

---

# Fitting Spectral Decay with the $k$ -Support Norm

---

**Andrew M. McDonald**  
University College London

**Massimiliano Pontil**  
Italian Institute of Technology  
University College London

**Dimitris Stamos**  
University College London

## Abstract

The spectral  $k$ -support norm enjoys good estimation properties in low rank matrix learning problems, empirically outperforming the trace norm. Its unit ball is the convex hull of rank  $k$  matrices with unit Frobenius norm. In this paper we generalize the norm to the spectral  $(k, p)$ -support norm, whose additional parameter  $p$  can be used to tailor the norm to the decay of the spectrum of the underlying model. We characterize the unit ball and we explicitly compute the norm. We further provide a conditional gradient method to solve regularization problems with the norm, and we derive an efficient algorithm to compute the Euclidean projection on the unit ball in the case  $p = \infty$ . In numerical experiments, we show that allowing  $p$  to vary significantly improves performance over the spectral  $k$ -support norm on various matrix completion benchmarks, and better captures the spectral decay of the underlying model.

## 1 Introduction

The problem of learning a sparse vector or a low rank matrix has generated much interest in recent years. A popular approach is to use convex regularizers which encourage sparsity, and a number of these have been studied with applications including image denoising, collaborative filtering and multitask learning, see for example, Buehlmann and van der Geer [2011], Wainwright [2014] and references therein.

Recently, the  $k$ -support norm was proposed by Argyriou et al. [2012], motivated as a tight relaxation of the set of  $k$ -sparse vectors of unit Euclidean norm. The

authors argue that as a regularizer for sparse vector estimation, the norm empirically outperforms the Lasso [Tibshirani, 1996] and Elastic Net [Zou and Hastie, 2005] penalties. Statistical bounds on the Gaussian width of the  $k$ -support norm have been provided by Chatterjee et al. [2014]. The  $k$ -support norm has also been extended to the matrix setting. By applying the norm to the vector of singular values of a matrix, McDonald et al. [2014] obtain the orthogonally invariant *spectral  $k$ -support norm*, reporting state of the art performance on matrix completion benchmarks.

Motivated by the performance of the  $k$ -support norm in sparse vector and matrix learning problems, in this paper we study a natural generalization by considering the  $\ell_p$ -norms (for  $p \in [1, \infty]$ ) in place of the Euclidean norm. These allow a further degree of freedom when fitting a model to the underlying data. We denote the ensuing norm the  $(k, p)$ -support norm. As we demonstrate in numerical experiments, using  $p = 2$  is not necessarily the best choice in all instances. By tuning the value of  $p$  the model can incorporate prior information regarding the singular values. When prior knowledge is lacking, the parameter can be chosen by validation, hence the model can adapt to a variety of decay patterns of the singular values. An interesting property of the norm is that it interpolates between the  $\ell_1$  norm (for  $k = 1$ ) and the  $\ell_p$ -norm (for  $k = d$ ). It follows that varying both  $k$  and  $p$  the norm allows one to learn sparse vectors which exhibit different patterns of decay in the non-zero elements.

A main goal of the paper is to study the proposed norm in matrix learning problems. The  $(k, p)$ -support norm is a symmetric gauge function hence it induces the orthogonally invariant *spectral  $(k, p)$ -support norm*. It interpolates between the trace norm (for  $k = 1$ ) and the Schatten  $p$ -norms (for  $k = d$ ) and its unit ball has a simple geometric interpretation as the convex hull of matrices of rank no greater than  $k$  and Schatten  $p$ -norm no greater than one. This suggests that the new norm favors low rank structure and the effect of varying  $p$  allows different patterns of decay in the spectrum. In the special case of  $p = \infty$ , the  $(k, p)$ -support

---

Appearing in Proceedings of the 19<sup>th</sup> International Conference on Artificial Intelligence and Statistics (AISTATS) 2016, Cadiz, Spain. JMLR: W&CP volume 51. Copyright 2016 by the authors.

norm is the dual of the Ky-Fan  $k$ -norm [Bhatia, 1997] and it encourages a flat spectrum when used as a regularizer.

In summary, the main contributions of the paper are: i) we propose the  $(k, p)$ -support norm as an extension of the  $k$ -support norm and we characterize in particular the unit ball of the induced orthogonally invariant matrix norm (Section 3); ii) we show that the norm can be computed efficiently and we discuss the role of the parameter  $p$  (Section 4); iii) we outline a conditional gradient method to solve the associated regularization problem for both vector and matrix problems (Section 5); and in the special case  $p = \infty$  we provide an  $\mathcal{O}(d \log d)$  computation of the projection operator (Section 5.1); finally, iv) we present numerical experiments on matrix completion benchmarks which demonstrate that the proposed norm offers significant improvement over previous methods, and we discuss the effect of the parameter  $p$  (Section 6). The supplementary material contains derivations of results which are sketched in or are omitted from the main body of the paper.

**Notation.** We use  $\mathbb{N}_n$  for the set of integers from 1 up to and including  $n$ . We let  $\mathbb{R}^d$  be the  $d$ -dimensional real vector space, whose elements are denoted by lower case letters. For any vector  $w \in \mathbb{R}^d$ , its *support* is defined as  $\text{supp}(w) = \{i \in \mathbb{R}^d : w_i \neq 0\}$ , and its *cardinality* is defined as  $\text{card}(w) = |\text{supp}(w)|$ . We let  $\mathbb{R}^{d \times m}$  be the space of  $d \times m$  real matrices. We denote the rank of a matrix as  $\text{rank}(W)$ . We let  $\sigma(W) \in \mathbb{R}^r$  be the vector formed by the singular values of  $W$ , where  $r = \min(d, m)$ , and where we assume that the singular values are ordered nonincreasing, that is  $\sigma_1(W) \geq \dots \geq \sigma_r(W) \geq 0$ . For  $p \in [1, \infty)$  the  $\ell_p$ -norm of a vector  $w \in \mathbb{R}^d$  is defined as  $\|w\|_p = (\sum_{i=1}^d |w_i|^p)^{1/p}$  and  $\|w\|_\infty = \max_{i=1}^d |w_i|$ . Given a norm  $\|\cdot\|$  on  $\mathbb{R}^d$  or  $\mathbb{R}^{d \times m}$ ,  $\|\cdot\|_*$  denotes the corresponding dual norm, defined by  $\|u\|_* = \sup\{\langle u, w \rangle : \|w\| \leq 1\}$ . The convex hull of a subset  $S$  of a vector space is denoted  $\text{co}(S)$ .

