation of the sum over $B \subseteq[n]$ of the first and third term in eq. 7 . After taking the sum, both terms become equal by symmetry. It suffices to bound the sum of the first term.

Notice that because the variables $\mathbf{x}_{1}, \ldots, \mathbf{x}_{n}$ are i.i.d., the distribution of $\Sigma_{B}$ depends only on the size of the set $B$. Using the same counting argument as in eq. 8 we arrive at:

$$
\begin{align*}
\mathbb{E} \sum_{B \subseteq[n]}|B|\left\|\Sigma_{B}-\Sigma\right\|_{*} & =\sum_{B \subseteq[n]}|B| \mathbb{E}\left\|\Sigma_{B}-\Sigma\right\|_{*} \\
& =\sum_{k=1}^{n} k\binom{n}{k} \underbrace{\mathbb{E}\left\|\Sigma_{k}-\Sigma\right\|_{*}}_{=M_{k}} \tag{9}
\end{align*}
$$

Combining the bounds in eq. 8 and eq. 9 , and dividing by $n 2^{n}$ we obtain the result.

Note that the proof of Lemma 3 follows from technical arguments but the final result is not to be found in the literature, in this form or a similar one. In particular, it is completely unrelated to the result by Vershynin (2011) with which it only shares the fairly trivial split of Rademacher random variables preceding eq. 6 .
Lemma 3 provides a way to derive Rademacher complexity bounds from a bound on the expected deviations $M_{k}$ defined in Lemma 3, and might be of independent interest. In the particular case where the norm in consideration is the nuclear-norm, this lemma will be used to establish Theorem 1 as a simple application of a well-known non-asymptotic bound for the convergence of the empirical second moment (covariance) matrix to the true second moment (covariance) matrix.

Proof of Theorem 1. Recall that the dual norm of the nuclear-norm is the spectral-norm. The value of $M_{k}$ in Lemma 3 measures the average deviation of the empirical second moment matrix $\Sigma_{k}$ to the true $\Sigma$, in spectral-norm. Our assumption that $\|\mathbf{x}\|_{2}^{2} \lesssim \mathbb{E}\|\mathbf{x}\|_{2}^{2}$ almost surely, implies the concentration result (Vershynin, 2018, Theorem 5.6.1), which concludes that

$$
\begin{equation*}
M_{k} \lesssim\left(\sqrt{\frac{r(\Sigma) \log d}{k}}+\frac{r(\Sigma) \log d}{k}\right)\|\Sigma\|_{2} \tag{10}
\end{equation*}
$$

Plug this in eq. 4. and use the bound $\sqrt{k} \leq \sqrt{n}$ for $k \leq n$ to obtain the inequality

$$
\begin{equation*}
\mathcal{R}_{n, \mu}\left(\mathcal{Q}_{\|\cdot\|_{\text {tr }}}\right) \lesssim\left(\sqrt{\frac{r(\Sigma) \log d}{n}}+\frac{r(\Sigma) \log d}{n}\right)\|\Sigma\|_{2} \tag{11}
\end{equation*}
$$

By assumption $n \gtrsim r(\Sigma) \log d$, so that the first term in eq. 11 is the largest. The second term is of smaller order and thus ends up hidden by the notation $\lesssim$. We conclude

$$
\begin{equation*}
\mathcal{R}_{n, \mu}\left(\mathcal{Q}_{\|\cdot\|_{\mathrm{tr}}}\right) \lesssim \sqrt{\frac{r(\Sigma) \log d}{n}}\|\Sigma\|_{2} \tag{12}
\end{equation*}
$$

Invoking Lemma 1 we obtain the desired result.

### 2.1 Improvement upon previous work

We now show how our derived upper bound improves over the current best known bound by Kakade et al. (2012). To the best of our knowledge, the Rademacher complexity of quadratic classifiers with a nuclear-norm constraint has not been previously analyzed in a direct manner, as we do. Instead, the only existing bound Corollary 1) appears as a particular case of Theorem 2
Theorem 2 (Kakade et al. (2012) page 1876). Let

$$
\begin{equation*}
\mathcal{G}_{\|\cdot\|_{\mathrm{tr}}, \lambda}:=\left\{g_{\mathbf{A}}(\mathbf{X}):=\langle\mathbf{A}, \mathbf{X}\rangle:\|\mathbf{A}\|_{\operatorname{tr}} \leq \lambda\right\} \tag{13}
\end{equation*}
$$

be the class of nuclear-norm constrained linear functions over square $d \times d$ matrices. Let $\mu$ be a distribution supported on $\mathcal{X} \subseteq \mathbb{R}^{d \times d}$. It holds that:

$$
\mathcal{R}_{n, \mu}\left(\mathcal{G}_{\|\cdot\|_{\mathrm{tr}}, \lambda}\right) \lesssim \lambda X_{\infty} \sqrt{\frac{\log d}{n}}, \quad X_{\infty}=\sup _{\mathbf{X} \in \mathcal{X}}\|\mathbf{X}\|_{2}
$$

Corollary 1. Let $\mathrm{x} \in \mathcal{X} \subseteq \mathbb{R}^{d}$ be a random variable with distribution $\mu$. The Rademacher complexity of the class in eq. 1 can be bounded as

$$
\mathcal{R}_{n, \mu}\left(\mathcal{Q}_{\|\cdot\|_{\mathrm{tr}}, \lambda}\right) \lesssim \lambda x_{\infty} \sqrt{\frac{\log d}{n}}, \quad x_{\infty}:=\sup _{\mathbf{x} \in \mathcal{X}}\|\mathbf{x}\|_{2}^{2}
$$

Corollary 1 is a consequence of the fact that a function of the form $f_{\mathbf{A}}(\mathbf{x})=\mathbf{x}^{T} \mathbf{A} \mathbf{x}$ can be written as a linear classifier on matrices, $g_{\mathbf{A}}(\mathbf{X})=\langle\mathbf{A}, \mathbf{X}\rangle$, where $\mathbf{X}=\mathbf{x x}^{T}$. In this case, it is easy to see that $X_{\infty}=\sup _{\mathbf{x} \in \mathcal{X}}\left\|\mathbf{x} \mathbf{x}^{T}\right\|=\sup _{\mathbf{x} \in \mathcal{X}}\|\mathbf{x}\|_{2}^{2}=: x_{\infty}$. The only difference with our bound in Theorem 1 is that the term $x_{\infty}$ appears in place of $\sqrt{r(\Sigma)}\|\Sigma\|_{2}$.
In order to understand the difference between the two bounds, we turn to the analysis of the quotient between the bound in Corollary 1] and our bound Theorem 1. In Lemma 4 we show that this quotient can differ drastically, depending on the distribution.
Remark 1. The variables $X_{\infty}$ and $x_{\infty}$ defined respectively in Theorem 2 and Corollary 1 respectively, correspond to the supremum of a random variable. Because the Rademacher complexity arises as an expectation, it is clear that such quantities can be (and should be) replaced by the closely related measure-theoretic notion of essential supremum (denoted by ess sup): the least upper bound that holds almost surely. In this way the bounds are only tighter, and we believe this was the true intended definition by Kakade et al. (2012). In the following we will compare our bound to this tighter, modified bound.
Lemma 4. Let x be a random variable supported on a set $\mathcal{X} \subseteq \mathbb{R}^{d}$, and such that $\|\mathrm{x}\|_{2}^{2} \lesssim \mathbb{E}\|\mathrm{x}\|_{2}^{2}$ almost surely, then:

$$
\begin{equation*}
\sqrt{r(\Sigma)} \lesssim \frac{\operatorname{ess}_{\sup _{\mathbf{x} \in \mathcal{X}}\|\mathbf{x}\|_{2}^{2}}^{\sqrt{r(\Sigma)}\|\Sigma\|_{2}} \lesssim \sqrt{r(\Sigma)}}{\sqrt{r(\Sigma)}} \tag{14}
\end{equation*}
$$

Proof. By definition of essential supremum it holds that $\mathbb{E}\|\mathbf{x}\|_{2}^{2} \leq \operatorname{ess} \sup _{\mathbf{x} \in \mathcal{X}}\|\mathbf{x}\|_{2}^{2}$. Further, our assumption clearly implies that ess $\sup _{\mathbf{x} \in \mathcal{X}}\|\mathbf{x}\|_{2}^{2} \lesssim \mathbb{E}\|\mathbf{x}\|_{2}^{2}$. The identity $\mathbb{E}\|\mathbf{x}\|_{2}^{2}=\operatorname{tr}(\Sigma)$ and Definition 2 imply the result.

