VECTOR AND MATRIX NORMS

David F. Gleich

August 21, 2023

Learning objectives
1. Examples of vector norms.
2. Examples of matrix norms.
3. The submultiplicative property of a matrix norm.
4. The property that *all norms are equivalent*

Norms are used to measure the size of vectors and matrices. They are generalizations of the scalar function $|x|$, which determines the size or magnitude of a scalar value. For instance, if x is close to y, then we have $|x - y|$ is close to zero.

So far, we have used the 2-norm of a vector. Let's work with them formally.

1 VECTOR NORMS

DEFINITION 1 *The Euclidean norm or 2-norm of a vector*[1] *is*

$$\|\mathbf{x}\| = \sqrt{\sum_{i=1}^{n} |x_i|^2} = \sqrt{\mathbf{x}^T\mathbf{x}}$$

[1] This definition includes absolute values. Yet, $x_i^2 \geq 0$ for all real values. We leave the absolute values because this then generalizes to complex values where we need a complex magnitude.

This can be generalized a p-norm.

DEFINITION 2 *The p-norm of a vector is*

$$\|\mathbf{x}\|_p = \left(\sum_{i=1}^{n} |x_i|^p\right)^{1/p}.$$

This isn't created to generalize for generalizations sake. One of the common uses of norms is to argue that a sequence of vectors

$$\mathbf{x}_k \to \mathbf{y}$$

which can be handled by showing

$$\|\mathbf{x}_k - \mathbf{y}\| \to 0.$$

Depending on the value of p, this can be easy or difficult. For instance, when $p = 1$, then this is simply a sum of absolute values:

$$\|\mathbf{x}\|_1 = \sum_{i=1}^{n} |x_i|$$

and $p = \infty$ can be defined via a limit:[2]

$$\|\mathbf{x}\|_\infty = \max_{i=1}^{n} |x_i|.$$

[2] It is a useful exercise to convince yourself that as $p \to \infty$, then the value of the norm will simply be the largest element by magnitude.

Now, we are going to define an extremely general notion of norm in order to state a few important results.

DEFINITION 3 *A vector norm on* $\mathbf{x} \in \mathbb{R}^n$ *is any function* $f(\mathbf{x}) \to \mathbb{R}$ *that satisfies:*

1. $f(\mathbf{x}) \geq 0$ (non-negative)
2. $f(\mathbf{x}) = 0$ *if and only if* $\mathbf{x} = 0$ (zero-sensitive)
3. $f(\alpha\mathbf{x}) = |\alpha| f(\mathbf{x})$ *for any scalar* α (1-homogeneous),
4. $f(\mathbf{x} + \mathbf{y}) \leq f(\mathbf{x}) + f(\mathbf{y})$ (triangle inequality).

Any p-norm with $p \geq 1$ satisfies these definitions. When $p < 1$, then we violate the triangle inequality.

There are some crazy norms too. For instance, the following function satisfies these three criteria:

$$f(\mathbf{x}) = \text{sum of largest two entries in } \mathbf{x} \text{ by magnitude.}$$

The following theorem guarantees that if $\mathbf{x}_k \to \mathbf{y}$ for any norm, then it will happen for all norms.

THEOREM 4 *Informally, all vector norms are equivalent. Formally, let $f(\mathbf{x})$ and $g(\mathbf{x})$ be any pair of vector norms on $\mathbb{R}^n$, then there exist positive constants $C_1 \leq C_2$ such that*

$$C_1 f(\mathbf{x}) \leq g(\mathbf{x}) \leq C_2 f(\mathbf{x}).$$

Note that these constants can depend on the dimension n.