## 2 Background and Previous Work

For every  $k \in \mathbb{N}_d$ , the  $k$ -support norm  $\|\cdot\|_{(k)}$  is defined as the norm whose unit ball is given by

$$\text{co}\{w \in \mathbb{R}^d : \text{card}(w) \leq k, \|w\|_2 \leq 1\}, \quad (1)$$

that is, the convex hull of the set of vectors of cardinality at most  $k$  and  $\ell_2$ -norm no greater than one [Argyriou et al., 2012]. We readily see that for  $k = 1$  and  $k = d$  we recover the unit ball of the  $\ell_1$  and  $\ell_2$ -norms respectively.

The  $k$ -support norm of a vector  $w \in \mathbb{R}^d$  can be expressed as an infimal convolution [Rockafellar, 1970,

p. 34],

$$\|w\|_{(k)} = \inf_{(v_g)} \left\{ \sum_{g \in \mathcal{G}_k} \|v_g\|_2 : \sum_{g \in \mathcal{G}_k} v_g = w \right\}, \quad (2)$$

where  $\mathcal{G}_k$  is the collection of all subsets of  $\mathbb{N}_d$  containing at most  $k$  elements and the infimum is over all vectors  $v_g \in \mathbb{R}^d$  such that  $\text{supp}(v_g) \subseteq g$ , for  $g \in \mathcal{G}_k$ . Equation (2) highlights that the  $k$ -support norm is a special case of the group lasso with overlap [Jacob et al., 2009], where the cardinality of the support sets is at most  $k$ . This expression suggests that when used as a regularizer, the norm encourages vectors  $w$  to be a sum of a limited number of vectors with small support. Due to the variational form of (2) computing the norm is not straightforward, however Argyriou et al. [2012] note that the dual norm has a simple form, namely it is the  $\ell_2$ -norm of the  $k$  largest components,

$$\|u\|_{(k),*} = \sqrt{\sum_{i=1}^k (|u|_i^\downarrow)^2}, \quad u \in \mathbb{R}^d, \quad (3)$$

where  $|u|^\downarrow$  is the vector obtained from  $u$  by reordering its components so that they are nonincreasing in absolute value. Note also from equation (3) that for  $k = 1$  and  $k = d$ , the dual norm is equal to the  $\ell_\infty$ -norm and  $\ell_2$ -norm, respectively, which agrees with our earlier observation regarding the primal norm.

A related problem which has been studied in recent years is learning a matrix from a set of linear measurements, in which the underlying matrix is assumed to have sparse spectrum (low rank). The trace norm, the  $\ell_1$ -norm of the singular values of a matrix, has been shown to perform well in this setting, see e.g. Argyriou et al. [2008], Jaggi and Sulovsky [2010]. Recall that a norm  $\|\cdot\|$  on  $\mathbb{R}^{d \times m}$  is called orthogonally invariant if  $\|W\| = \|UWV\|$ , for any orthogonal matrices  $U \in \mathbb{R}^{d \times d}$  and  $V \in \mathbb{R}^{m \times m}$ . A classical result by von Neumann establishes that a norm is orthogonally invariant if and only if it is of the form  $\|W\| = g(\sigma(W))$ , where  $\sigma(W)$  is the vector formed by the singular values of  $W$  in nonincreasing order, and  $g$  is a symmetric gauge function [Von Neumann, 1937]. In other words,  $g$  is a norm which is invariant under permutations and sign changes of the vector components, that is  $g(w) = g(Pw) = g(Jw)$ , where  $P$  is any permutation matrix and  $J$  is diagonal with entries equal to  $\pm 1$  [Horn and Johnson, 1991, p. 438].

Examples of symmetric gauge functions are the  $\ell_p$  norms for  $p \in [1, \infty]$  and the corresponding orthogonally invariant norms are called the Schatten  $p$ -norms [Horn and Johnson, 1991, p. 441]. In particular, those include the trace norm and Frobenius norm for  $p = 1$  and  $p = 2$  respectively. Regularization with Schatten

$p$ -norms has been previously studied by Argyriou et al. [2007] and a statistical analysis has been performed by Rohde and Tsybakov [2011]. As the set  $\mathcal{G}_k$  includes all subsets of size  $k$ , expression (2) for the  $k$ -support norm reveals that  $\|\cdot\|_{(k)}$  is a symmetric gauge function. McDonald et al. [2014] use this fact to introduce the spectral  $k$ -support norm for matrices, by defining  $\|W\|_{(k)} = \|\sigma(W)\|_{(k)}$ , for  $W \in \mathbb{R}^{d \times m}$  and report state of the art performance on matrix completion benchmarks.

### 3 The $(k, p)$ -Support Norm

In this section we introduce the  $(k, p)$ -support norm as a natural extension of the  $k$ -support norm. This follows by applying the  $\ell_p$ -norm, rather than the Euclidean norm, in the infimum convolution definition of the norm.

**Definition 1.** Let  $k \in \mathbb{N}_d$  and  $p \in [1, \infty]$ . The  $(k, p)$ -support norm of a vector  $w \in \mathbb{R}^d$  is defined as

$$\|w\|_{(k,p)} = \inf_{(v_g)} \left\{ \sum_{g \in \mathcal{G}_k} \|v_g\|_p : \sum_{g \in \mathcal{G}_k} v_g = w \right\}. \quad (4)$$

where the infimum is over all vectors  $v_g \in \mathbb{R}^d$  such that  $\text{supp}(v_g) \subseteq g$ , for  $g \in \mathcal{G}_k$ .

