

Matrix-Free Two-to-Infinity and One-to-Two Norms Estimation

Askar Tsyganov¹, Evgeny Frolov^{2,1}, Sergey Samsonov^{*1}, Maxim Rakhuba^{*1}

¹HSE University, Moscow, Russia

²Personalization Technologies, Moscow, Russia
{atsyganov, e.frolov, svsamsonov, mrakhuba}@hse.ru

Abstract

In this paper, we propose new randomized algorithms for estimating the two-to-infinity and one-to-two norms in a matrix-free setting, using only matrix-vector multiplications. Our methods are based on appropriate modifications of Hutchinson’s diagonal estimator and its Hutch++ version. We provide oracle complexity bounds for both modifications. We further illustrate the practical utility of our algorithms for Jacobian-based regularization in deep neural network training on image classification tasks. We also demonstrate that our methodology can be applied to mitigate the effect of adversarial attacks in the domain of recommender systems.

Code — <https://github.com/fallnlove/TwoToInfinity>

Extended version — <https://arxiv.org/abs/2508.04444>

1 Introduction

In recent years, there has been growing interest in randomized linear algebra techniques (Halko, Martinsson, and Tropp 2011; Martinsson and Tropp 2020) for estimating matrix functions without explicit access to the matrix entries, see e.g. (Hutchinson 1989; Bekas, Kokiopoulou, and Saad 2007). This setting, known as *matrix-free*, assumes access to an oracle that computes matrix-vector products with a matrix A , and in some cases also with its transpose A^\top . The goal is to approximate important characteristics of A , for example, norms, trace or spectral density, using only these products. Such a framework is essential in modern machine learning, where matrices such as the Jacobian of deep neural networks are prohibitively large to form explicitly but allow efficient computation of matrix-vector products via automatic differentiation (autograd), see e.g. (Baydin et al. 2018).

In this paper, we focus on matrix-free estimation of two operator norms: $\|\cdot\|_{2 \rightarrow \infty}$ and $\|\cdot\|_{1 \rightarrow 2}$. Formally, for a matrix $A \in \mathbb{R}^{d \times n}$, its two-to-infinity norm can be defined as

$$\|A\|_{2 \rightarrow \infty} = \sup_{x \neq 0} \frac{\|Ax\|_\infty}{\|x\|_2} = \max_{i \in [d]} \|A_i\|_2,$$

where A_i denotes the i -th row of A . Given the identity

$$\|A\|_{2 \rightarrow \infty} = \|A^\top\|_{1 \rightarrow 2},$$

^{*}These authors contributed equally.

it suffices to concentrate on the estimation of the two-to-infinity norm. Compared to classical matrix norms such as the spectral or Frobenius norm, the two-to-infinity norm provides finer control over the row-wise structure of a matrix. Indeed, when dealing with *tall* matrices, that is, matrices $A \in \mathbb{R}^{d \times n}$ with $d \gg n$, it is natural that $\|A\|_{2 \rightarrow \infty}$ does not scale with d , contrary to spectral and Frobenius norm of this matrix. In such cases, bounding the two-to-infinity norm ensures that the norm of each row is tightly controlled. This localized control is especially useful when studying theoretical properties of various algorithms (Cape, Tang, and Priebe 2019; Pensky 2024), see details in the related work section. In this paper, we focus on methodological and statistical aspects of estimating the two-to-infinity norm, rather than on its theoretical utility. Our main contributions are as follows:

- We introduce a novel randomized algorithm tailored specifically for the $\|\cdot\|_{1 \rightarrow 2}$ and $\|\cdot\|_{2 \rightarrow \infty}$ norms estimation under the matrix-free setting. Our method enjoys provable convergence guarantees and empirically demonstrates reliable performance.
- We apply suggested randomized estimators as regularizers in image classification and recommender systems problems. For image classification with deep neural networks, our method achieves better generalization performance compared to prior Jacobian regularization techniques (Hoffman, Roberts, and Yaida 2019; Roth, Kilcher, and Hofmann 2020). In the recommender systems domain, we consider UltraGCN-type architectures (Mao et al. 2021) and show that employing our regularizer increases the robustness of the algorithm to adversarial attacks, following the pipeline of (He et al. 2018).

The rest of the paper is structured as follows. We discuss related work in Section 2. Then, in Section 3, we present our main algorithm, TwINEst (see Algorithm 1), and analyze its sample complexity. In Section 4, we provide a variance-reduced version of the TwINEst algorithm and study the theoretical properties of the modified algorithm. Finally, we present numerical results in Section 5. Proofs of the theoretical results and additional numerical experiments are provided in the supplementary material.

Notations. For a vector $x \in \mathbb{R}^d$, $\|x\|_p = (\sum_{i=1}^d |x_i|^p)^{1/p}$ denotes the ℓ_p -norm ($p \geq 1$), $\|x\|_\infty = \max_i |x_i|$ denotes

the ℓ_∞ -norm. For a matrix $A \in \mathbb{R}^{d \times n}$, A_i denotes the i -th row of A , and $\|A\|_F = (\sum_{i=1}^d \sum_{j=1}^n A_{ij}^2)^{1/2}$ denotes its Frobenius norm. We define the induced norms as $\|A\|_{p \rightarrow q} := \sup_{x \neq 0} \|Ax\|_q / \|x\|_p$. For example, $\|A\|_{2 \rightarrow \infty}$ is equal to the maximum ℓ_2 norm of the rows and $\|A\|_{1 \rightarrow 2}$ is equal to the maximum ℓ_2 norm of the columns. We define the max-norm as $\|A\|_{\max} := \min_{U, V: A=UV^T} \|U\|_{2 \rightarrow \infty} \|V\|_{2 \rightarrow \infty}$. For matrices $A, B \in \mathbb{R}^{d \times n}$, $A \odot B$ denotes the Hadamard product (element-wise). For $d \in \mathbb{N}$, we denote $[d] = \{1, \dots, d\}$. A random variable ξ is called a Rademacher random variable if $\mathbb{P}(\xi = 1) = \mathbb{P}(\xi = -1) = \frac{1}{2}$. The vector $e_i := (0, \dots, 0, 1, 0, \dots, 0)^\top$, with 1 in position i and 0 elsewhere, denotes the i -th standard basis vector in \mathbb{R}^n .

