



On Euclidean norm approximations

M. Emre Celebi^{a,*}, Fatih Celiker^b, Hassan A. Kingravi^c

^a Department of Computer Science, Louisiana State University, Shreveport, LA, USA

^b Department of Mathematics, Wayne State University, Detroit, MI, USA

^c Department of Computer Science, Georgia Institute of Technology, Atlanta, GA, USA

ARTICLE INFO

Article history:

Received 25 May 2010

Accepted 23 August 2010

Keywords:

Euclidean norm

Approximation

ABSTRACT

Euclidean norm calculations arise frequently in scientific and engineering applications. Several approximations for this norm with differing complexity and accuracy have been proposed in the literature. Earlier approaches [1–3] were based on minimizing the maximum error. Recently, Seol and Cheun [4] proposed an approximation based on minimizing the average error. In this paper, we first examine these approximations in detail, show that they fit into a single mathematical formulation, and compare their average and maximum errors. We then show that the maximum errors given by Seol and Cheun are significantly optimistic.

© 2010 Elsevier Ltd. All rights reserved.

1. Introduction

The Minkowski (L_p) metric is inarguably one of the most commonly used quantitative distance (dissimilarity) measures in scientific and engineering applications. The Minkowski distance between two vectors $\mathbf{x}=(x_1,x_2,\dots,x_n)$ and $\mathbf{y}=(y_1,y_2,\dots,y_n)$ in the n -dimensional Euclidean space, \mathbb{R}^n , is given by

$$L_p(\mathbf{x},\mathbf{y}) = \left(\sum_{i=1}^n |x_i - y_i|^p \right)^{1/p}. \quad (1)$$

Three special cases of the L_p metric are of particular interest, namely, L_1 (city-block metric), L_2 (Euclidean metric), and L_∞ (chessboard metric). Given the general form (1), L_1 and L_2 can be defined in a straightforward fashion, while L_∞ is defined as

$$L_\infty(\mathbf{x},\mathbf{y}) = \max_{1 \leq i \leq n} |x_i - y_i|.$$

In many applications, the data space is Euclidean and therefore the L_2 metric is the natural choice. In addition, this metric has the advantage of being isotropic (rotation invariant). For example, when the input vectors stem from an isotropic vector field, e.g. a velocity field, the most appropriate choice is to use the L_2 metric so that all vectors are processed in the same way, regardless of their orientation [3].

The main drawback of L_2 is its high computational requirements due to the multiplications and the square root operation. As a result, L_1 and L_∞ are often used as alternatives. Although these metrics are computationally more efficient, they deviate

from L_2 significantly. The Minkowski metric is translation invariant, i.e. $L_p(\mathbf{x},\mathbf{y}) = L_p(\mathbf{x}+\mathbf{z},\mathbf{y}+\mathbf{z})$ for all $\mathbf{x},\mathbf{y},\mathbf{z} \in \mathbb{R}^n$, hence it suffices to consider $D_p(\mathbf{x}) = L_p(\mathbf{x},\mathbf{0})$, i.e. the distance from the point \mathbf{x} to the origin. Therefore, in the rest of the paper, we will consider approximations to $D_p(\mathbf{x})$ rather than $L_p(\mathbf{x},\mathbf{y})$.

In this paper, we examine several approximations to the Euclidean norm. The rest of the paper is organized as follows. In Section 2 we describe the Euclidean norm approximations that have appeared in the literature, and compare their average and maximum errors using numerical simulations. We then show that all of these methods fit into a single mathematical formulation. In Section 3 we examine the simulation results from a theoretical perspective. Finally, in Section 4 we provide our conclusions.

2. Euclidean norm approximations

For reasons explained in Section 1, we concentrate on approximations to the Euclidean norm D_2 on \mathbb{R}^n . Let \tilde{D} , defined on \mathbb{R}^n , be an approximation to D_2 . We assume that \tilde{D} is a continuous homogeneous function. We note that all variants of \tilde{D} we consider in this paper satisfy these assumptions. As a measure of the quality of the approximation of \tilde{D} to D_2 we define the maximum relative error (MRE) as

$$\varepsilon_{\max}^{\tilde{D}} = \sup_{\mathbf{x} \in \mathbb{R}^n \setminus \{\mathbf{0}\}} \frac{|\tilde{D}(\mathbf{x}) - D_2(\mathbf{x})|}{D_2(\mathbf{x})}. \quad (2)$$

Using the homogeneity of D_2 and \tilde{D} , (2) can be written as

$$\varepsilon_{\max}^{\tilde{D}} = \sup_{\mathbf{x} \in S_2^{n-1}} |\tilde{D}(\mathbf{x}) - 1|, \quad (3)$$

* Corresponding author.