In summary, Lemma 4 shows that the baseline bound of Kakade et al. (2012) is larger by a square-root factor of the intrinsic dimension of the distribution (modulo global constants), compared to our bound in Theorem 1 Such factor ranges between 1 and the square root of the ambient dimension.

Precisely, when the intrinsic dimension of the distribution is equal to the ambient dimension, our bound enjoys a reduced dimension complexity. This is the case, for example, for isotropic distributions i.e., distributions such that their second moment matrix is the identity matrix. The dependency of the Rademacher complexity on the intrinsic dimension of the distribution is not revealed by the more general proof of Theorem 2 (Kakade et al., 2012).
Remark 2. The logarithmic term in Theorem 1 can be removed under the more restrictive dimensionindependent L-subgaussianity assumption (Mendelson and Zhivotovskiy, 2018).

## 3 Computable Generalization error bounds

Let $\mathbf{y}=\left(y_{1}, \ldots, y_{n}\right) \in\{-1,1\}^{n}$ be the labels associated with the data sample, and let

$$
\begin{gather*}
L(f):=\mathbb{P}\{\operatorname{sign}(f(\mathbf{x})) \neq \mathbf{y}\}  \tag{15}\\
\hat{L}\left(f ; \mathbf{X}_{n}\right):=\frac{1}{n} \sum_{i=1}^{n} \min \left(1, \max \left(0,1-y_{i} f\left(\mathbf{x}_{i}\right)\right)\right) \tag{16}
\end{gather*}
$$

be the missclassification probability and the empirical margin loss of the classifier $f$, respectively. It is well-known (Mohri et al., 2018, Theorem 5.8.) that with probability at least $1-\delta$, for all $f \in \mathcal{Q}_{\|\cdot\|_{\text {tr }}, \lambda}$ :

$$
\begin{equation*}
L(f) \lesssim \hat{L}\left(f ; \mathbf{X}_{n}\right)+\mathcal{R}_{n, \mu}\left(\mathcal{Q}_{\|\cdot\|_{\mathrm{tr}}, \lambda}\right)+\sqrt{\frac{\log \frac{1}{\delta}}{2 n}} \tag{17}
\end{equation*}
$$

This bound, together with the results in section 2. allow high probability uniform bounds on the misclassification error of a nuclear-norm constrained quadratic classifier.
However, the bound derived in this way is not actually computable: the Rademacher complexity bound in Theorem 1 depends on the second moment matrix of the distribution, which is unknown in practical applications. In the rest of this section we will overcome this drawback.
Rewriting our bound in eq. 3 as:

$$
\begin{equation*}
\mathcal{R}_{n, \mu}\left(\mathcal{Q}_{\|\cdot\|_{\mathrm{tr}}}\right) \lesssim \sqrt{\frac{\operatorname{tr} \Sigma\|\Sigma\|_{2} \log d}{n}} \tag{18}
\end{equation*}
$$

we observe that we need to estimate the trace and the largest eigenvalue of the second order moment matrix. Our hope is that the empirical estimators, based on the empirical second moment matrix $\Sigma_{n}$, will provide a good approximation of the true values.
We arrive at the following high probability bound on the Rademacher complexity in eq. 18, which is readily computable from samples and a bound on the diameter of the support of the distribution:
Theorem 3. Suppose that $\|\mathbf{x}\|_{2}^{2} \lesssim \mathbb{E}\|\mathbf{x}\|_{2}^{2}$ almost surely. This implies that $\|\mathbf{x}\|_{2}^{2} \leq B$ for some $B>0$. Define

$$
\begin{equation*}
K_{1}:=\frac{\log d}{n-\sqrt{n d\left(\log d+\log \frac{1}{\delta}\right)}} \quad K_{2}:=\frac{\sqrt{2 \log \frac{1}{\delta}} \log d}{2 n\left(\sqrt{n}-\sqrt{d\left(\log d+\log \frac{1}{\delta}\right)}\right)} \tag{19}
\end{equation*}
$$

Let $\delta>0$ and $f_{\mathbf{A}} \in \mathcal{Q}_{\|\cdot\|_{\mathrm{tr}}, \lambda}$. Provided $n \gtrsim d\left(\log d+\log \frac{1}{\delta}\right)$, with probability at least $1-3 \delta$ it holds that

$$
\begin{equation*}
\mathcal{R}_{n, \mu}\left(\mathcal{Q}_{\|\cdot\|_{\mathrm{tr}}}\right) \lesssim \underbrace{\sqrt{K_{1} \operatorname{tr} \Sigma_{n}\left\|\Sigma_{n}\right\|_{2}+B K_{2}\left\|\Sigma_{n}\right\|_{2}}}_{:=M(n, d, \delta)} \tag{20}
\end{equation*}
$$

Moreover, with probability at least $1-4 \delta$, uniformly for all $f_{\mathrm{A}} \in \mathcal{Q}_{\|\cdot\|_{\mathrm{tr}}, \lambda}$ it holds that

$$
\begin{equation*}
L\left(f_{\mathbf{A}}\right) \lesssim \hat{L}\left(f_{\mathbf{A}} ; \mathbf{X}_{n}\right)+\lambda M(n, d, \delta)+\sqrt{\frac{\log \frac{1}{\delta}}{2 n}} \tag{21}
\end{equation*}
$$

Proof. See Appendix C

## 4 Frobenius vs Nuclear-norm constraint

Perhaps the most common way to use quadratic features is to use a support vector machine (SVM) with the quadratic kernel $K(\mathbf{x}, \mathbf{y}):=\langle\mathbf{x}, \mathbf{y}\rangle^{2}$. Lemma 5, which is folklore in the kernel methods literature, precisely states that the RKHS norm constraint is equivalent to a Frobenius-norm constraint on the matrix of coefficients $\mathbf{A}$ of the underlying quadratic polynomial. Its proof is included for completeness in Appendix D
Lemma 5. Let $\mathbb{H}$ be the Reproducing Kernel Hilbert Space associated to the symmetric, positive semidefinite polynomial kernel $K(\mathbf{x}, \mathbf{y})=\langle\mathbf{x}, \mathbf{y}\rangle^{2}$, and denote its induced norm by $\|\cdot\|_{\mathbb{H}}$. Then $f \in \mathbb{H}$ if and only if there exists a symmetric matrix $\mathbf{A}$ such that $f(\mathbf{x})=\mathbf{x}^{T} \mathbf{A} \mathbf{x}$ and $\|f\|_{\mathbb{H}}=\|\mathbf{A}\|_{F}$.

For this reason, we now turn to compare the qualities of nuclear-norm and Frobenius-norm constrained quadratic classifiers. As a consequence of our derived bound (eq. 3), we uncover a fundamental difference between both regularization schemes Corollary 2): as the dimension increases, the growth rate of the complexity of nuclear-norm constrained quadratics strongly depends on the intrinsic dimension of the distribution. In contrast, that of Frobenius-norm constrained quadratics is insensitive to it.
In order to derive this rate, we need a way to argue about distributions across different dimensions and express our generalization bounds only in terms of dimension and number of samples. To this end, we introduce a natural boundedness assumption on the data distribution:
Assumption 1. $\|\mathbf{x}\|_{2}^{2} \approx d$ almost surely.