2 Related Work

A classical problem in the matrix-free setting is the stochastic trace estimation (Hutchinson 1989), in which the trace of a matrix is approximated using a few matrix-vector products with random vectors. Several improvements on this approach have been developed, including variance-reduced methods such as Hutch++ (Meyer et al. 2021), which exploit low-rank structure to accelerate convergence, and dynamic algorithms (Dharangutte and Musco 2021; Woodruff, Zhang, and Zhang 2022), which adaptively allocate samples to achieve higher accuracy. Related work also addressed the problem of estimating the spectral norm of a matrix using structured estimators based on rank-one vectors, see (Bujanovic and Kressner 2021). Another closely related problem is the estimation of the matrix diagonal (Bekas, Kokiopoulou, and Saad 2007; Baston and Nakatsukasa 2022; Dharangutte and Musco 2023), for which recent works have provided algorithmic improvements and theoretical guarantees.

The two-to-infinity norm is widely used as a theoretical tool to study statistical guarantees for algorithms in various areas of high-dimensional statistics and learning theory. Particular applications include bandit problems with specific reward structures (Jedra et al. 2024), singular subspace recovery (Cape, Tang, and Priebe 2019), and clustering (Pensky 2024).

In machine learning algorithms, the two-to-infinity norm most often appears as part of the max-norm (see the notation section), an important tool for algorithm regularization. Its applications include, in particular, the matrix completion problem via matrix factorization (Srebro, Rennie, and Jaakkola 2004). In this setting, one can show that the max-norm serves as a surrogate for the rank of the approximator matrix. This idea was further developed in (Lee et al. 2010), where the authors introduce gradient-based algorithms, such as the projected gradient method and the proximal point method, for solving the regularized matrix factorization problem. Max-norm regularizers are also an important part of many online frameworks for such problems, often combined with additional sparsity-inducing constraints, as in (Shen, Xu, and Li 2014).

Estimating matrix operator norms $\|A\|_{p \rightarrow q}$ induced by different combinations of p and q has recently attracted a lot of attention (Higham and Tisseur 2000; Bujanovic and Kressner 2021; Bresch et al. 2024; Naumov et al. 2025). A founda-

tional contribution in this direction is the work (Boyd 1974), which proposes a general iterative algorithm for estimating $\|\cdot\|_{p \rightarrow q}$ norm for arbitrary p and q , except for the cases of $\|\cdot\|_{1 \rightarrow 2}$ and $\|\cdot\|_{2 \rightarrow \infty}$ norms. Their approach is based on a generalization of the classical power method. Papers (Higham 1992; Roth, Kilcher, and Hofmann 2020) extend the methodology of (Boyd 1974) and apply it to more specialized norms, including the $\|\cdot\|_{1 \rightarrow 2}$ and $\|\cdot\|_{2 \rightarrow \infty}$ norm cases. While (Boyd 1974) proves theoretical convergence results for certain types of matrices, the authors of (Higham 1992; Roth, Kilcher, and Hofmann 2020) do not provide any theoretical guarantees for the $\|\cdot\|_{1 \rightarrow 2}$ and $\|\cdot\|_{2 \rightarrow \infty}$ cases. Furthermore, we provide an example below demonstrating that, when estimating the $\|\cdot\|_{2 \rightarrow \infty}$ norm, this method diverges with positive probability even when applied to a simple diagonal matrix.

3 Two-to-Infinity Norm Estimation

Here, we briefly introduce the adaptive power method proposed by (Higham 1992; Roth, Kilcher, and Hofmann 2020). Full pseudocode is provided in the supplementary paper. The method begins with an initial random vector X_0 (typically drawn from a Gaussian distribution), followed by an iterative update process:

$$Y^i = \text{dual}_\infty(AX^{i-1}), \quad X^i = \text{dual}_2(A^\top Y^i),$$

where the dual operator is defined as

$$\text{dual}_p(x) = \begin{cases} \text{sign}(x) \odot |x|^{p-1} / \|x\|_p^{p-1}, & \text{if } p < \infty, \\ |\mathcal{I}|^{-1} \text{sign}(x) \odot \mathbf{1}_\mathcal{I}, & \text{if } p = \infty, \end{cases}$$

where $\mathcal{I} := \{i \in [d] : |x_i| = \|x\|_\infty\}$, and $\mathbf{1}_\mathcal{I} := \sum_{i \in \mathcal{I}} e_i$ is an indicator vector over the set \mathcal{I} . The final estimate after m iterations is obtained by computing $\|AX^m\|_\infty$.

Example 1 (Divergence of the Adaptive Power Method). *Let $A \in \mathbb{R}^{2 \times 2}$ be given by*

$$A = \begin{pmatrix} 2 & 0 \\ 0 & 1 \end{pmatrix}.$$

Then, with probability at least 0.295, the Adaptive Power Method diverges when applied to the matrix A .

Explanation. Let $X = AX^0$, where $X^0 \sim \mathcal{N}(0, I_2)$ is the initial vector, and $X = (X_1, X_2)^\top$. Then:

$$\begin{aligned} \mathbb{P}(|X_1| < |X_2|) &= \mathbb{P}(2|X_1^0| < |X_2^0|) = \mathbb{P}\left(\left|\frac{X_1^0}{X_2^0}\right| < \frac{1}{2}\right) \\ &= \int_{-1/2}^{1/2} \frac{1}{\pi(1+t^2)} dt = \frac{2}{\pi} \arctan(1/2) \approx 0.2951, \end{aligned}$$

since the random variable X_1^0/X_2^0 has a Cauchy distribution. By the definition of dual_∞ , we then have $Y^1 = \text{sign}(X) \odot (0, 1)^\top$ with probability at least 0.295. Then it is easy to check that

$$X^1 = \frac{A^\top Y^1}{\|A^\top Y^1\|_2} = \text{sign}(X) \odot (0, 1)^\top.$$

It is easy to verify that $X^i = Y^i = \text{sign}(X) \odot (0, 1)^\top$ for any $i \geq 1$. Consequently, the final estimate with probability ≥ 0.295 equals

$$L = \|AX^m\|_\infty = 1,$$

which is incorrect, as $\|A\|_{2 \rightarrow \infty} = 2$. \square

3.1 Hutchinson’s Diagonal Estimator

Unlike the iterative algorithms, such as the adaptive power method, we focus on a stochastic estimator of the two-to-infinity norm, which is accurate in the limit of large number of matrix-vector products. Our algorithms build upon a well-known technique for estimating the diagonal of a square matrix using only matrix-vector products. This technique is known as the Hutchinson diagonal estimator (Bekas, Kokiopoulou, and Saad 2007; Baston and Nakatsukasa 2022; Dharangutte and Musco 2023). In this section, we introduce the Hutchinson diagonal estimator and provide concentration inequalities that are useful for our subsequent analysis.