E-mail addresses: ecelebi@lsus.edu (M.E. Celebi), celiker@math.wayne.edu (F. Celiker), kingravi@gatech.edu (H.A. Kingravi).

where

$$S_2^{n-1} = \{\mathbf{x} \in \mathbb{R}^n : D_2(\mathbf{x}) = 1\}$$

is the unit hypersphere of \mathbb{R}^n with respect to the Euclidean norm. Furthermore, by the continuity of \tilde{D} , we can replace the supremum with maximum in (3) and write

$$e_{\max}^{\tilde{D}} = \max_{\mathbf{x} \in S_2^{n-1}} |\tilde{D}(\mathbf{x}) - 1|. \tag{4}$$

We will use (4) as the definition of MRE throughout.

In the trivial case where $\tilde{D} = D_2$ we have $e_{\max}^{\tilde{D}} = 0$. Hence, for nontrivial cases we wish to have a small $e_{\max}^{\tilde{D}}$ value. In other words, the smaller the value of $e_{\max}^{\tilde{D}}$, the better (more accurate) the corresponding approximation \tilde{D} . It can be shown that D_1 (city-block norm) overestimates D_2 and the corresponding MRE is given by $e_{\max}^{D_1} = \sqrt{n} - 1$ [1]. In contrast, D_∞ (chessboard norm) underestimates D_2 with MRE given by $e_{\max}^{D_\infty} = 1 - 1/\sqrt{n}$ [1]. More explicitly,

$$D_2(\mathbf{x}) \leq D_1(\mathbf{x}) \leq \sqrt{n}D_2(\mathbf{x}),$$

$$(1/\sqrt{n})D_2(\mathbf{x}) \leq D_\infty(\mathbf{x}) \leq D_2(\mathbf{x}) \tag{5}$$

for all $\mathbf{x} \in \mathbb{R}^n$. Therefore, it is natural to expect a suitable linear combination of D_1 and D_∞ to give an approximation to D_2 better than both D_1 and D_∞ [2].

2.1. Chaudhuri et al.'s approximation

Chaudhuri et al. [1] proposed the approximation¹

$$D_\lambda(\mathbf{x}) = |x_{i_{\max}}| + \lambda \sum_{\substack{i=1 \\ i \neq i_{\max}}}^n |x_i| \quad \text{with } \lambda = \frac{1}{n - \lfloor \frac{n-2}{2} \rfloor}.$$

Here i_{\max} is the index of the absolute largest component of \mathbf{x} , i.e. $i_{\max} = \operatorname{argmax}_{1 \leq i \leq n} (|x_i|)$, and $\lfloor x \rfloor$ is the floor function which returns the largest integer less than or equal to x . Since $D_\infty(\mathbf{x}) = |x_{i_{\max}}|$, by adding and subtracting the term $\lambda|x_{i_{\max}}|$, D_λ can be written as a linear combination of D_∞ and D_1 as

$$D_\lambda(\mathbf{x}) = (1-\lambda)D_\infty(\mathbf{x}) + \lambda D_1(\mathbf{x}). \tag{6}$$

It is easy to see that $D_\infty(\mathbf{x}) \leq D_\lambda(\mathbf{x}) \leq D_1(\mathbf{x})$ for all $\mathbf{x} \in \mathbb{R}^n$ since $0 < \lambda \leq 0.5$. It can also be shown that [1] for sufficiently large n , D_λ is closer to D_2 than both D_1 and D_∞ , i.e. $|D_\lambda(\mathbf{x}) - D_2(\mathbf{x})| \leq |D_1(\mathbf{x}) - D_2(\mathbf{x})|$ and $|D_\lambda(\mathbf{x}) - D_2(\mathbf{x})| \leq |D_2(\mathbf{x}) - D_\infty(\mathbf{x})|$ for all $\mathbf{x} \in \mathbb{R}^n$.