Proof This is just a sketch, but the essence of the result is here; it requires just a little bit more analysis to fully state. They key we look at how these functions map unit-vectors to get the extreme values. Everything else follows from straightforward (but not simple) analysis. The values in the theorem are:

$$C_1 = \begin{array}{ll} \underset{\mathbf{x}}{\text{maximize}} & f(\mathbf{x}) \\ \text{subject to} & g(\mathbf{x}) \leq 1 \end{array} \quad \text{and } C_2 = \begin{array}{ll} \underset{\mathbf{x}}{\text{maximize}} & g(\mathbf{x}) \\ \text{subject to} & f(\mathbf{x}) \leq 1 \end{array}.$$

Note that because of the scalar property, the extreme must occur on a boundary of the feasible set, i.e. where $f(x) = 1$. (If this isn't obvious, a quick proof by contradiction should help: If there is a point inside that gets the max, then we can scale it and make $f(\mathbf{x})$ (say) bigger and also $g(\mathbf{x})$ bigger, so it can't be optimal.) This is why we get the values of C_1 and C_2 in the above proof. ■

For instance, $\|\mathbf{x}\|_\infty \leq \|\mathbf{x}\|_1 \leq n\|\mathbf{x}\|_\infty$ is an instance with $C_1 = 1, C_2 = n$.

Quiz What are C_1 and C_2 such that $C_1\|\mathbf{x}\|_1 \leq \|\mathbf{x}\|_\infty \leq C_2\|\mathbf{x}\|_1$?

Consequently, suppose, for a vector norm $f(\mathbf{x})$, you show that $f(\mathbf{x}_k - \mathbf{y}) \to 0$. Then we know that $C_2 f(\mathbf{x}) \geq g(\mathbf{x})$ and also that $C_2 f(\mathbf{x}_k - \mathbf{y}) \to 0$. Since $g(\mathbf{x}) \geq 0$, then we must have $g(\mathbf{x}_k - \mathbf{y}) \to 0$ as well.

2 MATRIX NORMS

Vector norms measure the size or magnitude of a vector. Matrix norms do the same for a matrix. There are two important types of matrix norms: element-wise (or Frobenius norms) and operator norms. Just like vector norms, there is a general condition for all matrix norms.

DEFINITION 5 *A matrix norm on $\boldsymbol{X} \in \mathbb{R}^{m \times n}$ is any function $f(\boldsymbol{X}) \to \mathbb{R}$ that satisfies:*

1. $f(\boldsymbol{X}) \geq 0$ (non-negative)
2. $f(\boldsymbol{X}) = 0$ *if and only if* $\boldsymbol{X} = 0$ (zero-sensitive)
3. $f(\alpha\boldsymbol{X}) = |\alpha| f(\boldsymbol{X})$ *for any scalar* α (1-homogeneous)
4. $f(\boldsymbol{X} + \boldsymbol{Y}) \leq f(\boldsymbol{X}) + f(\boldsymbol{Y})$ (triangle inequality).

2.1 ELEMENT-WISE NORMS

Note that if vec $\boldsymbol{X}$ is any way of turning $\boldsymbol{X}$ into a vector by organizing the mn elements of $\boldsymbol{X}$ into a single array, then $f(\text{vec}(\boldsymbol{X}))$ is a matrix norm for any vector norm $f(\mathbf{x})$. These are called element-wise norms. The most common of which is the Frobenius norm.

DEFINITION 6 *The Frobenius norm of a matrix is*

$$\|\boldsymbol{X}\|_F = \sqrt{\sum_{ij} |X_{ij}|^2} = \|\text{vec}(\boldsymbol{X})\|_2 = \sqrt{trace(\boldsymbol{A}^T\boldsymbol{A})}.$$

Here, we used $\text{trace}(\boldsymbol{A}) = \sum_{i=1}^{\min(m,n)} A_{i,i}$, which is the sum of diagonal entries.

2.2 OPERATOR-INDUCED NORMS

Let $f(\mathbf{x})$ be any vector norm, then we can define a matrix norm via:

$$f(\boldsymbol{X}) = \underset{\mathbf{x}\neq 0}{\text{maximize}} \quad f(\boldsymbol{A}\mathbf{x})/f(\mathbf{x}) \ .$$