Definition 1. Let $X^1, \dots, X^m \in \{-1, 1\}^d$ be independent Rademacher random vectors. For a square matrix $A \in \mathbb{R}^{d \times d}$, the Hutchinson’s diagonal estimator $D^m(A) \in \mathbb{R}^d$ is defined as

$$D^m(A) := \frac{1}{m} \sum_{i=1}^m X^i \odot (AX^i).$$

Notably, in this definition, instead of Rademacher vectors, it is possible to use Gaussian vectors or any mean-zero random vectors with an identity covariance matrix. However, we prefer Rademacher vectors because they yield an estimate with minimal variance (see Proposition 1 in (Hutchinson 1989)).

This estimator is a natural extension of the classical Hutchinson method for trace estimation (Hutchinson 1989). It provides an unbiased estimate of the diagonal, moreover, its variance can be computed, for any $i \in [d]$, as

$$\text{Var}[D_i^1(A)] = \sum_{j \neq i} A_{ij}^2.$$

Furthermore, the estimator enjoys high-probability error bounds that quantify its deviation from the true diagonal. In particular, we rely on the following result from (Dharangutte and Musco 2023), which provides the following bound for the ℓ_2 norm of the Hutchinson’s estimator error:

Theorem 1 (Theorem 1 in (Dharangutte and Musco 2023)). Let $A \in \mathbb{R}^{d \times d}$, $m \in \mathbb{N}$, $\delta \in (0, 1]$. Then with probability at least $1 - \delta$:

$$\|D^m(A) - \text{diag}(A)\|_2 \leq c \sqrt{\frac{\log(2/\delta)}{m}} \|A - \text{diag}(A)\|_F,$$

where c is an absolute constant.

3.2 Our method

Now we describe our strategy for estimating $\|\cdot\|_{2 \rightarrow \infty}$ norm. The main idea is that the diagonal entries of the matrix AA^\top correspond to the squared ℓ_2 norms of the rows of A . Therefore, the $\|\cdot\|_{2 \rightarrow \infty}$ norm can be equivalently expressed as

$$\|A\|_{2 \rightarrow \infty}^2 = \max_{i \in [d]} \text{diag}(AA^\top)_i.$$

This identity suggests a natural strategy: instead of computing all row norms explicitly, we can estimate the diagonal of AA^\top using the Hutchinson method, which only requires matrix-vector products with A and A^\top . The final estimate of

Algorithm 1: TwINEst: Two-to-Infinity Norm Estimation

Input:

- Matrix-vector multiplication oracle for $A \in \mathbb{R}^{d \times n}$,
- Matrix-vector multiplication oracle for $A^\top \in \mathbb{R}^{n \times d}$,
- Positive integer $m \in \mathbb{N}$: number of random samples.

Output:

- An estimate of the $\|A\|_{2 \rightarrow \infty}$ norm.
 - 1: Sample m random Rademacher vectors X^1, \dots, X^m , where each $X^i \in \{-1, 1\}^d$
 - 2: **for** each $i = 1, 2, \dots, m$ **do**
 - 3: Compute $t_i = X^i \odot AA^\top X^i$
 - 4: **end for**
 - 5: Compute $D = \frac{1}{m} \sum_{i=1}^m t_i \in \mathbb{R}^d$
 $\triangleright D$ - estimate of the AA^\top diagonal
 - 6: Find $j = \arg \max_i D_i$
 - 7: Compute $L = \|A^\top e_j\|_2$
 $\triangleright e_j$ - is the j -th standard basis vector
 - 8: **return** L
-

the $\|A\|_{2 \rightarrow \infty}$ norm is then obtained by taking the maximum of the estimated diagonal.

However, estimating the maximum value through the direct application of Hutchinson’s method introduces high variance, leading to a noisy approximation. To mitigate this, we can eliminate one source of randomness in the final estimate. Namely, let D be the estimate of the diagonal of AA^\top . While the entries of D are typically noisy, we can reduce variance by avoiding direct use of $\max_i D_i$. Instead, we first identify the (random) index of the maximum estimated value, $j = \arg \max_i D_i$, and then compute $\|A^\top e_j\|_2 = \|A_j\|_2$. This approach significantly improves the quality of the estimate, which is supported by the ablation study carried out in Section 5.1. The outlined procedure is referred to as the TwINEst algorithm and is presented in Algorithm 1. Note that the algorithm steps given in lines 2 to 4 can be done either in parallel, or in a matrix form, if it is possible to store the matrix $X = [X^1, \dots, X^m] \in \mathbb{R}^{d \times m}$ and compute $X \odot AA^\top X$. This is the main practical advantage of our method over power-iteration based algorithms (Roth, Kilcher, and Hofmann 2020).

We now establish upper bounds on the sample complexity of TwINEst. In the context of randomized numerical linear algebra, sample complexity typically refers to the number of matrix-vector multiplications required to approximate some quantity within a specified error tolerance and failure probability. Beyond these bounds, we also derive an oracle complexity: the number of matrix–vector multiplications sufficient to return the correct answer with failure probability δ , expressed in terms of the gap Δ between the largest squared ℓ_2 -norm among the rows of A and the squared ℓ_2 -norm of the nearest non-maximal row. Formally, with $M = \max_i \|A_i\|_2^2$, we define

$$\Delta = M - \max_{i: \|A_i\|_2^2 < M} \|A_i\|_2^2.$$

A large Δ makes the largest ℓ_2 -norm easy to detect, reducing oracle complexity, whereas a small Δ requires more samples, since the top norms are close.

We now establish the oracle complexity of TwINEst, namely the number of samples required to recover the exact value of the matrix norm $\|\cdot\|_{2 \rightarrow \infty}$ with high probability.

Theorem 2 (TwINEst Oracle Complexity). *Let $A \in \mathbb{R}^{d \times n}$, $m \in \mathbb{N}$. Let $T^m(A)$ be the result of Algorithm 1 based on m random vectors. Then, it suffices to take*

$$m > \frac{8 \log(2d/\delta)}{\Delta^2} \|AA^\top - \text{diag}(AA^\top)\|_{2 \rightarrow \infty}^2$$

to ensure $T^m(A) = \|A\|_{2 \rightarrow \infty}$ with probability at least $1 - \delta$.

