

Improving Efficiency of Density-Based Shape Descriptors for 3D Object Retrieval

Ceyhun Burak Akgül^{1,2}, Bülent Sankur¹, Yücel Yemez³, and Francis Schmitt²

¹ Department of Electrical and Electronics Engineering, Boğaziçi University, Istanbul, Turkey

² GET - Télécom Paris - CNRS UMR 5141, Paris, France

³ Department of Computer Engineering, Koç University, Istanbul, Turkey

Abstract. We consider 3D shape description as a probability modeling problem. The local surface properties are first measured via various features, and then the probability density function (pdf) of the multidimensional feature vector becomes the shape descriptor. Our prior work has shown that, for 3D object retrieval, pdf-based schemes can provide descriptors that are computationally efficient and performance-wise on a par with or better than the state-of-the-art methods. In this paper, we specifically focus on discretization problems in the multidimensional feature space, selection of density evaluation points and dimensionality reduction techniques to further improve the performance of our density-based descriptors.

1 Introduction

A 3D shape descriptor can be viewed as a mapping from the space of 3D objects to some finite-dimensional vector space. The design objective of 3D shape descriptors is to maximize the shape discrimination ability with a low-dimensional and sparse feature set. Thus, while descriptors should enable effective retrieval, they should also satisfy practical requirements of computation and storage efficiency [1].

Density-based shape description is an analytical framework to extract 3D shape descriptors from local surface features characterizing the object geometry [2]. Processing the feature information with the kernel methodology for density estimation (KDE) [3], the probability density function (pdf) of the local feature is estimated at chosen target points. The shape descriptor vector is then simply a sampled version of this pdf. This density-based approach provides a mechanism to convert local shape evidences, using KDE, into a global shape description. Our recent work on density-based shape descriptors [2,4,5] for 3D object retrieval has proven that this scheme is both computationally rapid and performance-wise on a par with or better than the state-of-the-art methods [1,6,7,8].

In the present work, we depart from the pdf-based shape description scheme and investigate variations to improve its efficiency and effectiveness. In particular, we address the following issues: (i) the effect of the descriptor size;

(ii) the choice of the dissimilarity measure associated with the descriptor; (iii) the descriptor dimensionality reduction techniques.

The paper is organized as follows. In Sect. 2, we define the local features used in density-based shape description and present the methods to determine the feature range for target selection. In Sect. 3, we discuss issues related to 3D object retrieval and introduce dimensionality reduction techniques specifically for density-based descriptors: marginalization and probability density pruning. Experimental results are presented in Sect. 4 and conclusions drawn in Sect. 5.

2 Density-Based Shape Description: Target Selection Problem

A density-based descriptor of a 3D shape is defined as the sampled pdf of some surface feature, such as radial distance or direction. The feature is local to the surface patch and treated as a random variable. At each surface point, e.g., at each mesh vertex and/or triangle, one has a realization of this random variable. To fix the notation, let S be a random variable taking values within a subspace \mathcal{R}_S of \mathbb{R}^d and let $f(s|O_t)$ be the pdf of S evaluated on the surface of the object O_t . In the sequel, random variables appear as uppercase letters while their specific instances as lowercase. Suppose furthermore that we have specified a finite set of points $\overline{\mathcal{R}}_S = \{s_n \in \mathcal{R}_S : n = 1, \dots, N\}$, called *the target set*, within \mathcal{R}_S . The density-based descriptor $\mathbf{f}_{S|O_t}$ for the object O_t (with respect to the feature S) is then simply an N -dimensional vector whose entries consist of the pdf samples at the target set, that is, $\mathbf{f}_{S|O_t} = [f(s_1|O_t), \dots, f(s_N|O_t)]$.

Density-based shape description consists of three stages. First, we should choose a good local feature that accumulates to a global shape descriptor. We refer to this as the design problem. Second, we need an efficient computational scheme to estimate $f(s|O_t)$ at arbitrary points in the range \mathcal{R}_S . This issue is the generic computational problem that can be tackled in different ways [3]. In [2,4,5], we have used the kernel approach in conjunction with a fast algorithm, the fast Gauss transform [9]. Finally, in the target selection stage, we determine the sampling locations (set $\overline{\mathcal{R}}_S$) within \mathcal{R}_S , the final output being the object descriptor $\mathbf{f}_{S|O_t}$.