We now argue why this is a natural scaling order for this norm: if the entries of the random vector $\mathbf{x}$ are upper bounded as $\left|\mathbf{x}_{i}\right| \leq \alpha_{\max }$ and their average magnitude is lower bounded as $0<\alpha_{\text {min }} \leq \frac{1}{d} \sum_{i=1}^{d}\left|\mathbf{x}_{i}\right|$. Then

$$
\alpha_{\min } \sqrt{d} \leq \frac{1}{\sqrt{d}}\|\mathbf{x}\|_{1} \leq\|\mathbf{x}\|_{2} \leq \sqrt{d}\|\mathbf{x}\|_{\infty} \leq \sqrt{d} \alpha_{\max }
$$

and hence, Assumption 1 is satisfied. If we think about distributions of pixel (natural image) data of increasing resolution, we indeed have an upper bound on the intensity of each pixel. A lower bounded average pixel intensity only means that images are not arbitrarily dark, which often holds in practice. Any distribution of similar characteristics (e.g., sensor data) will probably satisfy our assumption.
We also introduce a growth condition on the intrinsic dimension, which states that $r(\Sigma) \approx d^{s}$ for some $0 \leq s \leq 1$. Rather than being an assumption, this condition helps to understand how our bounds change as the distribution falls between the two possible extremes given by the bound $1 \leq r(\Sigma) \leq d$. Traditionally, distributions that attain the lower bound $(s=0, r(\Sigma)=1$ ) or the upper bound ( $s=1$, $r(\Sigma)=d$ ) are called anisotropic or isotropic, respectively.
Corollary 2. Let Assumption 1 hold, and suppose that $r(\Sigma) \approx d^{s}$ for some $s \in[0,1]$ and $n \geq$ $r(\Sigma) \log d$. Then

$$
\begin{align*}
& \mathcal{R}_{n, \mu}\left(\mathcal{Q}_{\|\cdot\|_{\mathrm{tr}}, \lambda}\right) \lesssim \lambda \frac{d^{1-s / 2} \sqrt{\log d}}{\sqrt{n}}  \tag{22}\\
& \mathcal{R}_{n, \mu}\left(\mathcal{Q}_{\|\cdot\|_{F}, \lambda}\right) \approx \lambda \frac{d}{\sqrt{n}} \tag{23}
\end{align*}
$$

Proof. For the first inequality, note that $\|\mathbf{x}\|_{2}^{2} \approx d$ implies that $\operatorname{tr} \Sigma=\mathbb{E}\|\mathbf{x}\|_{2}^{2} \approx d$. Thus, $\frac{\operatorname{tr} \Sigma}{\|\Sigma\|_{2}}=$ $r(\Sigma) \approx d^{s} \Longrightarrow\|\Sigma\|_{2} \approx d^{1-s}$ Using these two identities in the inequality eq. 3 we obtain the first result.
For the second identity, Lemma 5 implies the class of Frobenius-norm constrained quadratic functions is equal to the ball of radius $\lambda$ in the RKHS corresponding to the quadratic kernel. By Mohri et al. (2018, Theorem 5.5) ${ }^{1}$, we have that

$$
\begin{equation*}
\hat{\mathcal{R}}\left(\mathcal{Q}_{\|\cdot\|_{F}, \lambda} ; \mathbf{X}_{n}\right) \approx \frac{\lambda}{n} \sqrt{\sum_{i=1}^{n} K\left(\mathbf{x}_{i}, \mathbf{x}_{i}\right)}=\frac{\lambda}{n} \sqrt{\sum_{i=1}^{n}\left\|\mathbf{x}_{i}\right\|_{2}^{4}} \approx \frac{\lambda d}{\sqrt{n}} \tag{24}
\end{equation*}
$$

where the final inequality comes from our assumption that $\|\mathbf{x}\|_{2}^{2} \approx d$. Taking expectation with respect to the sample we obtain that $\mathcal{R}_{n, \mu}\left(\mathcal{Q}_{\|\cdot\|_{F}, \lambda}\right) \approx \lambda \frac{d}{\sqrt{n}}$.

Corollary 2 then states that the nuclear norm constraint adapts to the intrinsic dimension of the distribution. In the worst case $(s=0)$ it grows linearly ${ }^{2}$ with dimension; In the best case $(s=1)$ it is much slower, and grows as the square root of the dimension.

Note that we can bring any distribution to approximate isotropic position, given a good approximation of $\Sigma$. Hence, it makes sense to expect performance gains if we perform such normalization procedure before training nuclear-norm constrained quadratic classifiers. In section 5we will put to test this claim in synthetic datasets, to illustrate the theory.

## 5 Experiments

The results derived in Corollary 2 pertain to the worst case generalization gap. This worst case gap is not guaranteed to be attained by the function found through empirical loss minimization. Therefore, in order to better test our results, we try to find the function that attains this worst case generalization gap.

[^0]The experiment we propose to illustrate the difference in intrinsic dimension sensitivity consists of computing the quadratic function $f$ such that $\mathbb{E}[\ell(f(\mathbf{x}, y))]-\frac{1}{n} \sum_{i=1}^{n} \ell\left(f\left(\mathbf{x}_{i}, y_{i}\right)\right)$ is maximized. The expectation in the term above being generally difficult to compute in closed form, we replace the objective by an empirical estimate

$$
\frac{1}{n_{\text {test }}} \sum_{k=1}^{n_{\text {test }}} \ell\left(f\left(\mathbf{x}_{k}, y_{k}\right)\right)-\frac{1}{n_{\text {train }}} \sum_{i=1}^{n_{\text {train }}} \ell\left(f\left(\mathbf{x}_{i}, y_{i}\right)\right)
$$

We expect this quantity to more faithfully track the upperbounds derived in Corollary 2 Indeed, what we evaluate here is the capacity of the class constraining $f$ to attain a low loss on $n_{\text {train }}$ samples while attaining a very high one on $n_{\text {test }}$ samples.
To illustrate that Frobenius and nuclear norm constrained classifiers exhibit different behaviors depending on the intrinsic dimension of the data, we compute the result of the maximization procedure on isotropic distributions on one hand and on anisotropic on the other. We set the radius $\lambda=1$ for both Nuclear and Frobenius norm constrained classifiers. We then observe the evolution of these quantities as the dimension increases.

Generating Isotropic Data To generate isotropic data satisfying our assumptions, we sample a standard Gaussian random vector in $\mathbb{R}^{d}$ and normalize to obtain i.i.d samples $\mathbf{x}_{i}$ that are uniformly distributed on the sphere of radius $\sqrt{d}$.

Generating Anistropic Data To generate anisotropic data, we first generate isotropic data features by uniformly sampling points on the sphere of radius $\sqrt{d}$ as described previously. We then transform the samples into more anisotropic ones by multiplying each feature vector coordinate-wise by a vector $\alpha^{(s)}$ defined as $\alpha_{i}^{(s)}=C_{s} \frac{1}{i^{s}}, i=1, \ldots, d$, for $s \in[0,1]$ where $C_{s}$ is chosen such that $\left\|\alpha^{(s)}\right\|_{2}=\sqrt{d}$. That, way, $s=0$ induces no change on the data, since $\alpha^{(1)}=[1,1, \ldots, 1]^{T}$, and larger $s$ implies that the data will be more squashed along the first few coordinates, inducing smaller intrinsic dimension. This transform does not affect the norm of the vectors since $\mathbb{E}\left[\left\|\alpha^{(s)} \times \mathbf{x}_{i}\right\|^{2}\right]=$ $\sum_{k=1}^{d}\left(\alpha_{k}^{(s)}\right)^{2} \mathbb{E}\left[\left[\mathbf{x}_{i}\right]_{k}^{2}\right]=\sum_{k=1}^{d}\left(\alpha_{k}^{(s)}\right)^{2}=\sqrt{d}$

Generating the labels We test two approaches for generating the labels: (1) We generate a random matrix $A$ with i.i.d standard Gaussian coordinates and set $y_{i}=\operatorname{sgn}\left(\mathbf{x}_{i}^{\top} A \mathbf{x}_{i}\right)$, (2) We set the labels randomly with $y_{i} \sim \operatorname{Bern}(0.5)$.