For sufficiently large n , D_λ underestimates D_2 and the corresponding MRE is given by

$$1 - \frac{1-\lambda(n-1)}{\sqrt{n}} \leq e_{\max}^{D_\lambda} \leq 1-\lambda \quad \text{for } n \geq 3.$$

Otherwise, D_λ overestimates D_2 and we have

$$e_{\max}^{D_\lambda} = \begin{cases} \sqrt{1 + \frac{4(n-1)}{(n+2)^2}} - 1 & \text{for } n = 2, 4, 6, \dots \\ \sqrt{1 + \frac{4(n-1)}{(n+3)^2}} - 1 & \text{for } n = 3, 5, 7, \dots \end{cases}$$

Proofs of these identities can be found in [1].

¹ Unfortunately the motivation behind this particular choice of λ is not given in the paper.

2.2. Rhodes' approximations

Rhodes [2] reformulated (6) as a maximum of linear functions

$$D_\lambda(\mathbf{x}) = \max_{1 \leq j \leq n} \left\{ (1-\lambda)|x_j| + \lambda \sum_{i=1}^n |x_i| \right\},$$

where $0 < \lambda < 1$ is taken as a free parameter. He determined the optimal value for λ by minimizing $e_{\max}^{D_\lambda} = \max_{\mathbf{x} \in S_2^{n-1}} |D_\lambda(\mathbf{x}) - 1|$ analytically. In particular, he showed that optimal λ' values for Chaudhuri et al.'s norm can be determined by solving the equation $1 - 2\sqrt{\lambda' - (\lambda')^2} = \sqrt{1 + (\lambda')^2(n-1)} - 1$ in the interval $(0, 1/2)$. This equation is a quartic (fourth order) in λ and can be solved using Ferrari's method [5]. It can be shown that this particular quartic equation has two real and two complex roots and the optimal λ' value is given by the smaller of the real roots. The corresponding MRE is given by [2]

$$e_{\max}^{D_\lambda} = 1 - 2\sqrt{\lambda' - (\lambda')^2}. \tag{7}$$

In the remainder of this paper, D_λ refers to this improved variant of Chaudhuri et al.'s norm.

Rhodes also investigated the two-parameter family of approximations given by

$$D_{\mu,\lambda}(\mathbf{x}) = (\mu-\lambda)D_\infty(\mathbf{x}) + \lambda D_1(\mathbf{x}), \tag{8}$$

where $0 < \lambda < \mu$. He proved that the optimal solution and its MRE in this case are given by

$$\lambda^* = \frac{2}{2n^{1/4} + \sqrt{2n+2\sqrt{n}}}, \quad \mu^* = (\sqrt{n}+1)\lambda^*,$$

$$e_{\max}^{D_{\mu,\lambda}} = 1 - 2\lambda^* n^{1/4}. \tag{9}$$

Finally, Rhodes investigated the $D_{\mu,\lambda}$ approximations with $0 \leq \mu < \lambda$. He proved that the optimal solution and its MRE are given by

$$\lambda^* = \frac{2}{1 + \sqrt{n-1}}, \quad \mu^* = 0, \quad e_{\max}^{D_{\mu,\lambda}} = 1 - \lambda^*.$$

This approximation will not be considered any further since its accuracy is inferior to even the single-parameter approximation D_λ .

It should be noted that Rhodes optimized D_λ and $D_{\mu,\lambda}$ over \mathbb{Z}^n . Therefore, these norms are in fact suboptimal on \mathbb{R}^n (see Section 2.5).

2.3. Barni et al.'s approximation

Barni et al. [3,6] formulated a generic approximation for D_2 as

$$D_B(\mathbf{x}) = \delta \sum_{i=1}^n \alpha_i x_{(i)},$$

where $x_{(i)}$ is the i -th absolute largest component of \mathbf{x} , i.e. $(x_{(1)}, x_{(2)}, \dots, x_{(n)})$ is a permutation of $(|x_1|, |x_2|, \dots, |x_n|)$ such that $x_{(1)} \geq x_{(2)} \geq \dots \geq x_{(n)}$. Here $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_n)$ and $\delta > 0$ are approximation parameters. Note that a non-increasing ordering and strict positivity of the component weights, i.e. $\alpha_1 \geq \alpha_2 \geq \dots \geq \alpha_n > 0$ is a necessary and sufficient condition for D_B to define a norm [6].