Let us note that the norm is well defined. Indeed, positivity, homogeneity and non degeneracy are immediate. To prove the triangle inequality, let  $w, w' \in \mathbb{R}^d$ . For any  $\epsilon > 0$  there exist  $\{v_g\}$  and  $\{v'_g\}$  such that  $w = \sum_g v_g$ ,  $w' = \sum_g v'_g$ ,  $\sum_g \|v_g\|_p \leq \|w\|_{(k,p)} + \epsilon/2$ , and  $\sum_g \|v'_g\|_p \leq \|w'\|_{(k,p)} + \epsilon/2$ . As  $\sum_g v_g + \sum_g v'_g = w + w'$ , we have

$$\begin{aligned} \|w + w'\|_{(k,p)} &\leq \sum_g \|v_g\|_p + \sum_g \|v'_g\|_p \\ &\leq \|w\|_{(k,p)} + \|w'\|_{(k,p)} + \epsilon, \end{aligned}$$

and the result follows by letting  $\epsilon$  tend to zero.

Note that, since a convex set is equivalent to the convex hull of its extreme points, Definition 1 implies that the unit ball of the  $(k, p)$ -support norm, denoted by  $C_k^p$ , is given by the convex hull of the set of vectors with cardinality no greater than  $k$  and  $\ell_p$ -norm no greater than 1, that is

$$C_k^p = \text{co} \{w \in \mathbb{R}^d : \text{card}(w) \leq k, \|w\|_p \leq 1\}. \quad (5)$$

Definition 1 gives the norm as the solution of a variational problem. Its explicit computation is not straightforward in the general case, however for  $p = 1$  the unit ball (5) does not depend on  $k$  and is always equal to the  $\ell_1$  unit ball. Thus, the  $(k, 1)$ -support norm

is always equal to the  $\ell_1$ -norm, and we do not consider further this case in this section. Similarly, for  $k = 1$  we recover the  $\ell_1$ -norm for all values of  $p$ . For  $p = \infty$ , from the definition of the dual norm it is not difficult to show that  $\|\cdot\|_{(k,p)} = \max\{\|\cdot\|_\infty, \|\cdot\|_1/k\}$ . We return to this in Section 4 when we describe how to compute the norm for all values of  $p$ .

Note further that in Equation (4), as  $p$  tends to  $\infty$ , the  $\ell_p$ -norm of each  $v_g$  is increasingly dominated by the largest component of  $v_g$ . As the variational formulation tries to identify vectors  $v_g$  with small aggregate  $\ell_p$ -norm, this suggests that higher values of  $p$  encourage each  $v_g$  to tend to a vector whose  $k$  entries are equal. In this manner varying  $p$  allows us adjust the degree to which the components of vector  $w$  can be clustered into (possibly overlapping) groups of size  $k$ .

As in the case of the  $k$ -support norm, the dual  $(k, p)$ -support norm has a simple expression. Recall that the dual norm of a vector  $u \in \mathbb{R}^d$  is defined by the optimization problem

$$\|u\|_{(k,p),*} = \max \{ \langle u, w \rangle : \|w\|_{(k,p)} = 1 \}. \quad (6)$$

**Proposition 2.** If  $p \in (1, \infty]$  then the dual  $(k, p)$ -support norm is given by

$$\|u\|_{(k,p),*} = \left( \sum_{i \in I_k} |u_i|^q \right)^{\frac{1}{q}}, \quad u \in \mathbb{R}^d,$$

where  $q = p/(p-1)$  and  $I_k \subset \mathbb{N}_d$  is the set of indices of the  $k$  largest components of  $u$  in absolute value. Furthermore, if  $p \in (1, \infty)$  and  $u \in \mathbb{R}^d \setminus \{0\}$  then the maximum in (6) is attained for

$$w_i = \begin{cases} \text{sign}(u_i) \left( \frac{|u_i|}{\|u\|_{(k,p),*}} \right)^{\frac{1}{p-1}} & \text{if } i \in I_k, \\ 0 & \text{otherwise.} \end{cases} \quad (7)$$

If  $p = \infty$  the maximum is attained for

$$w_i = \begin{cases} \text{sign}(u_i) & \text{if } i \in I_k, u_i \neq 0, \\ \lambda_i \in [-1, 1] & \text{if } i \in I_k, u_i = 0, \\ 0 & \text{otherwise.} \end{cases}$$

Note that for  $p = 2$  we recover the dual of the  $k$ -support norm in (3).

#### 3.1 The Spectral $(k, p)$ -Support Norm

From Definition 1 it is clear that the  $(k, p)$ -support norm is a symmetric gauge function. This follows since  $\mathcal{G}_k$  contains all groups of cardinality  $k$  and the  $\ell_p$ -norms only involve absolute values of the components. Hence we can define the spectral  $(k, p)$ -support norm as

$$\|W\|_{(k,p)} = \|\sigma(W)\|_{(k,p)}, \quad W \in \mathbb{R}^{d \times m}.$$

Since the dual of any orthogonally invariant norm is given by  $\|\cdot\|_* = \|\sigma(\cdot)\|_*$ , see e.g. Lewis [1995], we conclude that the dual spectral  $(k, p)$ -support norm is given by

$$\|Z\|_{(k,p),*} = \|\sigma(Z)\|_{(k,p),*}, \quad Z \in \mathbb{R}^{d \times m}.$$

The next result characterizes the unit ball of the spectral  $(k, p)$ -support norm. Due to the relationship between an orthogonally invariant norm and its corresponding symmetric gauge function, we see that the cardinality constraint for vectors generalizes in a natural manner to the rank operator for matrices.