2.1 Local Surface Features

We briefly present two of the local features as a reminder to facilitate the subsequent discussions. These features were already analyzed in [2,4]. In this work, 3D shapes are represented by triangular meshes and their centers of mass coincide with the origin. Let $\mathbf{Q} = (X, Y, Z)$ stand for a point Q , lying on some 3D surface, $\hat{\mathbf{N}} = (\hat{N}_x, \hat{N}_y, \hat{N}_z)$ for the unit surface normal vector at Q , and $\langle \cdot, \cdot \rangle$ for the usual dot product. A local geometric feature is a mapping S from the points of the surface into a d -dimensional space, usually constrained into a finite subspace \mathcal{R}_S of \mathbb{R}^d . Each dimension of this space corresponds to a specific geometric measure, characterizing the shape locally. We now consider two of these local geometric features.




Fig. 1. The components of radial and tangent plane based features at a surface point Q

The radial feature S_{rad} at the point Q is a 4-tuple defined as

$$S_{rad}(Q) \triangleq (R, \hat{R}_x, \hat{R}_y, \hat{R}_z) \triangleq (R, \hat{\mathbf{R}}) \text{ with } R \triangleq \|\hat{\mathbf{Q}}\| \text{ and } \hat{\mathbf{R}} \triangleq \mathbf{Q} / \|\hat{\mathbf{Q}}\|.$$

The *magnitude* component R measures the radial distance from the origin. The *direction* component $\hat{\mathbf{R}}$, pointing to the location of the surface point, is a 3-vector with unit-norm and lies on the unit 2-sphere \mathcal{S}^2 . Since the magnitude component is spread between zero and some r_{max} depending on the scale of the surface, S_{rad} takes values within the Cartesian product of $\mathcal{I}_R = (0, r_{max}) \subset \mathbb{R}$ and \mathcal{S}^2 , thus $\mathcal{R}_{S_{rad}} = \mathcal{I}_R \times \mathcal{S}^2$.

The tangent plane-based feature S_{tan} at the point Q is a 4-tuple defined as

$$S_{tan}(Q) \triangleq (D_{tan}, \hat{N}_x, \hat{N}_y, \hat{N}_z) \triangleq (D_{tan}, \hat{\mathbf{N}}) \text{ with } D_{tan} \triangleq \left| \langle \mathbf{Q}, \hat{\mathbf{N}} \rangle \right|.$$

Similar to the S_{rad} -feature, S_{tan} has a *magnitude* component D_{tan} taking values within $\mathcal{I}_{D_{tan}} = (0, d_{tan,max}) \subset \mathbb{R}$, which is the distance of the tangent plane to the origin, and a *direction* component $\hat{\mathbf{N}}$. The normal $\hat{\mathbf{N}}$ is a unit-norm vector by definition and lies on the unit 2-sphere \mathcal{S}^2 . Consequently, the range of S_{tan} is $\mathcal{R}_{S_{tan}} = \mathcal{I}_{D_{tan}} \times \mathcal{S}^2$. Fig. 1 illustrates these two local features. Note that the above features are neither scale- nor rotation-invariant. Accordingly, any method making use of them must assume prior scale and pose normalization of the mesh.