The minimization of (4) is equivalent to determining the weight vector α and the scale factor δ that solve the following minimax problem [6,7]:

$$\min_{\alpha, \delta} \max_{\mathbf{x} \in V} |D_B(\mathbf{x}) - 1|, \tag{10}$$

where $V = \{\mathbf{x} \in \mathbb{R}^n : x_1 \geq x_2 \geq \dots \geq x_n \geq 0, D_2(\mathbf{x}) = 1\}$.

The optimal solution and its MRE are given by

$$\alpha_i^* = \sqrt{i} - \sqrt{i-1}, \quad \delta^* = \frac{2}{1 + \sqrt{\sum_{i=1}^n \alpha_i^{*2}}}, \quad \varepsilon_{\max}^{D_B} = 1 - \delta^*. \quad (11)$$

It should be noted that a similar but less rigorous approach had been published earlier by Ohashi [8].

2.4. Seol and Cheun's approximation

Seol and Cheun [4] recently proposed an approximation of the form

$$D_{a,b}(\mathbf{x}) = aD_\infty(\mathbf{x}) + bD_1(\mathbf{x}), \quad (12)$$

where a and b are strictly positive parameters to be determined by solving the following 2×2 linear system:

$$aE(D_\infty^2) + bE(D_\infty D_1) = E(D_2 D_\infty),$$

$$aE(D_\infty D_1) + bE(D_1^2) = E(D_2 D_1),$$

where $E(\cdot)$ is the expectation operator.

Note that the formulation of $D_{a,b}$ is similar to that of $D_{\mu,\lambda}$ (8) in that they both approximate D_2 by a linear combination of D_∞ and D_1 . These approximations differ in their methodologies for finding the optimal parameters. Rhodes follows an analytical approach and derives theoretical values for the parameters and the maximum error. However, he achieves this by sacrificing maximization over \mathbb{R}^n , and maximizes only over \mathbb{Z}^n . Seol and Cheun follow an empirical approach where they approximate optimal parameters over \mathbb{R}^n , which causes them to sacrifice the ability to obtain analytical values for the parameters and the maximum error. They estimate the optimal values of a and b using 100,000 n -dimensional vectors whose components are independent and identically distributed, standard Gaussian random variables.

2.5. Comparison of the Euclidean norm approximations

It is easy to see that all of the presented approximations fit into the general form

$$\tilde{D}(\mathbf{x}) = \sum_{i=1}^n w_i x_{(i)},$$

which is a weighted D_1 norm.

The component weights for each approximation are given in Table 1. It can be seen that D_B has the most elaborate design in which each component is assigned a weight proportional to its ranking. However, this weighting scheme also presents a drawback in that a full ordering of the component absolute values is required (see Table 2).

Due to their formulations, the MREs for D_λ , $D_{\mu,\lambda}$, and D_B can be calculated analytically using (7), (9), and (11), respectively. In Fig. 1 we plot the theoretical errors for these norms for $n \leq 100$. It can be seen that D_B is not only more accurate than D_λ and $D_{\mu,\lambda}$, but it also scales significantly better. Although $D_{\mu,\lambda}$ is

Table 1
Weights for the approximate norms.

Norm	w_1	$w_{i \neq 1}$
D_λ	1	λ'
$D_{\mu,\lambda}$	μ^*	λ^*
D_B	δ^*	$\delta^* \alpha_i^*$
$D_{a,b}$	$a+b$	b

Table 2
Operation counts for the norms.

Norm	ABS	COMP	ADD	MULT	SQRT
D_∞	n	$n-1$	0	0	0
D_1	n	0	$n-1$	0	0
D_2	0	0	$n-1$	n	1
D_λ	n	$n-1$	$n-1$	1	0
$D_{\mu,\lambda}$	n	$n-1$	n	2	0
D_B	n	$\mathcal{O}(n \log n)$	$n-1$	n	0
$D_{a,b}$	n	$n-1$	n	2	0