Discussion. The proof of Theorem 2 is provided in the supplementary paper. Importantly, the structure of TwINEst allows us not only to bound the probability of deviating from the true value, but also to bound the probability that our algorithm returns the exact value. Notably, when $\Delta = 0$, identifying the maximum row norm becomes trivial, and the algorithm always returns the correct answer.

Remark 3 (TwINEst Sample Complexity). *Let $A \in \mathbb{R}^{d \times n}$, $m \in \mathbb{N}$, and $\varepsilon > 0$. Let $T^m(A)$ be the result of Algorithm 1 based on m random vectors. Then, it suffices to take*

$$m > \frac{2 \log(2d/\delta)}{\varepsilon^2} \|AA^\top - \text{diag}(AA^\top)\|_{2 \rightarrow \infty}^2$$

to ensure $|T^m(A) - \|A\|_{2 \rightarrow \infty}| < \varepsilon$ with probability at least $1 - \delta$.

Proof. Proof is completely analogous to the proof of Theorem 2. \square

Depending on Δ , either Theorem 3 or Theorem 2 may be more useful. Results similar to those in Theorem 3 have previously been obtained for the Hutchinson estimator (Roosta-Khorasani and Ascher 2015; Jiang et al. 2021), Hutch++ (Meyer et al. 2021), and several other methods. To the best of our knowledge, this is the first bound on the sample complexity of randomized estimation of the two-to-infinity norm.

4 Improved Algorithm

In this section, we present an improved version of our algorithm, which incorporates the variance reduction technique of the Hutch++ method (Meyer et al. 2021). Bearing similarity with the Hutch++ method, we call this modification TwINEst++. Rather than estimating the diagonal with the classical Hutchinson estimator, we adopt the unbiased Hutch++ diagonal modification (Han, Li, and Zhu 2024). The key idea of this technique is to approximate the dominant low-rank part of the matrix using a few random vectors, explicitly compute its diagonal, and estimate the diagonal of the residual term using Hutchinson’s estimator. More precisely, let $S \in \mathbb{R}^{d \times r}$ consist of r independent Rademacher vectors. Multiplying by AA^\top , we form the matrix $AA^\top S$, whose column span approximates the dominant eigenspace of AA^\top with high probability. To approximate this eigenspace, we compute the thin QR -decomposition of the matrix $AA^\top S$ and further rely on the factor $Q \in \mathbb{R}^{d \times r}$. Using the orthogonal projector $P = QQ^\top$, we decompose AA^\top into a low-rank approximation and a residual:

$$AA^\top = \underbrace{AA^\top P}_{\text{low-rank}} + \underbrace{AA^\top (I - P)}_{\text{residual term}}. \quad (1)$$

Algorithm 2: TwINEst++

Input:

Matrix-vector multiplication oracle for $A \in \mathbb{R}^{d \times n}$,
 Matrix-vector multiplication oracle for $A^\top \in \mathbb{R}^{n \times d}$,
 Positive integer $m \in \mathbb{N}$: total sampling budget.

Output:

An estimate of the $\|A\|_{2 \rightarrow \infty}$ norm.

- 1: Sample $\frac{m}{3}$ random Rademacher vectors $X^1, \dots, X^{\frac{m}{3}}$, where each $X^i \in \{-1, 1\}^d$
 - 2: Sample random Rademacher matrix $S \in \mathbb{R}^{d \times \frac{m}{3}}$, with i.i.d. $\{-1, 1\}$ entries
 - 3: Compute an orthonormal basis Q for $AA^\top S$
 \triangleright via QR decomposition
 - 4: **for** each $i = 1, 2, \dots, \frac{m}{3}$ **do**
 - 5: Compute $t_i = X^i \odot AA^\top (I - QQ^\top) X^i$
 - 6: **end for**
 - 7: Compute $\hat{D} = \frac{3}{m} \sum_{i=1}^{m/3} t_i \in \mathbb{R}^d$
 - 8: Compute $D = \hat{D} + \text{diag}(AA^\top QQ^\top)$
 $\triangleright D$ - estimate of the AA^\top diagonal
 - 9: Find $j = \arg \max_i D_i$
 - 10: Compute $L = \|A^\top e_j\|_2$
 $\triangleright e_j$ - is the j -th standard basis vector
 - 11: **return** L
-

The diagonal of the low-rank part can be computed explicitly:

$$\text{diag}(AA^\top P) = (A(A^\top Q) \odot Q) \cdot (1 \ 1 \ \dots \ 1)^\top,$$

this requires only $\mathcal{O}(dr)$ floating-point operations and $2r$ matrix-vector products with matrices A and A^\top . To estimate the diagonal of the matrix AA^\top , we combine this exact computation with a stochastic estimate of the residual term:

$$\text{diag}(AA^\top) \approx \text{diag}(AA^\top P) + D^m(AA^\top (I - P)).$$

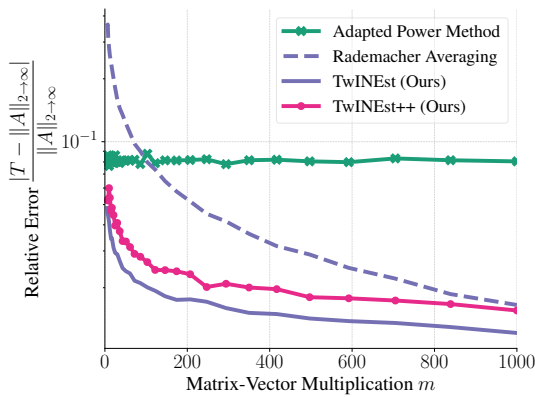
This hybrid estimator typically achieves a significantly lower variance than the standard Hutchinson method, which in turn leads to better identification of the maximum row. We follow the Hutch++ strategy and set $r = m/3$, where m is the whole budget for matrix-vector products (see ablation of different strategies for choosing r in the supplement materials). A detailed description of the algorithm is provided in Algorithm 2. Below we provide the counterpart of Theorem 2 for the TwINEst++ estimator.