2.2 Target Selection

The target selection problem was defined as sampling the range of the feature at which the pdf is evaluated. The ranges $\mathcal{R}_{S_{rad}}$ and $\mathcal{R}_{S_{tan}}$ of our local features

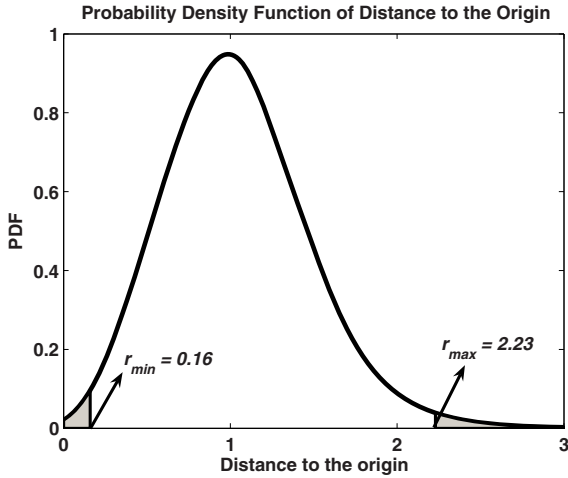


Fig. 2. Probability density function of R displayed with the lower and upper 1% intervals clipped

S_{rad} and S_{tan} both consist of the Cartesian product of an interval \mathcal{I} on the real line with the unit 2-sphere \mathcal{S}^2 . This decoupling allows us to sample \mathcal{I} and \mathcal{S}^2 independently into N_m and N_d points respectively, and to take the Cartesian product of the two resulting sets to yield the target set $\overline{\mathcal{R}}_S$ of size $N = N_m \times N_d$. In the following, we discuss the sampling of \mathcal{I} and \mathcal{S}^2 separately, providing two different methods for each.

Sampling the Magnitude Interval: For features that are not scale-invariant, 3D objects must be normalized. In [2,4,5], we rescaled all objects to unit mean radial distance. Vranic [7] has experimentally shown that this constitutes a good choice for retrieval. Other scale normalizers in the literature use scale measures ranging from the radius of the enclosing sphere to the mean coordinate variance [7,8]. The support of the mean-normalized radial distance, $\mathcal{I} = (r_{min}, r_{max})$, is determined by the 98-percentile of the distribution of R . Thus, we set the limits to satisfy $Pr\{R < r_{max}\} = \int_0^{r_{max}} f(r)dr = 0.99$ and $Pr\{r_{min} < R\} = 0.99$, where $f(r)$ is the pdf of R , calculated over all objects. More specifically, one has

$$f(r) = \sum_{t=1}^T f(r, O_t) = \sum_{t=1}^T f(r|O_t)Pr\{O_t\} = \frac{1}{T} \sum_{t=1}^T f(r|O_t), \quad (1)$$

assuming that all objects are equiprobable. In Fig. 2, we display the pdf of R and the clipping points delimiting the 98-percentile for the Princeton Shape Benchmark (PSB) [6]. The same method also applies for the tangent feature interval $\mathcal{I}_{D_{tan}}$.

Once the domain of the pdf is set, the N_m targets points remain to be determined. The two methods we consider are to either partition the interval into N_m equally spaced (uniform) regions and take the midpoints, or to partition the interval into N_m equal probability regions.

Sampling the Unit 2-Sphere: For the direction components of both S_{rad} and S_{tan} , we sample the unit 2-sphere \mathcal{S}^2 to obtain the targets of their pdfs. We experiment again with two different sampling methods. In the first one [10], we consider an octahedron circumscribed by the unit sphere, subdivide in four each of its eight triangles, radially project the new triangles on the unit sphere, and iterate a factor of a times the subdivision process. The barycenters of the resulting triangles become the target set for direction components. This leads to a uniform partitioning of the sphere. The recursion factor a determines the number of resulting points N_d ; e.g., for $a = 1$, we get $N_d = 8 \times 4 = 32$ points; for $a = 2$, we get $N_d = 8 \times 16 = 128$, and in general $N_d = 2^{2a+3}$.

In the second method, we parameterize the unit sphere in terms of spherical coordinates. Recall that any point on the sphere can be expressed as a 3-tuple given by $(\sin\theta\sin\phi, \cos\theta\sin\phi, \cos\phi)$ where $0 \leq \theta < 2\pi$ and $0 \leq \phi < \pi$. Uniformly sampling the θ - and ϕ -coordinates at N_θ and N_ϕ points, respectively, results in $N_d = N_\theta \times N_\phi$ points on the unit-sphere. This method, however, does not provide a uniform partitioning of the sphere.