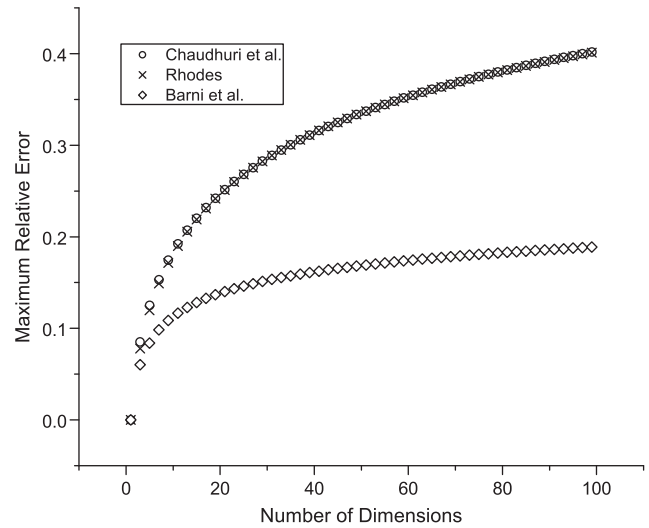


Fig. 1. Maximum relative errors for D_λ , $D_{\mu,\lambda}$, and D_B .

more accurate than D_λ when n is small, the difference between the two approximations becomes less significant as n is increased.

The operation counts for each norm are given in Table 2 (ABS: absolute value, COMP: comparison, ADD: addition, MULT: multiplication, SQRT: square root). The following conclusions can be drawn:

- D_B has the highest computational cost among the approximate norms due to its costly weighting scheme, which requires sorting of n numbers and n multiplications. For small values of n , sorting can be performed most efficiently by a sorting network [9]. For large values of n , sorting requires $\mathcal{O}(n \log n)$ comparisons, which is likely to exceed the cost of the square root operation [6]. Therefore, in high dimensional spaces, e.g. $n > 9$ [4], D_B provides no computational advantage over D_2 .
- D_λ has the lowest computational cost among the approximate norms. $D_{\mu,\lambda}$ and $D_{a,b}$ have the same computational cost, which is slightly higher than that of D_λ .
- A significant advantage of D_λ , $D_{\mu,\lambda}$, and $D_{a,b}$ is that they require a fixed number of multiplications (1 or 2) regardless of the value of n .
- D_λ , $D_{\mu,\lambda}$, and $D_{a,b}$ can be used to approximate D_2^2 (squared Euclidean norm) using an extra multiplication. On the other hand, the computational cost of D_B is higher than that of D_2^2 due to the extra absolute value and sorting operations involved.

In Table 3 we display the average and maximum errors for D_λ , $D_{\mu,\lambda}$, and D_B for $n \leq 10$. Average relative error (ARE) is defined as

$$\varepsilon_{\text{avg}}^{\tilde{D}} = \frac{1}{|S|} \sum_{\mathbf{x} \in S} |\tilde{D}(\mathbf{x}) - 1|, \quad (13)$$

Table 3
Average and maximum errors for D_λ , $D_{\mu,\lambda}$, and D_B .

n	D_λ			$D_{\mu,\lambda}$			D_B		
	ARE	MRE _e	MRE _t	ARE	MRE _e	MRE _t	ARE	MRE _e	MRE _t
2	0.0348	0.0551	0.0551	0.0276	0.0470	0.0470	0.0241	0.0396	0.0396
3	0.0431	0.0852	0.0852	0.0367	0.0778	0.0778	0.0300	0.0602	0.0602
4	0.0455	0.1074	0.1074	0.0420	0.1010	0.1010	0.0345	0.0739	0.0739
5	0.0460	0.1251	0.1251	0.0447	0.1197	0.1197	0.0377	0.0839	0.0839
6	0.0458	0.1400	0.1400	0.0462	0.1354	0.1354	0.0401	0.0919	0.0919
7	0.0454	0.1529	0.1529	0.0469	0.1489	0.1490	0.0418	0.0984	0.0984
8	0.0448	0.1641	0.1643	0.0471	0.1606	0.1609	0.0431	0.1039	0.1039
9	0.0442	0.1739	0.1745	0.0471	0.1709	0.1716	0.0440	0.1086	0.1086
10	0.0435	0.1827	0.1837	0.0469	0.1803	0.1812	0.0447	0.1128	0.1128

Table 4
Average and maximum errors for $D_{a,b}$.