Theorem 4 (TwINEst++ Oracle Complexity). *Let $A \in \mathbb{R}^{d \times n}$, $m \in \mathbb{N}$ and $c \in \mathbb{R}_+$ - some constant. Let $T_{++}^m(A)$ be the output of Algorithm 2 based on m samples. Then, it suffices to choose*


$$m > c \cdot \left(\frac{\sqrt{\log(2/\delta)}}{\Delta} \|A\|_F^2 + \log(1/\delta) \right)$$

to ensure that $T_{++}^m(A) = \|A\|_{2 \rightarrow \infty}$ with probability at least $1 - \delta$.

Proof. The structure of the proof follows the same reasoning as in the case of Theorem 2 for the base TwINEst algorithm. In particular, we again rely on a concentration inequality for



(a) Synthetic data with $\Delta = 10^{-2}$.



(b) Synthetic data with $\Delta = 10^{-1}$.

Figure 1: Comparison of methods for estimating the two-to-infinity norm on random square matrices. Shown is the relative error versus the number of matrix-vector multiplications, averaged over 500 trials.

the diagonal estimator. However, for TwINEst++, we employ a refined concentration result provided in Theorem 1, which yields a tighter control over the estimation error and enables the improved complexity result stated below. Full proof is provided in the supplement paper. \square

Theorem 4 shows that TwINEst++ achieves an improved oracle complexity compared to the original algorithm, particularly in challenging scenarios when $\Delta \rightarrow 0$, making identification of the correct row difficult. Specifically, the oracle complexity is reduced from $O(1/\Delta^2)$ in the original TwINEst algorithm to $O(1/\Delta)$ in TwINEst++. Moreover, we expect that TwINEst++ performs particularly well in low-rank settings, exploiting the low-rank structure of the matrices, as in recent methods (Meyer et al. 2021; Parkina and Rakhuba 2025). However, TwINEst’s oracle complexity depends on the norm of the matrix without its diagonal, while for TwINEst++ it depends on the norm of the whole matrix, so TwINEst may perform better for some matrices.

5 Experiments

In this section, we present an empirical evaluation of the proposed algorithms. We focus on comparing the algorithms on synthetic and real-world matrices (see Section 5.1 and Section 5.2), and applications of our methods to image classification and recommender systems tasks (see Section 5.3 and Section 5.4). All experiments are conducted on a single NVIDIA Tesla V100 GPU with 32GB memory.

5.1 Synthetic Data

In the following two sections, we compare the following methods:

- **Adaptive Power Method.** A modification of the power iteration method for estimating the two-to-infinity norm (see Algorithm 3 in the supplementary paper), based on (Higham 1992; Roth, Kilcher, and Hofmann 2020).
- **Rademacher Averaging.** A version of the TwINEst method that does not compute the exact norm of the candidate row with the maximum norm. Formally, for a matrix

$A \in \mathbb{R}^{d \times n}$ we output $\sqrt{\max_i D_i^m (AA^\top)}$ (see notations in Algorithm 1).

- **TwINEst.** Our algorithm introduced in Algorithm 1.
- **TwINEst++.** A variance-reduced version of TwINEst, described in Algorithm 2.

The Adaptive Power Method lacks theoretical guarantees and may diverge on certain matrices (see Example 1). Therefore, we hypothesize that our algorithms will outperform it. Experiments with Rademacher Averaging serve as a natural ablation study for the TwINEst method.

For our synthetic experiments, we generate random Gaussian matrices $A \in \mathbb{R}^{5000 \times 5000}$. Specifically, we fix a parameter $\Delta \in (0, 1)$ and sample values $c_3, \dots, c_{5000} \sim \mathcal{U}[0, 1]$, setting $c_1 = 1 + \Delta$, $c_2 = 1$. Each row of the matrix A is normalized such that its squared ℓ_2 -norm is equal to c_i , ensuring a gap of magnitude Δ between the largest and second-largest squared row norms. The singular value distributions of these matrices are quite similar for different values of Δ (see Figure 4a in the supplementary paper). Therefore, the parameter Δ may have a significant impact on the convergence rate.

The results of synthetic experiments for different values of Δ are illustrated in Figure 1. As anticipated, TwINEst consistently outperforms Rademacher Averaging. The Adaptive Power Method fails to converge, as evidenced by its flat performance line. For a relatively large gap $\Delta = 10^{-1}$, the TwINEst algorithm rapidly converges, achieving accurate results consistently within approximately 400 iterations.

5.2 WideResNet Jacobian

To validate our methods on real-world data, we evaluate them on the Jacobian matrix $J \in \mathbb{R}^{3 \cdot 32 \cdot 32 \times 100}$ of a WideResNet-16-10 (Zagoruyko and Komodakis 2016) trained on CIFAR-100 (Krizhevsky, Hinton et al. 2009). Given the structure of J (the number of columns is much smaller than the number of rows), we expect that the $AA^\top P$ term (see Equation (1)) in TwINEst++ will capture almost the entire span of AA^\top with high probability. Therefore, TwINEst++ should outperform the basic TwINEst method on this matrix.

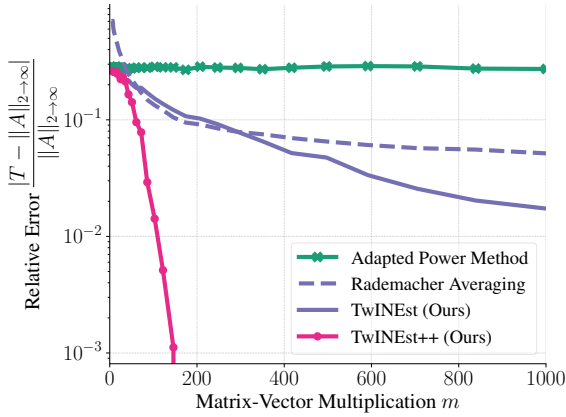


Figure 2: Comparison of methods for estimating the two-to-infinity norm of the Jacobian matrix of WideResNet-16-10 trained on CIFAR-100. The plot shows the relative error versus the number of matrix-vector multiplications, averaged over 500 trials.

Figure 2 confirms our hypothesis: the TwINEst++ algorithm achieves rapid convergence consistent with the rank of matrix J , significantly outperforming other algorithms. Again, the Adaptive Power Method fails to converge. For a comprehensive ablation, we provide a comparison between the relative error of the methods and their floating point operation counts (FLOPs) in Figure 6.