3 3D Object Retrieval: Improving Efficiency

3.1 Dissimilarity Measures

The classical dissimilarity measures between two pdfs, which are also our shape descriptors, are Minkowski metrics ($L^p, p = 1, 2, \infty$), symmetric Kullback-Leibler distance, Chi-Square distance and Bhattacharya distance [7,8]. For $\mathbf{f}_{S|O_t}$ and $\mathbf{f}_{S|O_u}$ representing discretized pdfs of some feature S on objects O_t and O_u , the continuous metric δ can be approximated as

$$\bar{\delta}(\mathbf{f}_{S|O_t}, \mathbf{f}_{S|O_u}) = \sum_{s_n \in \mathcal{R}_S} \eta(f(s_n|O_t), f(s_n|O_u)) \Delta s_n, \quad (2)$$

where η is a point-wise dissimilarity function, e.g., for L^1 , $\eta(\cdot, \cdot) = |\cdot - \cdot|$ and where Δs_n is the discretization step size. For uniform partitioning of \mathcal{R}_S , Δs_n is constant and can be dropped.

3.2 Dimensionality Reduction

We consider schemes to reduce the dimension of the descriptor vectors for computation and memory advantages, without sacrificing the retrieval performance. This is important because the descriptors may be quite high dimensional, in the order of thousands.

Marginalization: Features can be selectively removed from the multidimensional descriptor by marginalization, that is, by integrating out feature variables. To remove the component S_k from some d -dimensional feature $S = (S_1, S_2, \dots, S_d)$, we use

$$f_{S_{\setminus k}|O_t} \triangleq f(s_1, \dots, s_{k-1}, s_{k+1}, \dots, s_d|O_t) = \int_{S_k} f(s_1, \dots, s_k, \dots, s_d|O_t) ds_k. \quad (3)$$

This gives the pdf of a “reduced” feature vector $S_{\mathcal{Y}} \triangleq (S_1, \dots, S_{k-1}, S_{k+1}, \dots, S_d)$. Reducing the descriptor $f_{S|O_t}$ to $f_{S_{\mathcal{Y}}|O_t}$ saves us one dimension at the cost of any information brought by the component S_k . For instance, marginalizing the magnitude component R from the pdf of S_{rad} -feature vector, the size of the descriptor $\mathbf{f}_{S_{rad}, R|O_t}$ is N_m times smaller than that of $\mathbf{f}_{S_{rad}|O_t}$ since the target set for S_{rad} contains $N_m \times N_d$ points. We hope then to identify features that can be marginalized without compromising descriptor’s discrimination ability. An obvious instance is the case of a redundant component in the directional parts $\hat{\mathbf{R}}$ and $\hat{\mathbf{N}}$, respectively, of the local features S_{rad} and S_{tan} . For example, $\hat{\mathbf{R}}$ is unit norm with $\hat{R}_x^2 + \hat{R}_y^2 + \hat{R}_z^2 = 1$, hence given any two components, say \hat{R}_x and \hat{R}_y , the third one \hat{R}_z is completely determined up to the sign. Thus, we conjecture that \hat{R}_z can be marginalized out without deteriorating performance. We show this experimentally in Sect. 4.

Probability Density Pruning: The second approach to reduce a descriptor’s dimensionality involves pruning the unconditional pdf of features S , as defined in Eq. 1. The idea of pruning signifies removing negligible unconditional pdf bins from the descriptor. Thus, for a selected threshold λ , the new target set $\overline{\mathcal{R}}_S^\lambda$ is defined as $\overline{\mathcal{R}}_S^\lambda = \{s_n \in \overline{\mathcal{R}}_S : f(s_n) > \lambda\}$. The reduced descriptor for some object O_t becomes then $\mathbf{f}_{S|O_t}^\lambda = [f(s_n|O_t)]_{s_n \in \overline{\mathcal{R}}_S^\lambda}$. Notice that feature selection methods are not practical when the feature vector size is in the order of thousands. Pruning by suppressing small pdf values, albeit not tantamount to feature selection, still serves the goal by reducing the descriptor size.