Figure 1: Maximized Train/Test loss gaps for each dimension $d$. The results are averaged over 5 independent runs. The error bars correspond to the standard error.

The results We observe that the maximal gaps do indeed evolve as the theory predicts. We indeed have the nuclear norm constrained classifier on isotropic data not exhibiting the same growth as the others, indicating that the maximal generalization gap for nuclear norm constrained classifiers is indeed sensitive to isotropy. Additional experiments and details are provided in Appendix E

## 6 Discussion

Our result shows that given a fixed regularization parameter $\lambda$, nuclear-norm constrained classifiers can take advantage of the properties of the data distribution to shave off a $\sqrt{d}$ factor from the Rademacher complexity, whereas Frobenius-norm constrained ones cannot.

We rely on two main elements to show this : Hölder's inequality in equation (5) which introduces the dual norms and a bound on the expected deviations $M_{k}$, which quantifies how fast the empirical moment estimate converges to the true moment matrix. Both these elements admit extensions to higher order tensors which corresponds to polynomials of higher degree.

The extension of Hölder's inequality is immediate. The main difficulty of extending our result lies in establishing the convergence rate of the deviations $\mathbb{E}\left[\left\|\frac{1}{k} \sum_{i=1}^{k} x_{i}^{\otimes m}-\mathcal{M}\right\|_{o p}\right]$, where $\mathcal{M}=\mathbb{E}\left[\mathbf{x}^{\otimes m}\right]$ is the tensor of $m$-th moments. For matrices, the rate is obtained through the matrix Bernstein inequality. Therefore, an extension can be derived by using the recently proved tensor Bernstein inequality in Luo et al. (2019) obtained by flattening the tensors and recycling the results for matrices.

## Acknowledgements

This work is funded (in part) through a PhD fellowship of the Swiss Data Science Center, a joint venture between EPFL and ETH Zurich. This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement no. 725594 - time-data). This project was sponsored by the Department of the Navy, Office of Naval Research (ONR) under a grant number N62909-17-1-2111. Research was sponsored by the Army Research Office and was accomplished under Grant Number W911NF-19-1-0404. This work was supported by the Swiss National Science Foundation (SNSF) under grant number 407540_167319.

## References

Amit, Y., Fink, M., Srebro, N., and Ullman, S. (2007). Uncovering shared structures in multiclass classification. In Proceedings of the 24th international conference on Machine learning, pages 17-24.

Bai, Y. and Lee, J. D. (2020). Beyond linearization: On quadratic and higher-order approximation of wide neural networks. In International Conference on Learning Representations.

Bartlett, P. L. and Mendelson, S. (2003). Rademacher and gaussian complexities: Risk bounds and structural results. J. Mach. Learn. Res., 3(null):463-482.

Bradbury, J., Frostig, R., Hawkins, P., Johnson, M. J., Leary, C., Maclaurin, D., Necula, G., Paszke, A., VanderPlas, J., Wanderman-Milne, S., and Zhang, Q. (2018). JAX: composable transformations of Python+NumPy programs.

Buitinck, L., Louppe, G., Blondel, M., Pedregosa, F., Mueller, A., Grisel, O., Niculae, V., Prettenhofer, P., Gramfort, A., Grobler, J., Layton, R., VanderPlas, J., Joly, A., Holt, B., and Varoquaux, G. (2013). API design for machine learning software: experiences from the scikit-learn project. In ECML PKDD Workshop: Languages for Data Mining and Machine Learning, pages 108-122.

Caruana, R. (1997). Multitask Learning. Machine Learning, 28(1):41-75.
Du, S. and Lee, J. (2018). On the power of over-parametrization in neural networks with quadratic activation. In Dy, J. and Krause, A., editors, Proceedings of the 35th International Conference on Machine Learning, volume 80 of Proceedings of Machine Learning Research, pages 1329-1338. PMLR.

Duchi, J., Shalev-Shwartz, S., Singer, Y., and Chandra, T. (2008). Efficient projections onto the 1 1-ball for learning in high dimensions. In Proceedings of the 25th international conference on Machine learning, pages 272-279.

Fan, R.-E., Chang, K.-W., Hsieh, C.-J., Wang, X.-R., and Lin, C.-J. (2008). Liblinear: A library for large linear classification. J. Mach. Learn. Res., 9:1871-1874.

Ghorbani, B., Mei, S., Misiakiewicz, T., and Montanari, A. (2020). When do neural networks outperform kernel methods?

Jacot, A., Gabriel, F., and Hongler, C. (2018). Neural tangent kernel: convergence and generalization in neural networks. In Proceedings of the 32nd International Conference on Neural Information Processing Systems, pages 8580-8589.

Jayakumar, S. M., Czarnecki, W. M., Menick, J., Schwarz, J., Rae, J., Osindero, S., Teh, Y. W., Harley, T., and Pascanu, R. (2020). Multiplicative interactions and where to find them. In International Conference on Learning Representations.

Jiang*, Y., Neyshabur*, B., Mobahi, H., Krishnan, D., and Bengio, S. (2020). Fantastic generalization measures and where to find them. In International Conference on Learning Representations.

Kakade, S. M., Shalev-Shwartz, S., and Tewari, A. (2012). Regularization techniques for learning with matrices. Journal of Machine Learning Research, 13(Jun):1865-1890.

Koltchinskii, V. and Panchenko, D. (2002). Empirical margin distributions and bounding the generalization error of combined classifiers. Ann. Statist., 30(1):1-50.

Luo, Z., Qi, L., and Toint, P. L. (2019). Bernstein concentration inequalities for tensors via einstein products. arXiv preprint arXiv:1902.03056.

Mendelson, S. and Zhivotovskiy, N. (2018). Robust covariance estimation under $L_{4}-L_{2}$ norm equivalence. arXiv e-prints, page arXiv:1809.10462.

Mohri, M., Rostamizadeh, A., and Talwalkar, A. (2018). Foundations of Machine Learning. The MIT Press, 2nd edition.

Pontil, M. and Maurer, A. (2013). Excess risk bounds for multitask learning with trace norm regularization. In Shalev-Shwartz, S. and Steinwart, I., editors, Proceedings of the 26th Annual Conference on Learning Theory, volume 30 of Proceedings of Machine Learning Research, pages 55-76, Princeton, NJ, USA. PMLR.

Pope, P., Zhu, C., Abdelkader, A., Goldblum, M., and Goldstein, T. (2021). The intrinsic dimension of images and its impact on learning. In International Conference on Learning Representations.

Srebro, N. (2005). Learning with matrix factorizations. PhD. Thesis.
Srebro, N., Rennie, J., and Jaakkola, T. S. (2005). Maximum-margin matrix factorization. In Saul, L. K., Weiss, Y., and Bottou, L., editors, Advances in Neural Information Processing Systems 17, pages 1329-1336. MIT Press.

Srebro, N. and Shraibman, A. (2005). Rank, trace-norm and max-norm. In Auer, P. and Meir, R., editors, Learning Theory, volume 3559, pages 545-560. Springer Berlin Heidelberg. Series Title: Lecture Notes in Computer Science.

Tropp, J. A. (2015). An introduction to matrix concentration inequalities. Foundations and Trends ${ }^{\circledR}$ in Machine Learning, 8(1-2):1-230.

Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A. N., Kaiser, u., and Polosukhin, I. (2017). Attention is all you need. In Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS'17, page 6000-6010, Red Hook, NY, USA. Curran Associates Inc.

Vershynin, R. (2011). A simple decoupling inequality in probability theory (2011. http:// Www-personal.umich.edu/~romanv/papers/decoupling-simple.pdf [Online; accessed 1-February-2021].