5.3 Deep Learning Application

We study whether penalizing the two-to-infinity norm of the input-output Jacobian can improve the generalization ability and adversarial robustness of neural networks in image classification. Theorem 1 in (Roth, Kilcher, and Hofmann 2020) establishes a formal connection between $\ell_{p \rightarrow q}$ Jacobian norm regularization and adversarial training. This motivates our hypothesis that two-to-infinity regularization can enhance adversarial robustness of deep neural networks.

We compare our regularizer to the standard Jacobian-based penalties: Frobenius norm (Hoffman, Roberts, and Yaida 2019), spectral norm (Roth, Kilcher, and Hofmann 2020), and infinity norm (ℓ_∞) (Roth, Kilcher, and Hofmann 2020). For most image-classification datasets (e.g., CIFAR-100, TinyImageNet (Le and Yang 2015)), the number of output classes is much smaller than the number of input features (pixels). Consequently, the Jacobian of an image classifier trained on these datasets is a tall matrix. Hence, the two-to-infinity norm’s finer control over the Jacobian’s elements leads us to expect our regularizer to outperform standard norm penalties such as the Frobenius and spectral norms.

Formally, we minimize the following objective function:

$$\mathcal{L}(x, y) = \mathcal{L}_{\text{CE}}(f(x), y) + \lambda \cdot \|J_f(x)\|^2,$$

where \mathcal{L}_{CE} is the cross-entropy loss, $x \in \mathbb{R}^d$, $f(x) \in \mathbb{R}^M$ is the output of a deep neural network, $M \in \mathbb{N}$ is the number of classes, $J_f(x) \in \mathbb{R}^{d \times M}$ is the Jacobian of the logits $f(x)$ with respect to the input x , and $\|\cdot\|$ denotes one of the following norms: Frobenius, spectral, infinity, or two-to-infinity.

Explicit construction of the Jacobian of a deep network is computationally prohibitive in terms of both time and memory. Instead, we approximate each norm using the estimators described in the cited papers (see details in Appendix D.5). For the two-to-infinity norm, we employ TwINEst algorithm. All of these estimators require only Jacobian–vector and vector–Jacobian products, which are supported by the automatic-differentiation frameworks. However, such operations cost roughly as much as a single backward pass, and naively applying the regularizer at every training step can slow down learning. To provide a more comprehensive ablation, we show that updating the regularization term once every k iterations still outperforms other methods, while adding negligible wall-clock overhead relative to training without regularization. For more details, see Appendix D.6 in the supplementary paper.

Experiments are conducted on CIFAR-100 and TinyImageNet datasets using the WideResNet-16-10 architecture, implemented in PyTorch. The hyperparameters are provided in Appendix D.4. We evaluate each method by reporting the final test accuracy, the stable rank of the Jacobian (computed as $\|J\|_F^2 / \|J\|_2^2$), and adversarial metrics: accuracy after FGSM (Goodfellow, Shlens, and Szegedy 2015) and PGD (Madry et al. 2018) attacks with 2 steps.

As shown in Table 1, our method enhances both the generalization performance and the adversarial robustness of WideResNet-16-10 on the CIFAR-100 and TinyImageNet datasets. Other approaches show limited improvements over the baseline.

5.4 RecSys Application

In this section, we study the application of our two-to-infinity norm regularization to improve the adversarial robustness of recommender systems. We focus on the collaborative filtering task with d_u users and d_i items, and use the classic matrix factorization framework:

$$R \approx UV^\top, \quad (2)$$

where $R \in \mathbb{R}^{d_u \times d_i}$ is the user-item interaction matrix, and $U \in \mathbb{R}^{d_u \times r}$, $V \in \mathbb{R}^{d_i \times r}$ are users and items embedding matrices, respectively. The predicted interaction scores form the matrix $\hat{R} = UV^\top \in \mathbb{R}^{d_u \times d_i}$, where $\hat{R}_{u,i}$ denotes the predicted score for the interaction between the u -th user and the i -th item.

Some collaborative filtering models (e.g., (Rendle et al. 2009; Mao et al. 2021)) can be expressed within the framework described in Equation (2). We follow the adversarial setting introduced in (He et al. 2018), and adapt it to the UltraGCN model (Mao et al. 2021). The UltraGCN loss function is defined as:

$$\mathcal{L}(U, V) = \mathcal{L}_{\text{BCE}}(U, V) + \gamma \cdot \mathcal{L}_{\text{C}}(U, V) + \lambda \cdot (\|U\|_F^2 + \|V\|_F^2),$$

where \mathcal{L}_{BCE} is the binary cross-entropy loss, and \mathcal{L}_{C} is a constraint loss defined in (Mao et al. 2021).

In the adversarial setting, the goal is to find a small perturbation of the model parameters that significantly degrades prediction quality. Since only the binary cross-entropy term directly affects recommendation accuracy, we aim to find perturbations that maximize this component:

$$\Delta_U^*, \Delta_V^* \in \arg \max_{\|\Delta_U\|_2, \|\Delta_V\|_2 \leq \epsilon} \mathcal{L}_{\text{BCE}}(U + \Delta_U, V + \Delta_V),$$

Regularizer	CIFAR-100				TinyImageNet			
	Acc. \uparrow	FGSM \uparrow	PGD \uparrow	S. Rank \downarrow	Acc. \uparrow	FGSM \uparrow	PGD \uparrow	S. Rank \downarrow
No regularization	75.5 \pm 0.2	24.4 \pm 0.6	11.7 \pm 0.3	32.0 \pm 1.1	57.8 \pm 1.3	30.4 \pm 0.3	20.2 \pm 0.1	30.9 \pm 4.3
Frobenius	75.7 \pm 0.5	23.5 \pm 0.2	13.3 \pm 0.2	31.6 \pm 0.2	58.6 \pm 0.3	31.1 \pm 0.2	20.7 \pm 0.5	27.8 \pm 0.9
Spectral	75.7 \pm 0.3	23.3 \pm 0.7	11.3 \pm 0.4	32.0 \pm 1.0	57.4 \pm 0.8	30.0 \pm 1.1	20.0 \pm 0.5	28.2 \pm 0.3
Infinity	75.8 \pm 0.4	23.7 \pm 0.7	11.1 \pm 0.2	30.7 \pm 1.2	57.1 \pm 0.7	29.6 \pm 1.4	19.7 \pm 1.0	28.8 \pm 0.9
Two-to-Infinity (ours)	77.3 \pm 0.1	26.9 \pm 0.5	14.5 \pm 0.5	18.3 \pm 0.8	59.6 \pm 0.9	<u>31.0</u> \pm 1.0	23.4 \pm 0.7	24.9 \pm 0.3

Table 1: Comparison of Jacobian regularization methods on CIFAR-100 and TinyImageNet datasets using WideResNet-16-10. Metrics are averaged over 3 trials. Up arrow (\uparrow) indicates higher is better, while down arrow (\downarrow) indicates lower is better.