4 Experiments

We have tested our descriptors in a retrieval scenario on the training set of the Princeton Shape Benchmark (PSB) [6], which consists of 3D objects described as triangular meshes. PSB training set contains 907 models, categorized into 90 classes. The meshes contain 7460 triangles and 4220 vertices on the average. We have validated our results using the test set of PSB and on another database, Sculpteur (SCUdb) [2,8], which is fundamentally different from PSB in terms of both classification semantics and mesh quality. The performance figures are given in terms of discounted cumulative gain (DCG) and nearest neighbor (NN) scores [6]. Recall that NN is the percentage of the closest matches that belong to the query class and DCG is a statistic between 0 (worst) and 1 (best), weighting correct results, i.e., models of the same class as the query model, near the front of the list more than correct results later in the ranked list. Our descriptors have size $N_m \times N_d = 1024$. Specifically, the magnitude set for targets is chosen by uniform sampling ($N_m = 8$) and the direction set by octahedron subdivision (recursion factor $a = 2$, $N_d = 128$). We also note that all descriptors have been normalized prior to distance calculation.

Table 1. Percent DCG For Standard Dissimilarity Measures, $N_m = 8$ and $N_d = 128$

Feature	L^1	L^2	L^∞	KL	χ^2	B
S_{rad}	56.7	54.7	44.4	54.4	57.0	56.7
S_{tan}	59.9	55.3	47.1	58.2	61.1	59.4

Table 2. Percent DCG Values on PSB Test Set and SCUdb using Various Metrics

Database	Feature	$L^{0.6}$	L^1	L^2	L^∞	KL	χ^2	B
PSB Test	S_{rad}	55.0	54.4	52.0	43.6	53.4	54.7	54.9
	S_{tan}	58.9	58.0	54.5	47.0	55.8	58.9	58.6
SCUdb	S_{rad}	70.9	71.1	70.7	62.8	70.1	71.3	71.0
	S_{tan}	72.8	72.0	69.5	63.1	70.1	71.6	71.3

4.1 Effect of the Dissimilarity Measure

We test the impact of the chosen distance metric on the retrieval performance, namely, those of L^p , Kullback-Leibler (KL), Chi-Square (χ^2), and Bhattacharya (B) distances. For L^p , $p = 1, 2, \infty$, descriptors are rescaled to have unit L_p -norm. For KL , χ^2 , and B , on the other hand, we have rescaled the descriptors to unit L^1 -norm because these measures are designed for histograms and/or pdfs. We have observed that for all metrics, appropriate normalization invariably improves discrimination. In Table 1, we provide DCG scores using these metrics. For both the radial feature S_{rad} and the tangent plane-based feature S_{tan} , the measures L^1 , χ^2 , and B yield the best results while L^∞ has the poorest performance. Intrigued by the lower performance of L^2 with respect to L^1 , we explored the variability of DCG as a function of the p parameter of the L^p -metric within the interval $p \in]0, 2]$. As illustrated in Fig. 3, the performance degrades significantly for $p > 1$, and for $0 < p \leq 1$, DCG peaks around $p = 0.6$, even outperforming χ^2 and B , the best measures of the previous experiment (see Table 1). This $L^{0.6}$ peak is however minor with respect to L^1 , hence we preferred the computationally cheaper L^1 -metric in the following results. We remark also that the performance order of the considered metrics remain the same for PSB test set and SCUdb [2,8] as can be seen in Table 2.