Proof that $f(\boldsymbol{X})$ is a matrix norm

1. $f(\boldsymbol{X}) \geq 0$ because f is a vector norm.
2. If $f(\boldsymbol{X}) = 0$, then $f(\boldsymbol{A}\mathbf{x})/f(\mathbf{x}) = 0$ for all vectors $\mathbf{x} \neq 0$. Since $f(\mathbf{e}_i) > 0$, then we must have $f(\boldsymbol{A}\mathbf{e}_i) = 0$ for all $\mathbf{e}_i$, so the matrix is entirely empty. Also, if $\boldsymbol{A} = 0$, then $\boldsymbol{A}\mathbf{x} = 0$ for any $\mathbf{x}$, and so $f(\boldsymbol{A}) = 0$.
3. $f(\alpha\boldsymbol{X}) = \underset{\mathbf{x}\neq 0}{\text{maximize}} \quad f(\alpha\boldsymbol{A}\mathbf{x})/f(\mathbf{x}) \ = |\alpha| f(\boldsymbol{X})$.
4. Note that $f((\boldsymbol{X}+\boldsymbol{Y})\mathbf{x}) \leq f(\boldsymbol{X}\mathbf{x}) + f(\boldsymbol{Y}\mathbf{x})$ be the vector-norm triangle inequality. Hence,

$$\begin{aligned} f(\boldsymbol{X}+\boldsymbol{Y}) &= \underset{\mathbf{x}\neq 0}{\text{maximize}} \quad f((\boldsymbol{X}+\boldsymbol{Y})\mathbf{x})/f(\mathbf{x}) \ \leq \ \underset{\mathbf{x}\neq 0}{\text{maximize}} \quad f(\boldsymbol{X}\mathbf{x})/f(\mathbf{x}) + f(\boldsymbol{Y}\mathbf{x})/f(\mathbf{x}) \\ &\leq \ \underset{\mathbf{x}\neq 0}{\text{maximize}} \quad f(\boldsymbol{X}\mathbf{x})/f(\mathbf{x}) \ + \ \underset{\mathbf{x}\neq 0}{\text{maximize}} \quad f(\boldsymbol{Y}\mathbf{x})/f(\mathbf{x}) \\ &\leq f(\boldsymbol{X}) + f(\boldsymbol{Y}) \end{aligned}$$

The operator induced norms are harder to reason about.

Let $f(\mathbf{x}) = \|\mathbf{x}\|_1$, then

$$\|\boldsymbol{A}\|_1 = \max_{j=1}^{n} \sum_{i=1}^{m} |A_{ij}|$$

which is the maximum column 1-norm. If, instead, $f(\mathbf{x}) = \|\mathbf{x}\|_\infty$, then

$$\|\boldsymbol{A}\|_\infty = \max_{i=1}^{m} \sum_{j=1}^{n} |A_{ij}|$$

which is the maximum row 1-norm.

Here's my picture to remember these.

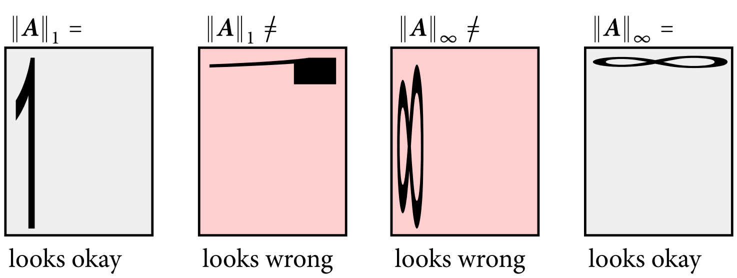

2.3 ADDITIONAL MATRIX NORMS

This section can be skipped on a first reading.

There is a wide additional class of norms defined in terms of the singular values of a matrix. See other sections on the singular values and their definitions.[3]

An $m \times n$ real-valued or complex-valued matrix has $\min(m,n)$ non-negative *real* singular values. Let $\sigma_1, \ldots, \sigma_{\min(m,n)}$ be the singular values of a $m \times n$ matrix with $m \geq n$.

DEFINITION 7 (The Nuclear Norm, the Trace Norm) *Let $\sigma_1, \ldots, \sigma_{\min(m,n)}$ be the singular values of an $m \times n$ matrix $\boldsymbol{A}$. Then the* nuclear norm *also called the* trace norm *is the matrix norm based on the function*

$$f(\boldsymbol{A}) = \sum_i \sigma_i \text{ commonly denoted } \|\boldsymbol{A}\|_*.$$

DEFINITION 8 (The Schatten Norms) *Let $\sigma_1, \ldots, \sigma_{\min(m,n)}$ be the singular values of an $m \times n$ matrix $\boldsymbol{A}$. Let $\mathbf{s}$ be the* vector *of singular values, ordered arbitrarily. Then the* Shatten p-norm *is the matrix norm based on the function*

$$f(\boldsymbol{A}) = \|\mathbf{s}\|_p.$$

[3] TODO Insert reference when assembled into bigger document.