**Proposition 3.** *The unit ball of the spectral  $(k, p)$ -support norm is the convex hull of the set of matrices of rank at most  $k$  and Schatten  $p$ -norm no greater than one.*

In particular, if  $p = \infty$ , the dual vector norm is given by  $u \in \mathbb{R}^d$ , by  $\|u\|_{(k,\infty),*} = \sum_{i=1}^k |u_i|^\downarrow$ . Hence, for any  $Z \in \mathbb{R}^{d \times m}$ , the dual spectral norm is given by  $\|Z\|_{(k,\infty),*} = \sum_{i=1}^k \sigma_i(Z)$ , that is the sum of the  $k$  largest singular values, which is also known as the Ky-Fan  $k$ -norm, see e.g. Bhatia [1997].

### 4 Computing the Norm

In this section we compute the norm, illustrating how it interpolates between the  $\ell_1$  and  $\ell_p$ -norms.

**Theorem 4.** *Let  $p \in (1, \infty)$ . For every  $w \in \mathbb{R}^d$ , and  $k \leq d$ , it holds that*

$$\|w\|_{(k,p)} = \left[ \sum_{i=1}^{\ell} (|w_i|^\downarrow)^p + \left( \frac{\sum_{i=\ell+1}^d |w_i|^\downarrow}{\sqrt[p]{k-\ell}} \right)^p \right]^{\frac{1}{p}} \quad (8)$$

where  $\frac{1}{p} + \frac{1}{q} = 1$ , and for  $k = d$ , we set  $\ell = d$ , otherwise  $\ell$  is the largest integer in  $\{0, \dots, k-1\}$  satisfying

$$(k-\ell)|w|_\ell^\downarrow \geq \sum_{i=\ell+1}^d |w_i|^\downarrow. \quad (9)$$

Furthermore, the norm can be computed in  $\mathcal{O}(d \log d)$  time.

*Proof.* Note first that in (8) when  $\ell = 0$  we understand the first term in the right hand side to be zero, and when  $\ell = d+1$  we understand the second term to be zero.

We need to compute

$$\|w\|_{(k,p)} = \max \left\{ \sum_{i=1}^d u_i w_i : \|u\|_{(k,p),*} \leq 1 \right\}$$

where the dual norm  $\|\cdot\|_{(k,p),*}$  is described in Proposition 2. Let  $z_i = |w_i|^\downarrow$ . The problem is then equivalent to

$$\max \left\{ \sum_{i=1}^d z_i u_i : \sum_{i=1}^k u_i^q \leq 1, u_1 \geq \dots \geq u_d \right\}. \quad (10)$$

This further simplifies to the  $k$ -dimensional problem

$$\max \left\{ \sum_{i=1}^{k-1} u_i z_i + u_k \sum_{i=k}^d z_i : \sum_{i=1}^k u_i^q \leq 1, u_1 \geq \dots \geq u_k \right\}.$$

Note that when  $k = d$ , the solution is given by the dual of the  $\ell_q$ -norm, that is the  $\ell_p$ -norm. For the remainder of the proof we assume that  $k < d$ . We can now attempt to use Holder's inequality, which states that for all vectors  $x$  such that  $\|x\|_q = 1$ ,  $\langle x, y \rangle \leq \|y\|_p$ , and the inequality is tight if and only if

$$x_i = \left( \frac{|y_i|}{\|y\|_p} \right)^{p-1} \text{sign}(y_i).$$

We use it for the vector  $y = (z_1, \dots, z_{k-1}, \sum_{i=k}^d z_i)$ . The components of the maximizer  $u$  satisfy  $u_i = \left( \frac{z_i}{M_{k-1}} \right)^{p-1}$  if  $i \leq k-1$ , and

$$u_k = \left( \frac{\sum_{i=\ell+1}^d z_i}{M_{k-1}} \right)^{p-1}.$$

where for every  $\ell \in \{0, \dots, k-1\}$ ,  $M_\ell$  denotes the r.h.s. in equation (8). We then need to verify that the ordering constraints are satisfied. This requires that

$$(z_{k-1})^{p-1} \geq \left( \sum_{i=k}^d z_i \right)^{p-1}$$

which is equivalent to inequality (9) for  $\ell = k-1$ . If this inequality is true we are done, otherwise we set  $u_k = u_{k-1}$  and solve the smaller problem

$$\max \left\{ \sum_{i=1}^{k-2} u_i z_i + u_{k-1} \sum_{i=k-1}^d z_i : \sum_{i=1}^{k-2} u_i^q + 2u_{k-1}^q \leq 1, u_1 \geq \dots \geq u_{k-1} \right\}.$$

We use again Hölder's inequality and keep the result if the ordering constraints are fulfilled. Continuing in this way, the generic problem we need to solve is

$$\max \left\{ \sum_{i=1}^{\ell} u_i z_i + u_{\ell+1} \sum_{i=\ell+1}^d z_i : \sum_{i=1}^{\ell} u_i^q + (k-\ell)u_{\ell+1}^q \leq 1, u_1 \geq \dots \geq u_{\ell+1} \right\}$$

where  $\ell \in \{0, \dots, k-1\}$ . Without the ordering constraints the maximum,  $M_\ell$ , is obtained by the change of variable  $u_{\ell+1} \mapsto (k-\ell)^{\frac{1}{q}} u_\ell$  followed by applying Hölder's inequality. A direct computation provides that the maximizer is  $u_i = \left(\frac{z_i}{M_\ell}\right)^{p-1}$  if  $i \leq \ell$ , and

$$(k-\ell)^{\frac{1}{q}} u_{\ell+1} = \left(\frac{\sum_{i=\ell+1}^d z_i}{(k-\ell)^{\frac{1}{q}} M_\ell^p}\right)^{p-1}.$$

Using the relationship  $\frac{1}{p} + \frac{1}{q} = 1$ , we can rewrite this as

$$u_{\ell+1} = \left(\frac{\sum_{i=\ell+1}^d z_i}{(k-\ell) M_\ell^p}\right)^{p-1}.$$

Hence, the ordering constraints are satisfied if

$$z_\ell^{p-1} \geq \left(\frac{\sum_{i=\ell+1}^d z_i}{(k-\ell)}\right)^{p-1},$$

which is equivalent to (9). Finally note that  $M_\ell$  is a nondecreasing function of  $\ell$ . This is because the problem with a smaller value of  $\ell$  is more constrained, namely, it solves (10) with the additional constraints  $u_{\ell+1} = \dots = u_d$ . Moreover, if the constraint (9) holds for some value  $\ell \in \{0, \dots, k-1\}$  then it also holds for a smaller value of  $\ell$ , hence we maximize the objective by choosing the largest  $\ell$ .