Vershynin, R. (2018). High-Dimensional Probability: An Introduction with Applications in Data Science. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press.

Wimalawarne, K., Tomioka, R., and Sugiyama, M. (2016). Theoretical and experimental analyses of tensor-based regression and classification. Neural Computation, 28(4):686-715. Publisher: MIT Press.

Yehudai, G. and Shamir, O. (2019). On the power and limitations of random features for understanding neural networks. 32 .

Yu, H.-F., Jain, P., Kar, P., and Dhillon, I. (2014). Large-scale multi-label learning with missing labels. In Xing, E. P. and Jebara, T., editors, Proceedings of the 31st International Conference on Machine Learning, volume 32 of Proceedings of Machine Learning Research, pages 593-601, Bejing, China. PMLR.

## Checklist

1. For all authors...
(a) Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope? [Yes]
(b) Did you describe the limitations of your work? [Yes]
(c) Did you discuss any potential negative societal impacts of your work? [No] The work is of theoretical nature and its societal impacts are difficult to foresee.
(d) Have you read the ethics review guidelines and ensured that your paper conforms to them? [Yes]
2. If you are including theoretical results...
(a) Did you state the full set of assumptions of all theoretical results? [Yes]
(b) Did you include complete proofs of all theoretical results? [Yes] They are given in the Appendix
3. If you ran experiments...
(a) Did you include the code, data, and instructions needed to reproduce the main experimental results (either in the supplemental material or as a URL)? [Yes]
(b) Did you specify all the training details (e.g., data splits, hyperparameters, how they were chosen)? [Yes]
(c) Did you report error bars (e.g., with respect to the random seed after running experiments multiple times)? [Yes]
(d) Did you include the total amount of compute and the type of resources used (e.g., type of GPUs, internal cluster, or cloud provider)? [No]
4. If you are using existing assets (e.g., code, data, models) or curating/releasing new assets...
(a) If your work uses existing assets, did you cite the creators? [Yes]
(b) Did you mention the license of the assets? [Yes]
(c) Did you include any new assets either in the supplemental material or as a URL? [Yes]
(d) Did you discuss whether and how consent was obtained from people whose data you're using/curating? [N/A]
(e) Did you discuss whether the data you are using/curating contains personally identifiable information or offensive content? [N/A]
5. If you used crowdsourcing or conducted research with human subjects...
(a) Did you include the full text of instructions given to participants and screenshots, if applicable? [N/A]
(b) Did you describe any potential participant risks, with links to Institutional Review Board (IRB) approvals, if applicable? [N/A]
(c) Did you include the estimated hourly wage paid to participants and the total amount spent on participant compensation? [N/A]

## A Appendix

## B Proof of Lemma 2

For completeness, we include the following result, taken from https://math.stackexchange com/questions/3209660/show-that-2m-choose-m-leq-frac22m-sqrt2m

Lemma 6. For $m \in \mathbb{N}$ we have

$$
\begin{equation*}
\binom{2 m}{m}<\frac{2^{2 m}}{\sqrt{2 m+1}}<\frac{2^{2 m}}{\sqrt{2 m}} \tag{25}
\end{equation*}
$$

Proof. By induction on $m \in \mathbb{N}$. For $m=1$ the inequality holds. Suppose it holds for some $m \in \mathbb{N}$

$$
\begin{equation*}
\frac{(2 m+2)!}{2^{2 m+2}((m+1)!)^{2}}=\frac{(2 m)!}{2^{2 m}(m!)^{2}} \cdot \frac{2 m+1}{2 m+2}<\frac{1}{\sqrt{2 m+1}} \cdot \frac{2 m+1}{2 m+2} \stackrel{?}{\leq} \frac{1}{\sqrt{2 m+3}} \tag{26}
\end{equation*}
$$

So it remains to prove that

$$
\begin{equation*}
4 m^{2}+8 m+3=(\sqrt{2 m+1} \sqrt{2 m+3})^{2} \leq(2 m+2)^{2}=4 m^{2}+8 m+4 \tag{27}
\end{equation*}
$$

which holds.

Lemma 7. For all $n \in \mathbb{N}$ even:

$$
\sum_{k=0}^{n}\binom{n}{k}|2 k-n|=n\binom{n}{\frac{n}{2}}
$$

Proof. Suppose that $n=2 m$ is even.

$$
\begin{aligned}
\sum_{k=0}^{n}\binom{n}{k}|2 k-n| & =\sum_{k=0}^{m}\binom{2 m}{k}(2 m-2 k)+\sum_{k=m+1}^{2 m}\binom{2 m}{k}(2 k-2 m) \\
& =2 \sum_{k=0}^{m}\binom{2 m}{k}(m-k)+\sum_{k=0}^{m-1}\binom{2 m}{2 m-k}(2(2 m-k)-2 m) \\
& =2 \sum_{k=0}^{m}\binom{2 m}{k}(m-k)+2 \sum_{k=0}^{m-1}\binom{2 m}{k}(m-k) \\
& =4 \sum_{k=0}^{m}\binom{2 m}{k}(m-k)
\end{aligned}
$$

We now show by induction that $\forall m \in \mathbb{N}$,

$$
\begin{equation*}
\sum_{k=0}^{m}\binom{2 m}{k} k=m 2^{2 m-1} \tag{28}
\end{equation*}
$$

The result holds trivially true for $m=0$. Suppose it is true for some $m \geq 0$ and let's prove it remains true for $m+1$ :

$$
\begin{aligned}
\sum_{k=0}^{m+1}\binom{2 m+2}{k} k & =\sum_{k=0}^{m+1}\left(\binom{2 m}{k}+2\binom{2 m}{k-1}+\binom{2 m}{k-2}\right) k \\
& =\sum_{k=0}^{m+1}\binom{2 m}{k} k+2 \sum_{k=1}^{m+1}\binom{2 m}{k-1} k+\sum_{k=2}^{m+1}\binom{2 m}{k-2} k \\
& =\sum_{k=0}^{m+1}\binom{2 m}{k} k+2 \sum_{k=0}^{m}\binom{2 m}{k}(k+1)+\sum_{k=0}^{m-1}\binom{2 m}{k}(k+2) \\
& =4 \sum_{k=0}^{m}\binom{2 m}{k} k+\binom{2 m}{m+1}(m+1)+2 \sum_{k=0}^{m}\binom{2 m}{k}-\binom{2 m}{m} m+2 \sum_{k=0}^{m-1}\binom{2 m}{k} \\
& =4 m 2^{2 m-1}+2 \sum_{k=0}^{2 m}\binom{2 m}{k} \\
& =m 2^{2(m+1)-1}+22^{2 m} \\
& =(m+1) 2^{2(m+1)-1} .
\end{aligned}
$$

We thus proved that eq. 28 is true for all $m \in \mathbb{N}$. Thus, we have:

$$
\begin{aligned}
\sum_{k=0}^{m}\binom{2 m}{k}(m-k) & =m \sum_{k=0}^{m}\binom{2 m}{k}-m 2^{2 m-1} \\
& =\frac{m}{2}\left(\sum_{k=0}^{m}\binom{2 m}{k}+\sum_{k=m}^{2 m}\binom{2 m}{k}\right)-m 2^{2 m-1} \\
& =\frac{m}{2}\left(\sum_{k=0}^{2 m}\binom{2 m}{k}+\binom{2 m}{m}\right)-m 2^{2 m-1} \\
& =\frac{m}{2}\left(2^{2 m}+\binom{2 m}{m}\right)-m 2^{2 m-1} \\
& =\frac{m}{2}\binom{2 m}{m}
\end{aligned}
$$

which concludes the proof.

Proof of Lemma 2. Combine Lemma 6and Lemma 7.