4.2 Effect of the Sampling Scheme and the Target Set Size

In Table 3, we provide the DCG and NN scores corresponding to all combinations of the sampling schemes presented in Sect. 2.2. We would like to point a few subtleties in target selection. For uniform sampling of the pdf, the step size factor Δs_n becomes constant and then irrelevant. When sampling the magnitude with equal probability intervals (same area under the pdf curve) and/or when sampling the unit sphere with equal spherical coordinate steps, neglecting the Δs_n -factor causes significant performance loss. When Δs_n -factor is included in Eq. 2, then the performances of equal area sampling and uniform (equal distance) sampling of the magnitude become virtually equal (within 1%). The same

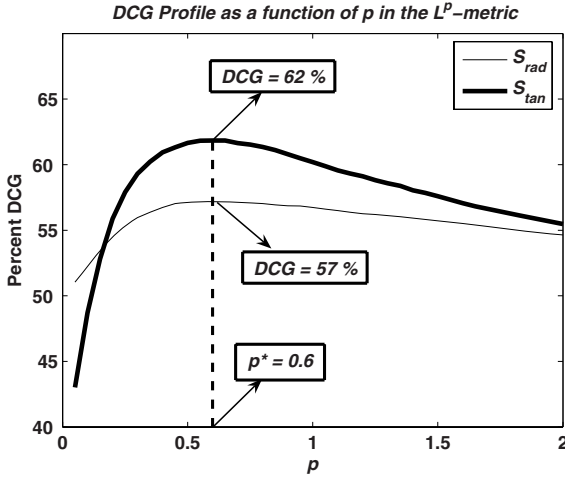


Fig. 3. DCG Performance vs the p parameter in the L^p -metric, N_m and $N_d = 128$

Table 3. DCG and NN Performances Using Different Sampling Schemes For Target Selection, N_m and $N_d = 128$, **1**: magnitude set: *uniform*, direction set: *by subdivision*, **2**: magnitude set: *uniform*, direction set: *by spherical coordinates*, **3**: magnitude set: *equal probability*, direction set: *by subdivision*, **4**: magnitude set: *equal probability*, direction set: *by spherical coordinates*

	S_{rad}				S_{tan}			
	1	2	3	4	1	2	3	4
%DCG	56.7	56.8	56.0	56.3	59.9	60.5	59.5	60.1
%NN	57.3	55.9	55.7	55.1	58.7	60.8	59.8	60.5

Table 4. Percent DCG For Various Target Set Sizes

	S_{rad}			S_{tan}		
	$N_d = 512$	$N_d = 128$	$N_d = 32$	$N_d = 512$	$N_d = 128$	$N_d = 32$
$N_m = 16$	57.2	56.6	52.2	60.7	60.6	58.3
$N_m = 8$	57.0	56.7	52.0	60.5	59.9	58.1
$N_m = 4$	55.8	55.4	49.9	57.3	57.1	53.4

observation also holds for sampling the unit sphere by octahedron subdivision and by spherical coordinates. Table 3 shows that all sampling schemes result in equivalent performances. We have also investigated the effect of changing the clipping level to determine the magnitude level. Our experiments with clipping levels 0.5%, 1%, 2%, and 5% yielded comparable results.

We have also analyzed the effect of changing the descriptor size N between 128 and 8192 under uniform sampling (i.e., uniform magnitude set and octahedron subdivision for direction set) using the L^1 -metric. We have tested the

combinations of various $N_m \in \{4, 8, 16\}$ and $N_d \in \{32, 128, 512\}$ values for the target set size. Table 4 reveals that to maintain adequate DCG performance, N_m should not be less than 8 and that N_d should be at least 128. We also remark that, for the combination $(N_m, N_d) = (8, 128)$, preprocessing, feature calculation and density estimation takes 0.4 second CPU time on the average on a Pentium M 1.86 GHz processor, 1 GB RAM.