The computational complexity stems from using the monotonicity of  $M_\ell$  with respect to  $\ell$ , which allows us to identify the critical value of  $\ell$  using binary search.  $\square$

Note that for  $k = d$  we recover the  $\ell_p$ -norm and for  $p = 2$  we recover the result in Argyriou et al. [2012], McDonald et al. [2014], however our proof technique is different from theirs.

**Remark 5** (Computation of the norm for  $p \in \{1, \infty\}$ ). *Since the norm  $\|\cdot\|_{(k,p)}$  computed above for  $p \in (1, \infty)$  is continuous in  $p$ , the special cases  $p = 1$  and  $p = \infty$  can be derived by a limiting argument. We readily see that for  $p = 1$  the norm does not depend on  $k$  and it is always equal to the  $\ell_1$ -norm, in agreement with our observation in the previous section. For  $p = \infty$  we obtain that  $\|w\|_{(k,\infty)} = \max(\|w\|_\infty, \|w\|_1/k)$ .*

## 5 Optimization

In this section, we describe how to solve regularization problems using the vector and matrix  $(k,p)$ -support norms. We consider the constrained optimization problem

$$\min \{f(w) : \|w\|_{(k,p)} \leq \alpha\}, \quad (11)$$

---

### Algorithm 1 Frank-Wolfe.

---

Choose  $w^{(0)}$  such that  $\|w^{(0)}\|_{(k,p)} \leq \alpha$   
**for**  $t = 0, \dots, T$  **do**  
    Compute  $g := \nabla f(w^{(t)})$   
    Compute  $s := \operatorname{argmin} \{ \langle s, g \rangle : \|s\|_{(k,p)} \leq \alpha \}$   
    Update  $w^{(t+1)} := (1 - \gamma)w^{(t)} + \gamma s$ , for  $\gamma := \frac{2}{t+2}$   
**end for**

---

where  $w$  is in  $\mathbb{R}^d$  or  $\mathbb{R}^{d \times m}$ ,  $\alpha > 0$  is a regularization parameter and the error function  $f$  is assumed to be convex and continuously differentiable. For example, in linear regression a valid choice is the square error,  $f(w) = \|Xw - y\|_2^2$ , where  $X$  is a matrix of observations and  $y$  a vector of response variables. Constrained problems of form (11) are also referred to as Ivanov regularization in the inverse problems literature [Ivanov et al., 1978].

A convenient tool to solve problem (11) is provided by the *Frank-Wolfe* method [Frank and Wolfe, 1956], see also Jaggi [2013] for a recent account. The method is outlined in Algorithm 1, and it has worst case convergence rate  $\mathcal{O}(1/T)$ . The key step of the algorithm is to solve the subproblem

$$\operatorname{argmin} \{ \langle s, g \rangle : \|s\|_{(k,p)} \leq \alpha \}, \quad (12)$$

where  $g = \nabla f(w^{(t)})$ , that is the gradient of the objective function at the  $t$ -th iteration. This problem involves computing a subgradient of the dual norm at  $g$ . It can be solved exactly and efficiently as a consequence of Proposition 2. We discuss here the vector case and postpone the discussion of the matrix case to Section 5.2. By symmetry of the  $\ell_p$ -norm, problem (12) can be solved in the same manner as the maximum in Proposition 2 and the solution is given by  $s_i = -\alpha w_i$ , where  $w_i$  is given by (7).

Specifically, letting  $I_k \subset \mathbb{N}_d$  be the set of indices of the  $k$  largest components of  $g$  in absolute value, for  $p \in (1, \infty)$  we have

$$s_i = \begin{cases} \alpha \operatorname{sign}(g_i) \left(\frac{|g_i|}{\|g\|_{(k,p),*}}\right)^{\frac{1}{p-1}}, & \text{if } i \in I_k \\ 0, & \text{if } i \notin I_k \end{cases} \quad (13)$$

and, for  $p = \infty$  we choose the subgradient

$$s_i = \begin{cases} \alpha \operatorname{sign}(g_i) & \text{if } i \in I_k, g_i \neq 0, \\ 0 & \text{otherwise.} \end{cases} \quad (14)$$

### 5.1 Projection Operator

An alternative method to solve (11) in the vector case is to consider the equivalent problem

$$\min \left\{ f(w) + \delta_{\{\|\cdot\|_{(k,p)} \leq \alpha\}}(w) : w \in \mathbb{R}^d \right\}, \quad (15)$$

where  $\delta_C(\cdot)$  is the indicator function of convex set  $C$ . Proximal gradient methods can be used to solve optimization problems of the form  $\min \{f(w) + \lambda g(w) : w \in \mathbb{R}^d\}$ , where  $f$  is a convex loss function with Lipschitz continuous gradient,  $\lambda > 0$  is a regularization parameter, and  $g$  is a convex function for which the proximity operator can be computed efficiently, see Beck and Teboulle [2009], Combettes and Wajs [2005], Nesterov [2007] and references therein. The proximity operator of  $g$  with parameter  $\rho > 0$  is defined as  $\text{prox}_{\rho g}(w) = \text{argmin}\{\frac{1}{2}\|x - w\|^2 + \rho g(x) : x \in \mathbb{R}^d\}$ . The proximity operator for the squared  $k$ -support norm was computed by Argyriou et al. [2012] and McDonald et al. [2014], and for the  $k$ -support norm by Chatterjee et al. [2014].