## C Proof of Theorem 3

Lemma 8. Suppose that $n \gtrsim d(\log (d)+\log (1 / \delta))$ and that $\|\mathbf{x}\|_{2}^{2} \lesssim \mathbb{E}\|\mathbf{x}\|_{2}^{2}$ almost surely. With probability at least $1-2 \delta$ :

$$
\begin{equation*}
\|\Sigma\|_{2} \lesssim\left(1-\sqrt{\frac{d\left(\log d+\log \frac{1}{\delta}\right)}{n}}\right)^{-1}\left\|\Sigma_{n}\right\|_{2} \tag{29}
\end{equation*}
$$

Proof. The lemma is an application of the matrix Bernstein inequality. Recall that $\Sigma_{n}=\sum_{i=1}^{n} \mathbf{x}_{i} \mathbf{x}_{i}^{T}$ where $\mathbf{x}_{i}$ for $i \in\{1, \ldots, n\}$ are independent, identically distributed and, by assumption, satisfy the inequality $\left\|\mathbf{x}_{i}\right\|_{2}^{2} \lesssim \mathbb{E}\left\|\mathbf{x}_{i}\right\|_{2}^{2}=\operatorname{tr}(\Sigma)$ almost surely. Spelling out the definition of $\lesssim$, we have that there exists a $C \geq 1$ such that

$$
\forall i \in\{1, \ldots, n\}, \quad\left\|\mathbf{x}_{i} \mathbf{x}_{i}^{T}\right\|_{2}=\left\|\mathbf{x}_{i}\right\|_{2}^{2} \leq C \operatorname{tr}(\Sigma) \quad \text { almost surely } .
$$

Consequently, by matrix Bernstein, for any $t \geq 0$,

$$
\mathbb{P}\left(\left\|\sum_{i=1}^{n}\left(\mathbf{x}_{i} \mathbf{x}_{i}^{T}-\Sigma\right)\right\|_{2} \geq t\right) \leq 2 d \exp \left(-\frac{\frac{t^{2}}{2}}{\sigma^{2}+K \frac{t}{3}}\right)
$$

where $K=C \operatorname{tr}(\Sigma)$ and $\sigma^{2}=\left\|\sum_{i=1}^{n} \mathbb{E}\left[\left(\mathbf{x}_{i} \mathbf{x}_{i}^{T}-\Sigma\right)^{2}\right]\right\|_{2}$.
Now for any $0<\delta \leq 1$, we take

$$
t=\sqrt{2}\left(\sigma \sqrt{\log d+\log \frac{1}{\delta}}+K\left(\log d+\log \frac{1}{\delta}\right)\right)
$$

The term inside the exponential on the right hand side of the Bernstein bound then becomes

$$
\begin{aligned}
-\frac{\frac{t^{2}}{2}}{\sigma^{2}+K \frac{t}{3}} & =-\frac{\sigma^{2}\left(\log d+\log \frac{1}{\delta}\right)+2 K \sigma \sqrt{\log d+\log \frac{1}{\delta}}\left(\log d+\log \frac{1}{\delta}\right)+K^{2}\left(\log d+\log \frac{1}{\delta}\right)^{2}}{\sigma^{2}+\frac{\sqrt{2} K}{3} \sigma \sqrt{\log d+\log \frac{1}{\delta}}+\frac{\sqrt{2} K^{2}}{3}\left(\log d+\log \frac{1}{\delta}\right)} \\
& =-\frac{\left(\log d+\log \frac{1}{\delta}\right)\left(\sigma^{2}+2 K \sigma \sqrt{\log d+\log \frac{1}{\delta}}+K^{2}\left(\log d+\log \frac{1}{\delta}\right)\right)}{\sigma^{2}+\frac{\sqrt{2} K}{3} \sigma \sqrt{\log d+\log \frac{1}{\delta}}+\frac{\sqrt{2} K^{2}}{3}\left(\log d+\log \frac{1}{\delta}\right)} \\
& \leq-\left(\log d+\log \frac{1}{\delta}\right) \quad \quad\left(\text { since } \frac{\sqrt{2}}{3} \leq 1 \text { and } \log d+\log \frac{1}{\delta} \geq 0\right)
\end{aligned}
$$

Consequently, the event

$$
\begin{equation*}
\left\|\sum_{i=1}^{n}\left(\mathbf{x}_{i} \mathbf{x}_{i}^{T}-\Sigma\right)\right\|_{2} \leq \sqrt{2}\left(\sigma \sqrt{\log d+\log \frac{1}{\delta}}+K\left(\log d+\log \frac{1}{\delta}\right)\right) \tag{30}
\end{equation*}
$$

holds with probability at least $1-2 \delta$. Now, by proceding just like Theorem 5.6.1 of Vershynin (2018), we can bound $\sigma^{2}$ :

$$
\begin{array}{rlr}
\sigma^{2} & =\left\|\sum_{i=1}^{n} \mathbb{E}\left[\left(\mathbf{x}_{i} \mathbf{x}_{i}^{T}-\Sigma\right)^{2}\right]\right\|_{2} & \\
& =n\left\|\mathbb{E}\left[\left(\mathbf{x}_{1} \mathbf{x}_{1}^{T}-\Sigma\right)^{2}\right]\right\|_{2} & \\
& =n\left\|\mathbb{E}\left[\left(\mathbf{x}_{1} \mathbf{x}_{1}^{T}\right)^{2}\right]-\Sigma^{2}\right\|_{2} & \text { (by expanding the square) } \\
& \leq n\left\|\mathbb{E}\left[\left(\mathbf{x}_{1} \mathbf{x}_{1}^{T}\right)^{2}\right]\right\|_{2} & \text { (since } \Sigma \succeq 0 \text { ) } \\
& =n\left\|\mathbb{E}\left[\left\|\mathbf{x}_{1}\right\|_{2}^{2}\left(\mathbf{x}_{1} \mathbf{x}_{1}^{T}\right)\right]\right\|_{2} & \\
& \leq n C \operatorname{tr}(\Sigma)\left\|\mathbb{E}\left[\mathbf{x}_{1} \mathbf{x}_{1}^{T}\right]\right\|_{2} & \text { (by assumption) } \\
& =n C \operatorname{tr}(\Sigma)\|\Sigma\|_{2} .
\end{array}
$$

Plugging this bound into and dividing by $n$, we find that

$$
\left\|\Sigma_{n}-\Sigma\right\|_{2} \leq \sqrt{2}\left(\sqrt{n C \operatorname{tr}(\Sigma)\|\Sigma\|_{2}} \frac{\sqrt{\log d+\log \frac{1}{\delta}}}{n}+C \operatorname{tr}(\Sigma) \frac{\left(\log d+\log \frac{1}{\delta}\right)}{n}\right)
$$

Factorizing by $\|\Sigma\|_{2}$ on the RHS and using the assumption that $C d\left(\log d+\log \frac{1}{\delta}\right) \lesssim n$, we find

$$
\left\|\Sigma_{n}-\Sigma\right\|_{2} \lesssim \sqrt{\frac{r(\Sigma)\left(\log d+\log \frac{1}{\delta}\right)}{n}}\|\Sigma\|_{2}
$$

With this in hand, we simply use the triangle inequality to obtain

$$
\begin{aligned}
\|\Sigma\|_{2} & =\left\|\Sigma-\Sigma_{n}+\Sigma_{n}\right\|_{2} \\
& =\left\|\Sigma_{n}-\Sigma\right\|_{2}+\left\|\Sigma_{n}\right\|_{2} \\
& \lesssim \sqrt{\frac{r(\Sigma)\left(\log d+\log \frac{1}{\delta}\right)}{n}}\|\Sigma\|_{2}+\left\|\Sigma_{n}\right\|_{2}
\end{aligned}
$$

Isolating $\|\Sigma\|_{2}$ on the left hand side, we find that

$$
\|\Sigma\|_{2} \lesssim\left(1-\sqrt{\frac{r(\Sigma)\left(\log d+\log \frac{1}{\delta}\right)}{n}}\right)^{-1}\left\|\Sigma_{n}\right\|_{2}
$$