4.3 Dimensionality Reduction

We now report the outcome of the dimensionality reduction experiments via *marginalization* and *probability density pruning*. We define the efficiency ϵ as the ratio of DCG after reducing dimensionality to the baseline DCG (i.e., the DCG of the “non-reduced” descriptor), concretely, $\epsilon = \text{DCG}_{S, \text{reduced}} / \text{DCG}_{S, \text{full}}$. Fig. 4 summarizes the effect of marginalization. On the left plot in Fig. 4, we see that, when only one feature is marginalized, then one of the direction components, i.e., \hat{R}_x, \hat{R}_y or \hat{R}_z for S_{rad} , and \hat{N}_x, \hat{N}_y or \hat{N}_z for S_{tan} , can be sacrificed. We do not incur into any loss in marginalizing one of the direction components and the descriptor size is halved. This should not be a surprise since the redundancy of a directional component given the other two was already pointed out in Sect. 3.2. Even more impressive economies can be attained by marginalizing two components, as shown on the right of Fig. 4. As long as we retain the magnitude component and keep only one of the direction components, we can still achieve DCG efficiency at around 95% with 83% reduction in descriptor size (i.e., the size is reduced by a factor of 6). It is also worth noting that, with these results at our disposal, we can directly estimate the pdf of the most informative components and reduce the computational overhead beforehand. Although, due to space restrictions, we cannot provide our validation results on the effect of marginalization using PSB test set and SCUdb, we note that the above observations hold invariably for other databases.

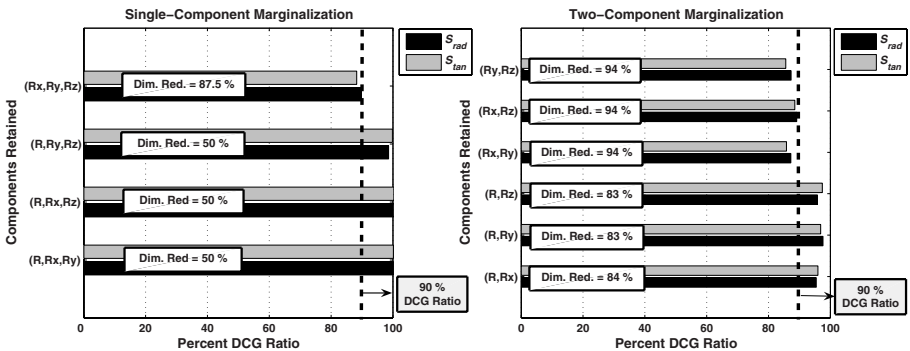


Fig. 4. Effect of dimensionality reduction by marginalization on DCG performance: horizontal axis stands for percent ratio of DCG after marginalization to the original DCG obtained using the full descriptor of size $N = 8 \times 128 = 1024$

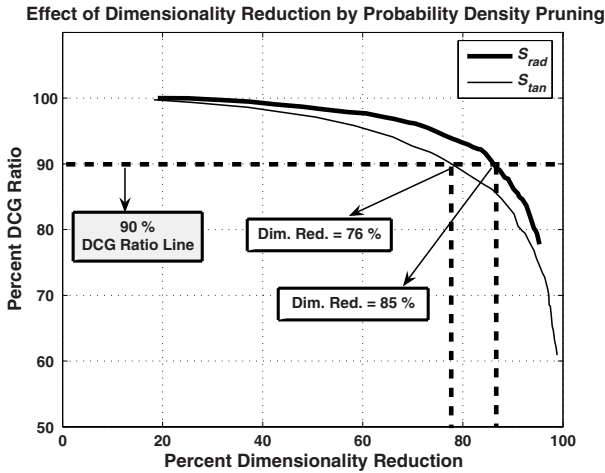


Fig. 5. Effect of reducing dimensionality by probability density pruning on DCG performance

Finally, we experiment with the probability density pruning technique by varying the threshold parameter λ (Sect. 3.2). In Fig. 5, we display the DCG efficiency as a function of percent reduction in descriptor size. We observe that DCG efficiency profile remains nearly flat at around 100% up to 40% reduction in descriptor size and starts to degrade afterwards. Comparing Figs. 4 and 5, at the 90% DCG efficiency level, descriptor simplification with pruning is slightly less effective (85%) than with marginalization (94%). We want to point out that for density pruning, unlike marginalization, the reduction in size is controllable as a function of the performance loss incurred. Furthermore, insignificant targets can be eliminated off from density estimation stage, reducing the computational overhead beforehand.