In the special case that  $g(w) = \delta_C(w)$ , where  $C$  is a convex set, the proximity operator reduces to the projection operator onto  $C$ . For the  $(k, p)$ -support norm, for the case  $p = \infty$  we can compute the projection onto its unit ball using the following result.

**Proposition 6.** *For every  $w \in \mathbb{R}^d$ , the projection  $x$  of  $w$  onto the unit ball of the  $(k, \infty)$ -norm is given by*

$$x_i = \begin{cases} \text{sign}(w_i)(|w_i| - \beta), & \text{if } ||w_i| - \beta| \leq 1, \\ \text{sign}(w_i), & \text{if } ||w_i| - \beta| > 1, \end{cases} \quad (16)$$

where  $\beta = 0$  if  $\|w\|_1 \leq k$ , otherwise  $\beta \in (0, \infty)$  is chosen to maximize  $\sum_{i=1}^d |x_i|$  subject to the constraint  $\sum_{i=1}^d |x_i| \leq k$ . Furthermore, the projection can be computed in  $\mathcal{O}(d \log d)$  time.

*Proof.* (Sketch) We assume without loss of generality that the components of  $w$  are non zero. We solve the optimization problem

$$\min_{x \in \mathbb{R}^d} \left\{ \sum_{i=1}^d (x_i - w_i)^2 : |x_i| \leq 1, \sum_{i=1}^d |x_i| \leq k \right\}. \quad (17)$$

We consider two cases. If  $\sum_{i=1}^d |w_i| \leq k$ , then the problem decouples and we solve it componentwise. If  $\sum_{i=1}^d |w_i| > k$ , we solve problem (17) by minimizing the Lagrangian function  $\mathcal{L}(x, \beta) = \sum_{i=1}^d (x_i - w_i)^2 + 2\beta(\sum_{i=1}^d |x_i| - k)$  with nonnegative multiplier  $\beta$ , which can be done componentwise. Finally, both cases can be combined into the form of (16). The complexity follows by taking advantage of the monotonicity of  $x_i(\beta)$ .  $\square$

We can use Proposition 6 to project onto the unit ball of radius  $\alpha > 0$  by a rescaling argument (see supplementary material for details).

## 5.2 Matrix Problem

Given data matrix  $X \in \mathbb{R}^{d \times m}$  for which we observe a subset of entries, we consider the constrained optimization problem

$$\min_{W \in \mathbb{R}^{d \times m}} \{ \|\Omega(X) - \Omega(W)\|_F : \|W\|_{(k,p)} \leq \alpha \} \quad (18)$$

where the operator  $\Omega$  applied to a matrix sets unobserved values to zero. As in the vector case, the Frank-Wolfe method can be applied to the matrix problems. Algorithm 1 is particularly convenient in this case as we only need to compute the largest  $k$  singular values, which can result in a computationally efficient algorithm. The result is a direct consequence of Proposition 2 and von Neumann’s trace inequality, see e.g. Marshall and Olkin [1979, Ch. 9 Sec. H.1.h].

Specifically, we solve the optimization problem

$$\min_S \{ \langle S, G \rangle : \|S\|_{(k,p)} \leq \alpha \},$$

where  $G$  is the gradient of the objective function at the  $t$ -th iteration. Equivalently we consider

$$\max_R \{ \langle R, G \rangle : \|R\|_{(k,p)} \leq \alpha \},$$

where  $R = -S$ . The solution is given by

$$R = U_k \text{diag}(r) V_k^T$$

where  $U_k$  and  $V_k$  are the top  $k$  left and right singular vectors of the gradient  $G$  of the objective function in (18) evaluated at the current solution, whose singular values we denote by  $g$ , and  $r$  is obtained from  $g$  as

$$r_i = \begin{cases} \alpha \left( \frac{g_i}{\|g\|_{(k,p),*}} \right)^{\frac{1}{p-1}}, & \text{if } i \in I_k \\ 0, & \text{if } i \notin I_k \end{cases} \quad (19)$$

and, for  $p = \infty$

$$r_i = \begin{cases} \alpha & \text{if } i \in I_k, g_i \neq 0, \\ 0 & \text{otherwise} \end{cases} \quad (20)$$

for  $p \in (1, \infty)$  and  $p = \infty$ , respectively, where  $I_k = [1, \dots, k]$ . It follows that a subgradient for the matrix problem is  $R = -U_k \text{diag}(s) V_k^T$ .

Note also that the proximity operator of the norm and the Euclidean projection on the associated unit ball both require the full singular value decomposition to be performed. Indeed, the proximity operator of an orthogonally invariant norm  $\|\cdot\| = g(\sigma(\cdot))$  at  $W \in \mathbb{R}^{d \times m}$  is given by  $\text{prox}_{\|\cdot\|}(W) = U \text{diag}(\text{prox}_g(\sigma(W))) V^T$ , where  $U$  and  $V$  are the matrices formed by the left and right singular vectors of  $W$ , see e.g. Argyriou et al. [2011, Prop. 3.1], and this requires the full decomposition.