Finally we have that $r(\Sigma) \leq d$.
Lemma 9. Suppose that $\|\mathbf{x}\|_{2}^{2} \leq B$ almost surely. With probability at least $1-\delta$ :

$$
\begin{equation*}
\operatorname{tr} \Sigma \leq \Sigma_{n}+\sqrt{\frac{B^{2} \log (1 / \delta)}{2 n}} \tag{31}
\end{equation*}
$$

Proof. We define

$$
\begin{equation*}
\sigma\left(\mathbf{x}_{1}, \ldots, \mathbf{x}_{n}\right)=\frac{1}{n} \sum_{i=1}^{n} \operatorname{tr}\left(\mathbf{x}_{i} \mathbf{x}_{i}^{T}\right)=\operatorname{tr} \Sigma_{n} \tag{32}
\end{equation*}
$$

The empirical second moment matrix is an unbiased estimator of the second moment matrix $\Sigma=\mathbb{E} \Sigma_{n}$. Hence

$$
\begin{equation*}
\operatorname{tr} \Sigma=\operatorname{tr} \mathbb{E} \Sigma_{n}=\mathbb{E} \operatorname{tr} \Sigma_{n}=\mathbb{E} \sigma\left(\mathbf{x}_{1}, \ldots, \mathbf{x}_{n}\right) \tag{33}
\end{equation*}
$$

We observe that the quantity we want to bound $\left|\operatorname{tr} \Sigma-\operatorname{tr} \Sigma_{n}\right|$ is precisely the deviation of $\sigma$ to its expected value.
In order to obtain a high-probability bound on this deviation, we will make use of McDiarmid's Inequality. We first show that $\sigma$ satisfies the bounded difference inequality
$\left|\sigma\left(\mathbf{x}_{1}, \ldots, \mathbf{x}_{i}, \ldots, \mathbf{x}_{n}\right)-\sigma\left(\mathbf{x}_{1}, \ldots, \hat{\mathbf{x}}_{i}, \ldots, \mathbf{x}_{n}\right)\right|=\frac{1}{n}\left|\operatorname{tr}\left(\mathbf{x}_{i} \mathbf{x}_{i}^{T}-\hat{\mathbf{x}}_{i} \hat{\mathbf{x}}_{i}^{T}\right)\right|=\frac{1}{n}\left|\left\|\mathbf{x}_{i}\right\|^{2}-\left\|\hat{\mathbf{x}}_{i}^{2}\right\| \|\right| \leq \frac{B}{n}$
Hence by McDiarmid's inequality we conclude that

$$
P\left(\operatorname{tr} \Sigma_{n}-\operatorname{tr} \Sigma<-t\right) \leq \exp \left(\frac{-2 n t^{2}}{B^{2}}\right)
$$

Letting $\delta=\exp \left(\frac{-2 n t^{2}}{B^{2}}\right)$ we have $t=\sqrt{\frac{B^{2} \log (1 / \delta)}{2 n}}$ so that with probability at least $1-\delta$

$$
\begin{equation*}
\operatorname{tr} \Sigma \leq \Sigma_{n}+\sqrt{\frac{B^{2} \log (1 / \delta)}{2 n}} \tag{36}
\end{equation*}
$$

Proof of Theorem 3. The first result follows after bounding the Rademacher Complexity using Lemma 8 and Lemma 9 in the inequality 18 . The probability is obtained with a union bound.

From this result, the final computable uniform generalization bound in eq. 21 follows from eq. 17 and a union bound.

## D Proof of Lemma 1and Lemma 5

Proof of Lemma 1. Dividing by $n$ and taking expectation on the following inequality yields the result.

$$
\begin{equation*}
\sup _{\|A\| \leq \lambda} \sum_{i=1}^{n} \sigma_{i} x_{i}^{T} A x_{i}=\sup _{\|A\| \leq \lambda}\|A\| \sum_{i=1}^{n} \sigma_{i} x_{i}^{T}\left(\|A\|^{-1} A\right) x_{i} \leq \lambda \sup _{\|A\| \leq 1} \sum_{i=1}^{n} \sigma_{i} x_{i}^{T} A x_{i} \tag{37}
\end{equation*}
$$

Proof of Lemma 5. By the Moore-Aaronszajn Theorem, the RKHS corresponding to the kernel $K(x, y)=\langle x, y\rangle^{2}$, denoted by $\mathbb{H}$, is built upon the vector space $\mathbb{H}_{0}$ corresponding to the linear span of functions of the form $K_{y}=K(\cdot, y)=\langle\cdot, y\rangle^{2}$ for $y \in \mathbb{R}^{d}$. We equip $\mathbb{H}_{0}$ with the inner product:

$$
\begin{equation*}
\left\langle\sum_{i=1}^{n} a_{i} K_{y_{i}}, \sum_{i=j}^{m} b_{j} K_{z_{j}}\right\rangle:=\sum_{i=1}^{n} \sum_{i=j}^{m} a_{i} b_{j} K\left(y_{i}, z_{j}\right) \tag{38}
\end{equation*}
$$

$\mathbb{H}$ is precisely the completion of $\mathbb{H}_{0}$, under the metric induced by this inner product. However, because any $K_{y}$ is a homogenous polynomial of second degree, $\mathbb{H}_{0}$ is finite-dimensional, hence closed and $\mathbb{H}=\mathbb{H}_{0}$. This shows that $\mathbb{H}$ is a subspace of the space of homogenous polynomials of second degree so $f \in \mathbb{H}$ can be represented as $f(x)=x^{T} A x$ for some symmetric matrix $A$.
On the other hand, for any symmetric matrix we can represent the function $f(x)=x^{T} A x$ as

$$
\begin{equation*}
f(x)=x^{T} A x=x^{T} U S U^{T} x=\sum_{i=1}^{d} S_{i i}\left\langle x, u_{i}\right\rangle^{2}=\sum_{i=1}^{d} S_{i i} K\left(x, u_{i}\right) \tag{39}
\end{equation*}
$$

where $A=U S U^{T}$ is the SVD (orthogonal diagonalization) of the symmetric matrix A , and $u_{i}$ are the columns of $U$. So any homogeneous polynomial is in the linear span of the functions of the form $K_{y}=K(\cdot, y)$. We conclude that, as a set, $\mathbb{H}$ is equal to the space of homogeneous polynomials of second degree.

Now we show that its norm is equal to the Frobenius norm of the associated matrix of coefficients. Let $f(x)=x^{T} A x \in \mathbb{H}$, where $A$ is a symmetric matrix. Let $A=U S U^{T}$ be the SVD of $A$. We know that $f=\sum_{i=1}^{d} S_{i i} K_{u_{i}}$ so that by definition of the RKHS norm

$$
\begin{align*}
\|f\|_{\mathbb{H}}^{2} & =\left\langle\sum_{i=1}^{d} S_{i i} K_{u_{i}}, \sum_{i=1}^{d} S_{i i} K_{u_{i}}\right\rangle_{\mathbb{H}}=\sum_{i=1}^{d} \sum_{j=1}^{d} S_{i i} S_{j j}\left\langle K_{u_{i}}, K_{u_{j}}\right\rangle \\
& =\sum_{i=1}^{d} \sum_{j=1}^{d} S_{i i} S_{j j} K\left(u_{i}, u_{j}\right)=\sum_{i=1}^{d} \sum_{j=1}^{d} S_{i i} S_{j j}\left\langle u_{i}, u_{j}\right\rangle^{2}=\sum_{i=1}^{d} S_{i i}^{2}=\|A\|_{F}^{2} \tag{40}
\end{align*}
$$