5 Conclusion

In this work, we have investigated density-based 3D shape descriptors [1, 2, 3] from the perspective of retrieval efficiency and effectiveness. While standard dissimilarity measures for histogram and/or pdf matching, such as L^1 , χ^2 , and B , provide similar and satisfactory retrieval performances, the best metric has proven to be the non-classical $L^{0.6}$. More generally, the DCG performance profile of the L^p -metric as a function of p , strictly increases from $p = \infty$ to $p = 0.6$, and then decreases. Nevertheless, we advise the use of the L^1 -metric as it provides adequate discrimination with the least computational effort. The experiments on target selection demonstrate that, as long as discretization step size is taken into account in dissimilarity computation, all sampling schemes are equally effective on retrieval performance. We did not gain much by increasing the descriptor size after $N = 1024$ and the best (N_m, N_d) -configuration has been (8,128).

We have also introduced two flexible dimensionality reduction techniques specifically for density-based descriptors. Our marginalization tool exploits possible redundancy of different components in the local features and can provide significant reduction in descriptor size with no or little loss of discrimination ability ($\epsilon_{marg} = 100\%$ at 50% reduction). Marginalization can also be used to explore the effectiveness of different local features by selectively removing one or multiple components at a time. Probability density pruning, maybe not as efficient as marginalization in dimensionality reduction for fixed DCG efficiency ($\epsilon_{pruning} = 100\%$ at 40% reduction), provides a machinery that controls the amount of performance loss as a function of reduction in descriptor size and vice versa. In marginalization, reduction in descriptor size occurs in discrete steps; in density pruning, the parameter λ enables us to reduce the descriptor size arbitrarily. Finally, while marginalization tells us which feature *components* are more effective for retrieval; probability density pruning, although not to the full extent, guides us in choosing more informative feature *locations*, i.e., targets.

Our future work will concentrate on enlarging the repository of local features for use in the density-based shape description framework. Furthermore, we will consider joining the radial and tangent plane-based features and perform density estimation on this “augmented” space in order to exploit the full surface information probed by these features.

References

1. Bustos, B., Keim, D.A., Saupe, D., Schreck, T., Vranic, D.V.: Feature-based similarity search in 3D object databases. *ACM Comput. Surv.* **37** (2005) 345–387
2. Akgül, C.B., Sankur, B., Yemez, Y., Schmitt, F.: Density-based 3d shape descriptors. *EURASIP Journal on Advances in Signal Processing* **2007** (2007) Article ID 32503, 16 pages doi:10.1155/2007/32503.
3. Härdle, W., Müller, M., Sperlich, S., Werwatz, A.: *Nonparametric and Semiparametric Models*. Springer Series in Statistics. Springer (2004)
4. Akgül, C.B., Sankur, B., Schmitt, F., Yemez, Y.: Density-based shape descriptors for 3D object retrieval. In: *International Workshop on Multimedia Content Representation, Classification and Security (MRCSS '06)*, Istanbul, Turkey (2006)
5. Akgül, C.B., Sankur, B., Yemez, Y., Schmitt, F.: A framework for histogram-induced 3D descriptors. In: *European Signal Processing Conference (EUSIPCO '06)*, Florence, Italy (2006)
6. Shilane, P., Min, P., Kazhdan, M., Funkhouser, T.: The Princeton shape benchmark. In: *Proc. of the Shape Modeling International 2004 (SMI '04)*, Genoa, Italy (2004) 167–178
7. Vranić, D.V.: *3D Model Retrieval*. PhD thesis, University of Leipzig (2004)
8. Tung, T.: *Indexation 3D de bases de données d'objets par graphes de Reeb améliorés*. PhD thesis, Ecole Nationale Supérieure des Télécommunications (ENST), Paris, France (2005)
9. Yang, C., Duraiswami, R., Gumerov, N.A., Davis, L.: Improved fast Gauss transform and efficient kernel density estimation. *ICCV* **1** (2003) 464
10. Zaharia, T., Prêteux, F.: Shape-based retrieval of 3D mesh models. In: *Proc. of the IEEE International Conference on Multimedia and Expo (ICME'2002)*, Lausanne, Switzerland (2002)