Figure 3: Comparison of different regularization methods for adversarial robustness of UltraGCN (Mao et al. 2021). Shown is the NDCG@10 metric (higher is better) versus the magnitude of the attack. Metric is averaged over 5 trials.

where ϵ controls the attack magnitude. A practical way to obtain such perturbations is via:

$$\Delta_U^{\text{adv}} = \epsilon \cdot \frac{\Gamma_U}{\|\Gamma_U\|_2}, \quad \text{and} \quad \Delta_V^{\text{adv}} = \epsilon \cdot \frac{\Gamma_V}{\|\Gamma_V\|_2},$$

where Γ_U and Γ_V are the gradients of \mathcal{L}_{BCE} with respect to U and V , respectively. To improve the adversarial robustness of the UltraGCN model, we modify its loss function by replacing the standard weight decay term ($\|U\|_F^2 + \|V\|_F^2$) with a regularization term based on the two-to-infinity norm of the score matrix. This regularizer provides better control over the embedding matrices by taking into account the impact of every user embedding and penalizing the largest ones. Intuitively, this should prevent the model from overfitting to a small subset of popular users. Specifically, our loss function becomes:

$$\mathcal{L}_{\text{ours}}(U, V) = \mathcal{L}_{\text{BCE}}(U, V) + \gamma \cdot \mathcal{L}_{\text{C}}(U, V) + \lambda \cdot \|\hat{R}\|_{2 \rightarrow \infty}^2.$$

Note that computing the full score matrix \hat{R} at every training iteration is computationally prohibitive due to its large size and memory requirements. To address this, we approximate the two-to-infinity norm using our TwINEst algorithm. We compare our two-to-infinity regularizer with the standard weight decay used in the UltraGCN model, as well as with the two-to-infinity norm of the factors ($\|U\|_{2 \rightarrow \infty}^2 + \|V\|_{2 \rightarrow \infty}^2$), which is closely related to the max-norm regularizer popular in recommender systems (Srebro, Rennie, and Jaakkola 2004).

Experiments are conducted on the MovieLens-1M (Harper and Konstan 2015), Yelp-2018 (Asghar 2016), and CiteULike

(Wang, Chen, and Li 2013) datasets, using the UltraGCN architecture implemented in PyTorch. The hyperparameter configurations and the dataset preprocessing pipeline are provided in Appendix D.7. We evaluate methods by reporting the NDCG@10 metric (which rewards ranking relevant items near the top) (Järvelin and Kekäläinen 2002). Additional results using other metrics are provided in Appendix D.8.

Results are shown in Figure 3. On every dataset, our regularization demonstrates improved robustness under adversarial attacks of moderate to high magnitude. While our method achieves comparable metrics to the baselines under no attack, it exhibits greater stability under strong perturbations.

6 Conclusion

In this paper, we propose two novel matrix-free stochastic algorithms for estimating the two-to-infinity and one-to-two norms, and provide a theoretical analysis of their behavior. Our empirical results demonstrate that the proposed methods outperform existing approaches in terms of both accuracy and computational efficiency. Furthermore, we show that our algorithms can be easily integrated into deep learning pipelines and applied to improve robustness to adversarial attacks in recommender systems. Possible directions for future research include establishing lower bounds on the sample and oracle complexity of the two-to-infinity norm estimation or exploring additional applications of two-to-infinity norm estimation in the matrix-free setting.

Acknowledgements

The work was supported by the grant for research centers in the field of AI provided by the Ministry of Economic Development of the Russian Federation in accordance with the agreement 000000C313925P4E0002 and the agreement with HSE University № 139-15-2025-009. This research was supported in part through computational resources of HPC facilities at HSE University (Kostenetskiy, Chulkevich, and Kozyrev 2021).