## E Additional experiments

## E. 1 A synthetic experiment maximizing the train/test loss gap

The results derived in Corollary 2 pertain to the worst case generalization gap. This worst case gap is not guaranteed to be attained by the function found through empirical loss minimization. Therefore, in order to better test our results, we try to find the function that attains this worst case generalization gap.
The experiment we propose consists of computing the quadratic function $f$ such that

$$
\mathbb{E}[\ell(f(\mathbf{x}, y))]-\frac{1}{n} \sum_{i=1}^{n} \ell\left(f\left(\mathbf{x}_{i}, y_{i}\right)\right)
$$

is maximized. The expectation in the term above being generally difficult to compute in closed form, we replace the objective by an empirical estimate

$$
\frac{1}{n_{\text {test }}} \sum_{k=1}^{n_{\text {test }}} \ell\left(f\left(\mathbf{x}_{k}, y_{k}\right)\right)-\frac{1}{n_{\text {train }}} \sum_{i=1}^{n_{\text {train }}} \ell\left(f\left(\mathbf{x}_{i}, y_{i}\right)\right)
$$

We expect this quantity to more faithfully track the upperbounds derived in Corollary 2 Indeed, what we evaluate here is the capacity of the class constraining $f$ to attain a low loss on $n_{\text {train }}$ samples while attaining a very high one on $n_{\text {test }}$ samples.
To illustrate that Frobenius and nuclear norm constrained classifiers exhibit different behaviors depending on the intrinsic dimension of the data, we compute the result of the maximization procedure on isotropic distributions on one hand and on anisotropic on the other. We set the radius $\lambda=1$ for both Nuclear and Frobenius norm constrained classifiers. We then observe the evolution of these quantities as the dimension increases.

Generating Isotropic Data To generate isotropic data satisfying our assumptions, we sample a standard Gaussian random vector in $\mathbb{R}^{d}$ and normalize it as follows :

$$
\mathbf{x}_{i}=\sqrt{d} \times \frac{z_{i}}{\left\|z_{i}\right\|} \quad \text { where } z_{i} \sim \mathcal{N}\left(0, I_{d}\right)
$$

The $z_{i}$ 's are sampled independently. The procedure generates i.i.d samples $\mathbf{x}_{i}$ that are uniformly distributed on the sphere of radius $\sqrt{d}$.

Generating Anistropic Data To generate anisotropic data, we first generate isotropic data features by uniformly sampling points on the sphere of radius $\sqrt{d}$ as described previously. We then transform the samples into more anisotropic ones.
The transform consists of multiplying each feature vector coordinate-wise by a vector $\alpha^{(s)}$ defined as $\alpha_{i}^{(s)}=C_{s} \frac{1}{i^{s}}, i=1, \ldots, d$, for $s \in[0,1]$ where $C_{s}$ is chosen such that $\left\|\alpha^{(s)}\right\|_{2}=\sqrt{d}$. That, way, $s=0$ induces no change on the data, since $\alpha^{(1)}=[1,1, \ldots, 1]^{T}$, and larger $s$ implies that the data will be more squashed along the first few coordinates, inducing smaller intrinsic dimension. This transform does not affect the norm of the vectors since

$$
\mathbb{E}\left[\left\|\alpha^{(s)} \times \mathbf{x}_{i}\right\|^{2}\right]=\sum_{k=1}^{d}\left(\alpha_{k}^{(s)}\right)^{2} \mathbb{E}\left[\left[\mathbf{x}_{i}\right]_{k}^{2}\right]=\sum_{k=1}^{d}\left(\alpha_{k}^{(s)}\right)^{2}=\sqrt{d}
$$

Generating the labels We test two approaches for generating the labels:

1. We generate a random matrix $A$ with i.i.d standard Gaussian coordinates and set $y_{i}=$ $\operatorname{sgn}\left(\mathbf{x}_{i}^{\top} A \mathbf{x}_{i}\right)$.
2. We set the labels randomly with $y_{i} \sim \operatorname{Bern}(0.5)$.


Figure 2: Maximized Train/Test loss gaps for each dimension $d$. The results are averaged over 5 independent runs. The error bars correspond to the standard error.

The results We observe that the maximal gaps do indeed evolve as the theory predicts. We indeed have the nuclear norm constrained classifier on isotropic data not exhibiting the same growth as the others, indicating that the maximal generalization gap for nuclear norm constrained classifiers is indeed sensitive to isotropy.

## E. 2 Additional details on the isotropic normalization effects on real data

Here we give additional details on the experimental setup. As mentioned in the main text, each dataset is randomly split with an 80/20 ratio to define training and testing sets. The training set is used to compute the normalizing factors $\Sigma_{n}$ and $\nu_{n}$. These factors are then used to normalize the entire dataset.

The nuclear- and Frobenius norm (or SVM) classifiers each have a hyperparameter $C$ and $\lambda$ that needs to be set before training. We search for the optimal hyperparameter among the grids [0.0000001, $0.000001,0.00001,0.0001,0.001,0.01,0.1,1,10]$ for $C$, and $[0.0001,0.001,0.01,0.1,1,10,50$, $100,500,1000,10000$ ] for $\lambda$. The best hyperparameter is determined through 4-fold cross-validation on the training set.

The SVM classifier results are computed using scikit learn's Buitinck et al. 2013 ${ }^{3}$ liblinearFan et al. $(2008)^{4}$ wrapper. Details on the nuclear norm classifier are given below.

## E. 3 Computational details on the Nuclear Norm Constrained Classifier

The classifier solves the following optimization problem

$$
\begin{array}{lr}
\min _{A \in \mathbb{R}^{d \times d}} & \frac{1}{n} \sum_{i=1}^{n} \ell\left(\mathbf{x}_{i}^{\top} A \mathbf{x}_{i}, y_{i}\right) \\
\text { subject to } & \|A\|_{n u c} \leq \lambda
\end{array}
$$

using accelerated projected gradient descent. The projection step makes use of the simplex projection algorithm proposed in Duchi et al. $\left.2008\right|^{5}$. The loss function used is the smoothed hinge loss, and its smoothness constant is upper bounded by $L=\frac{1}{n} \sum_{i=1}^{n}\left\|\mathbf{x}_{i}\right\|_{2}^{4}$. We use $1 / L$ as the stepsize for the accelerated projected gradient descent. The gradients are computed with JAX Bradbury et al. (2018). For the real dataset experiments, additional variables $b \in \mathbb{R}^{n}$ and $c \in \mathbb{R}$ are added in order to optimize over non-homogeneous quadratic polynomials like the SVM classifier.

[^1]
[^0]:    ${ }^{1}$ See remarks following the proof
    ${ }^{2}$ for simplicity we ignore the logarithmic factor

[^1]:    ${ }^{3}$ Buitinck, L., Louppe, G., Blondel, M., Pedregosa, F.,Mueller, A., Grisel, O., Niculae, V., Prettenhofer, P.,Gramfort, A., Grobler, J., Layton, R., VanderPlas, J.,Joly, A., Holt, B., and Varoquaux, G. API design for ma-chine learning software: experiences from the scikit-learnproject. InECML PKDD Workshop: Languages for DataMining and Machine Learning, pp. 108-122, 2013.
    ${ }^{4}$ Fan, R.-E., Chang, K.-W., Hsieh, C.-J., Wang, X.-R., andLin, C.-J. Liblinear: A library for large linear classifi-cation.J. Mach. Learn. Res., 9:1871-1874, June 2008.ISSN 1532-4435.
    ${ }^{5}$ Duchi, J., Shalev-Shwartz, S., Singer, Y., and Chandra, T.Efficient projections onto the 1 1-ball for learning in highdimensions. InProceedings of the 25th internationalconference on Machine learning, pp. 272-279, 2008
    ${ }^{6}$ Bradbury, J., Frostig, R., Hawkins, P., Johnson, M. J., Leary,C., Maclaurin, D., Necula, G., Paszke, A., VanderPlas, J.,Wanderman-Milne, S., and Zhang, Q. JAX: composabletransformations of Python+NumPy programs, 2018. URLhttp://github.com/google/jax.