n	Seol and Cheun		This study	
	ARE	MRE _e	ARE	MRE _e
2	0.0200	0.0526	0.0200	0.0525
3	0.0239	0.0991	0.0239	0.0998
4	0.0257	0.1342	0.0257	0.1363
5	0.0268	0.1420	0.0268	0.1649
6	0.0273	0.1674	0.0273	0.1871
7	0.0276	0.1772	0.0276	0.1968
8	0.0277	0.1753	0.0277	0.2076
9	0.0277	0.1711	0.0277	0.2120
10	0.0276	0.1526	0.0276	0.2156

where S is a finite subset of the unit hypersphere S_2^{n-1} , and $|S|$ denotes the number of elements in S . An efficient way to pick a random point on S_2^{n-1} is to generate n independent Gaussian random variables x_1, x_2, \dots, x_n with zero mean and unit variance. The distribution of the unit vectors

$$\left\{ \mathbf{y} = (y_1, y_2, \dots, y_n) : y_i = x_i / \left(\sum_{i=1}^n x_i^2 \right)^{1/2}, i = 1, 2, \dots, n \right\}$$

will then be uniform over the surface of the hypersphere [10]. For each approximate norm, the ARE and MRE values were calculated over an increasing number of points, $2^{20}, 2^{21}, \dots, 2^{32} - 1$ (that are uniformly distributed on the hypersphere) until the error values converge, i.e. the error values do not differ by more than $\epsilon = 10^{-5}$ in two consecutive iterations. Note that for each norm, two types of maximum error were considered: empirical maximum error (MRE_e), which is calculated numerically over S and the theoretical maximum error (MRE_t), which is calculated analytically using (7), (9), or (11). It can be seen that for D_B , the empirical and maximum errors agree in all cases, which demonstrates the validity of the presented iterative error calculation scheme. This is not the case for D_λ and $D_{\mu,\lambda}$ since these norms are optimized over \mathbb{Z}^n instead of \mathbb{R}^n . Therefore, a perfect agreement between the empirical and theoretical results should not be expected. Nevertheless, the empirical error is always less than the maximum error, which is expected because we are maximizing over a smaller set.

Table 4 shows the average and maximum errors for $D_{a,b}$. The error values under the column “Seol and Cheun” are taken from [4] (where the simulations were performed on a set of 100,000 n -dimensional vectors whose components are independent and identically distributed, zero mean, and unit variance Gaussian random variables), whereas those under the column “this study” were obtained using the aforementioned iterative scheme. It can be seen that the maximum errors obtained by Seol and Cheun are

lower than those that we obtained and the discrepancy between the outcomes of the two error calculation schemes increases as n is increased. The optimistic maximum error values given by Seol and Cheun are due to the fact that 100,000 vectors are not enough to cover the surface of the hypersphere in higher dimensions. This is investigated further in the following section. On the other hand, the average error values agree perfectly in both calculation schemes.

By examining Tables 3 and 4 the following observations can be made regarding the maximum error:

- D_B is the most accurate approximation in all cases. This is because this norm is designed to minimize the maximum error and it has a more sophisticated weighting scheme than the other two approximations, i.e. D_λ and $D_{\mu,\lambda}$, that are based on the same optimality criterion.
- As is also evident from Fig. 1, $D_{\mu,\lambda}$ is slightly more accurate than D_λ especially for small values of n , in accordance with the greater degrees of freedom it is afforded.
- $D_{a,b}$ is the least accurate approximation except for $n=2$. This was expected since this norm is designed to minimize the mean squared error rather than the maximum error.
- As n is increased, the error increases in all approximations. However, as can be seen from Fig. 1, the error grows faster in some approximations than others.

On the other hand, with respect to average error we can see that:

- As expected, $D_{a,b}$ is the most accurate approximation.
- As n is increased, the error increases consistently for the D_B norm. This is not the case for the D_λ and $D_{\mu,\lambda}$ norms. This inconsistent average error behavior is not surprising given the fact that these norms are designed to minimize the maximum error.
- Interestingly, D_λ is more accurate than $D_{\mu,\lambda}$ for $n > 5$. A possible explanation to this phenomenon is that both approximations are optimized for the maximum error. Since the minimization of the maximum and average errors are conflicting objectives, it is likely that $D_{\mu,\lambda}$ sacrifices the average error to obtain better (lower) maximum error. The same relationship holds between $D_{a,b}$ and D_B .