Table 1: Matrix completion on synthetic datasets generated with decaying spectrum. The improvement of the  $(k, p)$ -support norm over the  $k$ -support and trace norms is statistically significant at a level  $< 0.001$ .

dataset	norm	test error	$k$	$p$
rank 5 $\rho=10\%$	trace	0.8184 (0.03)	-	-
	$k$ -supp	0.8036 (0.03)	3.6	-
	$kp$ -supp	0.7831 (0.03)	1.8	7.3
rank 5 $\rho=20\%$	trace	0.4085 (0.03)	-	-
	$k$ -supp	0.4031 (0.03)	3.1	-
	$kp$ -supp	0.3996 (0.03)	2.0	4.7
rank 10 $\rho=20\%$	trace	0.6356 (0.03)	-	-
	$k$ -supp	0.6284 (0.03)	4.4	-
	$kp$ -supp	0.6270 (0.03)	2.0	4.4

## 6 Numerical Experiments

In this section we apply the spectral  $(k, p)$ -support norm to matrix completion (collaborative filtering problems), in which we want to recover a low rank, or approximately low rank, matrix from a small sample of its entries, see e.g. Jaggi and Sulovsky [2010]. One prominent method of solving this problem is trace norm regularization: we look for a matrix which closely fits the observed entries and has a small trace norm (sum of singular values) [Jaggi and Sulovsky, 2010, Mazumder et al., 2010, Toh and Yun, 2011]. We apply the  $(k, p)$ -support norm to this framework and we investigate the impact of varying  $p$ , and we compare the spectral  $(k, p)$ -support norm to the trace norm and the spectral  $k$ -support norm ( $p = 2$ ) in both synthetic and real datasets. In each case we solve the optimization problem (18) using the Frank-Wolfe method as outlined in Section 5. We determine the values of  $k$  and  $p \geq 1$  by validation, averaged over a number of trials.

**Impact of  $p$ .** A key motivation for the additional parameter  $p$  is that it allows us to tune the norm to the decay of the singular values of the underlying matrix. In particular the variational formulation of (4) suggests that as the spectrum of the true low rank matrix flattens out, larger values of  $p$  should be preferred.

We ran the method on a set of  $100 \times 100$  matrices of rank 12, with decay of the non zero singular values  $\sigma_\ell$  proportional to  $\exp(-\ell a)$ , for 26 values of  $a$  between  $10^{-6}$  and 0.18, and we determined the corresponding optimal value of  $p$ . Figure 1 illustrates the optimum value of  $p$  as a function of  $a$ . We clearly observe the negative slope, that is the steeper the slope the smaller the optimal value of  $p$ . Figure 2 shows the spectrum and the optimal  $p$  for several decay values.

Note that  $k$  is never equal to 1, which is a special case




Figure 1: Optimal  $p$  vs. decay  $a$ .




Figure 2: Optimal  $p$  fitted to Matrix spectra with various decays.

in which the norm is independent of  $p$ , and is equal to the trace norm. In each case the improvement of the spectral  $(k, p)$ -support norm over the  $k$ -support and trace norms is statistically significant at a level  $< 0.001$ .

Figure 3 illustrates the impact of the curvature  $p$  on the test error on synthetic and real datasets. We observe that the error levels off as  $p$  tends to infinity, so for these specific datasets the major gain is to be had for small values of  $p$ . The optimum value of  $p$  for both the real and synthetic datasets is statistically different from  $p = 2$  ( $k$ -support norm), and  $p = 1$  (trace norm).

**Simulated Data.** We replicated the experimental setting of McDonald et al. [2014] for synthetic matrix completion. Each  $100 \times 100$  matrix is generated as  $W = UV^\top + E$ , where  $U, V \in \mathbb{R}^{100 \times r}$ ,  $r \ll 100$ , and the entries of  $U, V$  and  $E$  are i.i.d. standard Gaussian.

Table 2: Matrix completion on real datasets. The improvement of the  $(k, p)$ -support norm over the  $k$ -support and trace norms is statistically significant at a level  $< 0.001$ .

dataset	norm	test error	$k$	$p$
MovieLens 100k	trace	0.2017	-	-
	k-supp	0.1990	1.9	-
	kp-supp	0.1971	2.0	3.0
Jester 1	trace	0.1752	-	-
	k-supp	0.1739	6.4	-
	kp-supp	0.1731	2.0	6.5
Jester 3	trace	0.1959	-	-
	k-supp	0.1942	2.1	-
	kp-supp	0.1924	5.0	5.0

The error is measured as  $\|\text{true} - \text{predicted}\|^2 / \|\text{true}\|^2$ , standard deviations are shown in brackets and the mean values of  $k$  and  $p$  are selected by validation. Table 1 illustrates the results. We found that the  $(k, p)$ -support norm outperformed the standard  $k$ -support norm, as well as the trace norm, at a statistically significant level.

We note that although Frank Wolfe method for the  $(k, p)$ -support norm does not generally converge as quickly as proximal methods (which are available in the case of  $k$ -support norm [McDonald et al., 2014, Chatterjee et al., 2014]), the computational cost can be mitigated using a continuation method. Specifically given an ordered sequence of parameter values for  $p$  we can proceed sequentially, initializing its value based on the previously computed value. Empirically we tried this approach for a range of values of  $p$  and found that the total computation time increased only moderately.

**Real Data.** Finally, we applied the norms to real collaborative filtering datasets. We observe a subset of the (user, rating) entries of a matrix and predict the unobserved ratings, with the assumption that the true matrix is likely to have low rank. We report on the MovieLens 100k dataset (<http://grouplens.org/datasets/movielens/>), which consists of ratings of movies, and the Jester 1 and 3 datasets (<http://goldberg.berkeley.edu/jester-data/>), which consist of ratings of jokes. Following McDonald et al. [2014], Toh and Yun [2011], for MovieLens for each user we uniformly sampled  $\rho = 50\%$  of available entries for training, and for Jester 1 and Jester 3 we sampled 20, respectively 8 ratings per user, using 10% for validation. We used normalized mean absolute error,

$$\text{NMAE} = \frac{\|\text{true} - \text{predicted}\|^2}{\#\text{obs.}/(r_{\max} - r_{\min})},$$

where  $r_{\min}$  and  $r_{\max}$  are lower and upper bounds for




Figure 3: Test error vs curvature ( $p$ ). Left axis: synthetic data (blue crosses); right axis: Jester 1 dataset (red circles).

the ratings [Toh and Yun, 2011], and we implemented a final thresholding step as in McDonald et al. [2014].