DEFINITION 9 (The Ky-Fan Norms) *Let $\sigma_1, \ldots, \sigma_{\min(m,n)}$ be the singular values of an $m \times n$ matrix $\boldsymbol{A}$ where $\sigma_1 \ge \sigma_2 \ge \ldots \ge \sigma_{\min(m,n)}$ by convention (that is, the elements are ordered in decreasing order in most conventions). Then the* Ky-Fan p-norm *is the matrix norm based on the function*

$$f(\boldsymbol{A}) = \sum_{i=1}^{p} \sigma_i.$$

Note that both Shatten and Ky-Fan norms are *vector norms* applied to the vector of singular values $\mathbf{s}$. For Shatten norms, it is a p-norm. For Ky-Fan norms, it is the sum of the largest p elements. Indeed, any vector norm applied to the singular values of a matrix is a valid matrix norm.

3 NORM PROPERTIES

3.1 ORTHOGONAL INVARIANCE

An important property of a norm is that it is orthogonally invariant. This property is a realization of two ideas:

- *norms measure lengths*
- *orthogonal matrices generalize rotations*

When we rotate a vector, we simple change its orientation, but not its length. Consequently, we have the definition:

DEFINITION 10 (orthogonally invariant) *Let $\boldsymbol{Q}$ be a square orthogonal matrix. Then a vector norm $f(\mathbf{x})$ is* orthogonally invariant *when*

$$f(\boldsymbol{Q}\mathbf{x}) = f(\mathbf{x}) \quad \text{or written as} \quad \|\boldsymbol{Q}\mathbf{x}\| = \|\mathbf{x}\|.$$

Let $\boldsymbol{A}$ be an $m \times n$ matrix. Let $\boldsymbol{U}$ be a square $m \times m$ orthogonal matrix and let $\boldsymbol{V}$ be a square $n \times n$ orthogonal matrix. Then a matrix norm $f(\boldsymbol{A})$ is orthogonally invariant *when*

$$f(\boldsymbol{U}\boldsymbol{A}\boldsymbol{V}) = f(\boldsymbol{A}) \quad \text{or written as} \quad \|\boldsymbol{U}\boldsymbol{A}\boldsymbol{V}\| = \|\boldsymbol{A}\|.$$

3.2 SUBMULTIPLICATIVE

Note that operator-induced matrix norms satisfy the property that:

$$f(\boldsymbol{A}\mathbf{x}) \le f(\boldsymbol{A}) f(\mathbf{x})$$

which is handy for studying iterative algorithms! This property has the special name: *sub-multiplicative*.

DEFINITION 11 *A matrix-norm $f(\boldsymbol{A})$ is sub-multiplicative if:*

$$f(\boldsymbol{A}\boldsymbol{B}) \le f(\boldsymbol{A}) f(\boldsymbol{B}).$$

As you'll see on the homework, not all norms are sub-multiplicative. But we can always scale a norm to be sub-multiplicative.

4 EXERCISES

1. Let $\mathbf{x} \in \mathbb{C}^n$. Decompose $\mathbf{x}$ into the real and imaginary parts: $\mathbf{x} = \mathbf{y} + i\mathbf{z}$ where $\mathbf{y} \in \mathbb{R}^n$ and $\mathbf{z} \in \mathbb{R}^n$. Show that $\|\mathbf{x}\|_2 = \left\| \begin{bmatrix} \mathbf{x} \\ \mathbf{y} \end{bmatrix} \right\|_2$.
2. Let $\boldsymbol{P}$ be a permutation matrix. So $\boldsymbol{P}\mathbf{x}$ reorders the elements of $\mathbf{x}$. Find a vector-norm function on length 2 vectors where $\|\mathbf{x}\| \ne \|\boldsymbol{P}\mathbf{x}\|$.
3. (This requires knowledge of the SVD.) Show that the Schatten and Ky-Fan norms are orthogonally invariant.