3. Sampling on the unit hypersphere

In this section, we demonstrate why a fixed number of samples from the unit hypersphere (i.e. the approach advocated in [4]) can give biased estimates for the maximum error. The basic reason behind this is the fact that a fixed number of samples fail to suffice as the dimension of the space increases. The following provides a

plausibility argument as to why this is the case. To this end, we need to consider the notion of covering a sphere “sufficiently”. We begin with some definitions.

A closed n -ball of radius r with respect to the Euclidean norm, denoted $B_2^n(r)$, is the set of points whose Euclidean norm is less than or equal to r . That is,

$$B_2^n(r) = \{\mathbf{x} \in \mathbb{R}^n : D_2(\mathbf{x}) \leq r\}.$$

Note that, in particular, the unit hypersphere S_2^{n-1} of \mathbb{R}^n is the boundary of $B_2^n(1)$.

Given an $\varepsilon > 0$, we say that a set \mathcal{C} of points on S_2^{n-1} is an ε -dense covering of S_2^{n-1} if for any \mathbf{x} in \mathcal{C} , there exists at least one $\hat{\mathbf{x}}$ in \mathcal{C} (different than \mathbf{x}) such that $D_2(\mathbf{x}-\hat{\mathbf{x}}) < \varepsilon$. Essentially, our main purpose here is to give a rough estimate of the number of points in \mathcal{C} , where \mathcal{C} is an ε -dense covering of S_2^{n-1} . We would then argue that, if ε is sufficiently small then \mathcal{C} is a fine-enough representation of points on S_2^{n-1} . Therefore, we can restrict any computation that needs to be performed on S_2^{n-1} to the finite set \mathcal{C} .

The basic idea behind the proof is to approximate S_2^{n-1} by $B_2^{n-1}(\varepsilon)$, that is, approximate the unit hypersphere of \mathbb{R}^n by $(n-1)$ -balls of radius ε . This is the same principle as approximating a circle (S_2^1) in \mathbb{R}^2 by tiny line segments ($B_2^1(\varepsilon)$), or the surface of a sphere (S_2^2) in \mathbb{R}^3 by tiny discs ($B_2^2(\varepsilon)$). It is easy to see that, if we choose ε small enough, then the approximation is satisfactory for most practical purposes.

To proceed further, we need a lemma from elementary probability theory, which is known as the coupon collector’s problem [11].

Lemma 1. Given a collection of c distinct objects, the expected number of independent random trials needed to sample each one of the c objects is $\mathcal{O}(c \log c)$.

We can now prove the following result.

Theorem 1. The expected number of uniformly distributed samples needed to generate an ε -dense covering of S_2^{n-1} is $\mathcal{O}(N \log N)$ where $N = n/\varepsilon^{n-1}$.

Proof. Let $\varepsilon > 0$ be given. We will first count the number of identical copies of $B_2^{n-1}(\varepsilon)$ -balls that are needed to approximate S_2^{n-1} in the sense described above. By elementary calculus one can compute the volume of an $B_2^n(r)$ to be

$$V_n(r) := \frac{\pi^{n/2}}{\Gamma(\frac{n}{2}+1)} r^n := C_n r^n,$$

where Γ is the gamma function. The surface area of this ball is equal to the derivative with respect to r of its volume:

$$A_{n-1}(r) = \frac{d}{dr} V_n(r) = n C_n r^{n-1}.$$

Note that $A_{n-1}(1)$ is equal to the surface area of S_2^{n-1} . The approximate number of $B_2^{n-1}(\varepsilon)$ -balls needed to cover the surface of S_2^{n-1} is the ratio of the surface area of S_2^{n-1} to the volume of $B_2^{n-1}(\varepsilon)$, i.e.,

$$\begin{aligned} \frac{A_{n-1}(1)}{V_{n-1}(\varepsilon)} &= \frac{n C_n}{C_{n-1} \varepsilon^{n-1}} = \frac{n \pi^{n/2}}{\pi^{(n-1)/2}} \frac{\Gamma(\frac{n-1}{2}+1)}{\Gamma(\frac{n}{2}+1)} \frac{1}{\varepsilon^{n-1}} \\ &\approx \frac{n \pi^{1/2}}{\varepsilon^{n-1}} = \mathcal{O}\left(\frac{n}{\varepsilon^{n-1}}\right). \end{aligned}$$

The result now follows once we apply Lemma 1 with $c = n/\varepsilon^{n-1}$. \square

In the light of the following result, we see that the actual number of samples required does not deviate significantly from the value provided by Theorem 1.