The results are outlined in Table 2. The spectral  $(k, p)$ -support outperformed the trace norm and the spectral  $k$ -support norm, and the improvement is statistically significant at a level  $< 0.001$  (the standard deviations, not shown here, are of the order of  $10^{-5}$ ). In summary the experiments suggest that the additional flexibility of the  $p$  parameter indeed allows the model to better fit both the sparsity and the decay of the true spectrum.

## 7 Conclusion

We presented a generalization of the  $k$ -support norm, the  $(k, p)$ -support norm, where the additional parameter  $p$  is used to better fit the decay of the components of the underlying model. We determined the dual norm, characterized the unit ball and computed an explicit expression for the norm. As the norm is a symmetric gauge function, we further described the induced spectral  $(k, p)$ -support norm. We adapted the Frank-Wolfe method to solve regularization problems with the norm, and in the particular case  $p = \infty$  we provided a fast computation for the projection operator. In numerical experiments we considered synthetic and real matrix completion problems and we showed that varying  $p$  leads to significant performance improvements. Future work could include deriving statistical bounds for the performance of the norms, and situating the norms in the framework of other structured sparsity norms which have recently been studied [Micchelli et al., 2013, Obozinski and Bach, 2012].



## References

- A. Argyriou, C. Micchelli, M. Pontil, and Y. Ying. A spectral regularization framework for multi-task structure learning. *Advances in Neural Information Processing Systems 20*, pages 25–32, 2007.
- A. Argyriou, T. Evgeniou, and M. Pontil. Convex multi-task feature learning. *Machine Learning*, 73(3):243–272, 2008.
- A. Argyriou, C. A. Micchelli, M. Pontil, L. Shen, and Y. Xu. Efficient first order methods for linear composite regularizers. *CoRR*, arXiv/1104.1436, 2011.
- A. Argyriou, R. Foygel, and N. Srebro. Sparse prediction with the k-support norm. *Advances in Neural Information Processing Systems 25*, pages 1466–1474, 2012.
- A. Beck and M. Teboulle. A fast iterative shrinkage-thresholding algorithm for linear inverse problems. *SIAM J. Imaging Sciences*, 2(1):183–202, 2009.
- R. Bhatia. *Matrix Analysis*. Springer, 1997.
- P. Buehlmann and S. A. van der Geer. *Statistics for High-Dimensional Data*. Springer, 2011.
- S. Chatterjee, S. Chen, and A. Banerjee. Generalized Dantzig selector: application to the k-support norm. In *Advances in Neural Information Processing Systems 28*, pages 1934–1942, 2014.
- P. L. Combettes and V. R. Wajs. Signal recovery by proximal forward-backward splitting. *Multiscale Model. Simul.*, 2005.
- M. Frank and P. Wolfe. An algorithm for quadratic programming. *Naval Research Logistics Quarterly*, 3 (1-2):95–110, 1956.
- R. A. Horn and C. R. Johnson. *Topics in Matrix Analysis*. Cambridge University Press, 1991.
- V.K. Ivanov, V. V. Vasin, and V.P. Tanana. *Theory of Linear Ill-Posed Problems and its Applications*. De Gruyter, 1978.
- L. Jacob, G. Obozinski, and J.-P. Vert. Group lasso with overlap and graph lasso. *Proceedings of the 26th International Conference on Machine Learning*, pages 433–440, 2009.
- M. Jaggi. Revisiting Frank-Wolfe: Projection-free sparse convex optimization. *Proceedings of the 30th International Conference on Machine Learning*, pages 427–435, 2013.
- M Jaggi and M. Sulovsky. A simple algorithm for nuclear norm regularized problems. *Proceedings of the 27th International Conference on Machine Learning*, pages 471–478, 2010.
- A. S. Lewis. The convex analysis of unitarily invariant matrix functions. *Journal of Convex Analysis*, 2: 173–183, 1995.
- A. W. Marshall and I. Olkin. *Inequalities: Theory of Majorization and its Applications*. Academic Press, 1979.
- R. Mazumder, T. Hastie, and R. Tibshirani. Spectral regularization algorithms for learning large incomplete matrices. *Journal of Machine Learning Research*, 11:2287–2322, 2010.
- A. M. McDonald, M. Pontil, and D. Stamos. New perspectives on k-support and cluster norms. *Journal of Machine Learning Research*, Forthcoming.
- A. M. McDonald, M. Pontil, and D. Stamos. Spectral k-support regularization. In *Advances in Neural Information Processing Systems 28*, pages 3644–3652, 2014.
- C. A. Micchelli, J. M. Morales, and M. Pontil. Regularizers for structured sparsity. *Advances in Comp. Mathematics*, 38:455–489, 2013.
- Y. Nesterov. Gradient methods for minimizing composite objective function. *Center for Operations Research and Econometrics*, 76, 2007.
- G. Obozinski and F. Bach. Convex relaxation for combinatorial penalties. *CoRR*, arXiv/1205.1240, 2012.
- R. T. Rockafellar. *Convex Analysis*. Princeton University Press, 1970.
- A. Rohde and A.B. Tsybakov. Estimation of high-dimensional low rank matrices. *Annals of Statistics*, 39:887–930, 2011.
- W. Rudin. *Functional Analysis*. McGraw Hill, 1991.
- R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society*, 58:267–288, 1996.
- K.-C. Toh and S. Yun. An accelerated proximal gradient algorithm for nuclear norm regularized least squares problems. *SIAM J. on Img. Sci.*, 4:573–596, 2011.
- J. Von Neumann. *Some matrix-inequalities and metrization of matrix-space*. Tomsk. Univ. Rev. Vol I, 1937.
- M. Wainwright. Structured regularizers for high-dimensional problems. *Annual Review of Statistics and Its Application*, 1:233–253, 2014.
- H. Zou and T. Hastie. Regularization and variable selection via the elastic net. *Journal of the Royal Statistical Society, Series B*, 67(2):301–320, 2005.