References

- Asghar, N. 2016. Yelp dataset challenge: Review rating prediction. *arXiv preprint arXiv:1605.05362*.
- Baston, R. A.; and Nakatsukasa, Y. 2022. Stochastic diagonal estimation: probabilistic bounds and an improved algorithm. *arXiv preprint arXiv:2201.10684*.
- Baydin, A. G.; Pearlmutter, B. A.; Radul, A. A.; and Siskind, J. M. 2018. Automatic differentiation in machine learning: a survey. *Journal of machine learning research*, 18(153): 1–43.
- Bekas, C.; Kokiopoulou, E.; and Saad, Y. 2007. An estimator for the diagonal of a matrix. *Applied numerical mathematics*, 57(11-12): 1214–1229.
- Boyd, D. W. 1974. The power method for ℓ_p norms. *Linear Algebra and its Applications*, 9: 95–101.
- Bresch, J.; Lorenz, D. A.; Schneppe, F.; and Winkler, M. 2024. Matrix-free stochastic calculation of operator norms without using adjoints. *arXiv preprint arXiv:2410.08297*.
- Bujanovic, Z.; and Kressner, D. 2021. Norm and trace estimation with random rank-one vectors. *SIAM Journal on Matrix Analysis and Applications*, 42(1): 202–223.
- Cape, J.; Tang, M.; and Priebe, C. E. 2019. The two-to-infinity norm and singular subspace geometry with applications to high-dimensional statistics. *The Annals of Statistics*, 47(5): 2405–2439.
- Dharangutte, P.; and Musco, C. 2021. Dynamic trace estimation. *Advances in Neural Information Processing Systems*, 34: 30088–30099.
- Dharangutte, P.; and Musco, C. 2023. A tight analysis of hutchinson’s diagonal estimator. In *Symposium on Simplicity in Algorithms (SOSA)*, 353–364. SIAM.
- Goodfellow, I. J.; Shlens, J.; and Szegedy, C. 2015. Explaining and Harnessing Adversarial Examples. In Bengio, Y.; and LeCun, Y., eds., *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*.
- Halko, N.; Martinsson, P.-G.; and Tropp, J. A. 2011. Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions. *SIAM review*, 53(2): 217–288.
- Han, Z.; Li, W.; and Zhu, S. 2024. Stochastic diagonal estimation with adaptive parameter selection. *arXiv preprint arXiv:2410.11613*.
- Harper, F. M.; and Konstan, J. A. 2015. The movielens datasets: History and context. *Acm transactions on interactive intelligent systems (tiis)*, 5(4): 1–19.
- He, X.; He, Z.; Du, X.; and Chua, T.-S. 2018. Adversarial personalized ranking for recommendation. In *The 41st International ACM SIGIR conference on research & development in information retrieval*, 355–364.
- Higham, N. J. 1992. Estimating the matrix p-norm. *Numerische Mathematik*, 62(1): 539–555.
- Higham, N. J.; and Tisseur, F. 2000. A block algorithm for matrix 1-norm estimation, with an application to 1-norm pseudospectra. *SIAM Journal on Matrix Analysis and Applications*, 21(4): 1185–1201.
- Hoffman, J.; Roberts, D. A.; and Yaida, S. 2019. Robust learning with jacobian regularization. *arXiv preprint arXiv:1908.02729*.
- Hutchinson, M. F. 1989. A stochastic estimator of the trace of the influence matrix for Laplacian smoothing splines. *Communications in Statistics-Simulation and Computation*, 18(3): 1059–1076.
- Järvelin, K.; and Kekäläinen, J. 2002. Cumulated gain-based evaluation of IR techniques. *ACM Transactions on Information Systems (TOIS)*, 20(4): 422–446.
- Jedra, Y.; Réveillard, W.; Stojanovic, S.; and Proutiere, A. 2024. Low-Rank Bandits via Tight Two-to-Infinity Singular Subspace Recovery. In *International Conference on Machine Learning*, 21430–21485. PMLR.
- Jiang, S.; Pham, H.; Woodruff, D.; and Zhang, R. 2021. Optimal sketching for trace estimation. *Advances in Neural Information Processing Systems*, 34: 23741–23753.
- Kostenetskiy, P.; Chulkevich, R.; and Kozyrev, V. 2021. HPC resources of the higher school of economics. In *Journal of Physics: Conference Series*, volume 1740, 012050. IOP Publishing.
- Krizhevsky, A.; Hinton, G.; et al. 2009. Learning multiple layers of features from tiny images.
- Le, Y.; and Yang, X. 2015. Tiny imagenet visual recognition challenge. *CS 231N*, 7(7): 3.
- Lee, J. D.; Recht, B.; Srebro, N.; Tropp, J.; and Salakhutdinov, R. R. 2010. Practical Large-Scale Optimization for Max-norm Regularization. In Lafferty, J.; Williams, C.; Shawe-Taylor, J.; Zemel, R.; and Culotta, A., eds., *Advances in Neural Information Processing Systems*, volume 23. Curran Associates, Inc.
- Madry, A.; Makelov, A.; Schmidt, L.; Tsipras, D.; and Vladu, A. 2018. Towards Deep Learning Models Resistant to Adversarial Attacks. In *International Conference on Learning Representations*.
- Mao, K.; Zhu, J.; Xiao, X.; Lu, B.; Wang, Z.; and He, X. 2021. UltraGCN: ultra simplification of graph convolutional networks for recommendation. In *Proceedings of the 30th ACM international conference on information & knowledge management*, 1253–1262.
- Martinsson, P.-G.; and Tropp, J. A. 2020. Randomized numerical linear algebra: Foundations and algorithms. *Acta Numerica*, 29: 403–572.
- Meyer, R. A.; Musco, C.; Musco, C.; and Woodruff, D. P. 2021. Hutch++: Optimal stochastic trace estimation. In

Symposium on Simplicity in Algorithms (SOSA), 142–155. SIAM.

Naumov, A.; Rakhuba, M.; Ryapolov, D.; and Samsonov, S. 2025. On the Upper Bounds for the Matrix Spectral Norm. *arXiv preprint arXiv:2506.15660*.

Parkina, U.; and Rakhuba, M. 2025. COALA: Numerically Stable and Efficient Framework for Context-Aware Low-Rank Approximation. In *Advances in Neural Information Processing Systems*.

Pensky, M. 2024. Davis-Kahan Theorem in the two-to-infinity norm and its application to perfect clustering. *arXiv preprint arXiv:2411.11728*.

Rendle, S.; Freudenthaler, C.; Gantner, Z.; and Schmidt-Thieme, L. 2009. BPR: Bayesian personalized ranking from implicit feedback. In *Proceedings of the Twenty-Fifth Conference on Uncertainty in Artificial Intelligence, UAI '09*, 452–461. Arlington, Virginia, USA: AUAI Press. ISBN 9780974903958.

Roosta-Khorasani, F.; and Ascher, U. 2015. Improved bounds on sample size for implicit matrix trace estimators. *Foundations of Computational Mathematics*, 15(5): 1187–1212.

Roth, K.; Kilcher, Y.; and Hofmann, T. 2020. Adversarial training is a form of data-dependent operator norm regularization. *Advances in Neural Information Processing Systems*, 33: 14973–14985.

Shen, J.; Xu, H.; and Li, P. 2014. Online Optimization for Max-Norm Regularization. In Ghahramani, Z.; Welling, M.; Cortes, C.; Lawrence, N.; and Weinberger, K., eds., *Advances in Neural Information Processing Systems*, volume 27. Curran Associates, Inc.

Srebro, N.; Rennie, J.; and Jaakkola, T. 2004. Maximum-Margin Matrix Factorization. In Saul, L.; Weiss, Y.; and Bottou, L., eds., *Advances in Neural Information Processing Systems*, volume 17. MIT Press.

Wang, H.; Chen, B.; and Li, W.-J. 2013. Collaborative topic regression with social regularization for tag recommendation. In *IJCAI*, volume 13, 2719–2725.

Woodruff, D.; Zhang, F.; and Zhang, R. 2022. Optimal query complexities for dynamic trace estimation. *Advances in Neural Information Processing Systems*, 35: 35049–35060.

Zagoruyko, S.; and Komodakis, N. 2016. Wide Residual Networks. In Richard C. Wilson, E. R. H.; and Smith, W. A. P., eds., *Proceedings of the British Machine Vision Conference (BMVC)*, 87.1–87.12. BMVA Press. ISBN 1-901725-59-6.