Theorem 2. Let X be the number of samples observed before obtaining one in each region. Then, for any constant $s > 0$ we have $P(X > c \ln c + sc) < e^{-s}$.

Proof. The probability of not obtaining the i th region after $c \ln c + sc$ steps is

$$\left(1 - \frac{1}{c}\right)^{c(\ln c + s)} < e^{-(\ln c + s)} = \frac{1}{e^s c}.$$

By a union bound, the probability that a region has not been obtained after $c \ln c + sc$ steps is only e^{-s} . \square

Note that one can use a Chernoff bound to obtain an even tighter bound in Theorem 2 since

$$\lim_{c \rightarrow \infty} P(X > c \ln c + sc) = 1 - e^{-e^{-s}}.$$

See [11] for details.

We should note that in order to apply Lemma 1 the patches used to cover S_2^{n-1} should be disjoint which is clearly not the case since we have used $B_2^{n-1}(\varepsilon)$ for this purpose. This leads to an overestimate of the samples needed to obtain a dense covering, and thus the argument presented in this section is only a rough estimate. However, as empirically demonstrated in the previous section, a fixed number of samples as in [4] is definitely not sufficient either. To come up with a tight estimate of the number of sample points needed, one has to express S_2^{n-1} as a disjoint union of small patches. The delicacy lies in the requirement that this has to be achieved through a constructive process in a way that the surface area of each patch can be explicitly computed as a function of the dimension n , and a characteristic measure ε . To the best of the authors’ knowledge there is no systematic method in the literature to achieve this.

4. Conclusions

In this paper, we investigated the theoretical and practical aspects of several Euclidean norm approximations in the literature and showed that these are in fact special cases of the weighted city-block norm. We evaluated the average and maximum errors of these norms using numerical simulations. Finally, we demonstrated that the maximum errors given in a recent study [4] are significantly optimistic.

The implementations of the approximate norms described in this paper will be made publicly available at <http://www.lsus.edu/faculty/~ecelebi/research.htm>.

Acknowledgments

This work was supported by a grant from the Louisiana Board of Regents (LEQSF2008-11-RD-A-12). The authors are grateful to Changkyu Seol for clarifying various points about his paper.

References

- [1] D. Chaudhuri, C.A. Murthy, B.B. Chaudhuri, A modified metric to compute distance, Pattern Recognition 25 (7) (1992) 667–677.
- [2] F. Rhodes, On the metrics of Chaudhuri, Murthy and Chaudhuri, Pattern Recognition 28 (5) (1995) 745–752.
- [3] M. Barni, F. Bartolini, F. Buti, V. Cappellini, Optimum linear approximation of the Euclidean norm to speed up vector median filtering, in: Proceedings of the

- 2nd IEEE International Conference on Image Processing (ICIP'95), 1995, pp. 362–365.
- [4] C. Seol, K. Cheun, A low complexity Euclidean norm approximation, *IEEE Transactions on Signal Processing* 56 (4) (2008) 1721–1726.
- [5] R.B. King, *Beyond the Quartic Equation*, Birkhäuser, Boston, 2008.
- [6] M. Barni, F. Buti, F. Bartolini, V. Cappellini, A quasi-Euclidean norm to speed up vector median filtering, *IEEE Transactions on Image Processing* 9 (10) (2000) 1704–1709.
- [7] V.F. Dem'yanov, V.N. Malozemov, *Introduction to Minimax*, Dover Publications, 1990.
- [8] Y. Ohashi, Fast linear approximations of Euclidean distance in higher dimensions, in: P. Heckbert (Ed.), *Graphics Gems IV*, Academic Press, 1994.
- [9] T.H. Cormen, C.E. Leiserson, R.L. Rivest, C. Stein, *Introduction to Algorithms*, The MIT Press, 2009.
- [10] M.E. Muller, A note on a method for generating points uniformly on N-dimensional spheres, *Communications of the ACM* 2 (4) (1959) 19–20.
- [11] M. Mitzenmacher, E. Upfal, *Probability and Computing: Randomized Algorithms and Probabilistic Analysis*, Cambridge University Press, 2005.