Second Order 3D Shape Features: An Exhaustive Study

Marco Reisert [a 1] and Hans Burkhardt [a]

[a] *University of Freiburg, Computer Science Department, Georges-Köhler-Allee 52, 79110 Freiburg i.Br., Germany*

Abstract

In recent years many techniques for 3D shape retrieval and classification were proposed. Most of them follow the feature vector paradigm, i.e. the shape is extended with some compact feature representation, on which basis the objects are compared and similarity measures are computed. A main demand of such similarity measures is their invariance to Euclidean motion. There are three main direction to obtain such invariances: Matched Filter Approaches, Pose Normalization or Transformation Group Integration. Among those, registration approaches perform best in retrieval accuracy, while their computational expense however is rather high. In contrast, representation obtained by group integration are fast to compute and compare, but show bad retrieval performance due to its loss in information. In this article we try to close this gap and show that it is also possible to obtain meaningful representations of surface models by group integration approaches.

Key words: Shape search and retrieval, Shape Distribution, Spherical Harmonics, Polygonal Mesh, geometric matching

1 Introduction

With the growing amount of freely available 3D models the need for methods to retrieve and search in large databases of 3D models have become more and more important. Usual 3D shape representations lack of intrinsic and concise information about the object. There is a need for small, compact and discriminative representations of the objects to perform fast and reliable queries. The main demand for reliable similarity measurements is the invariance to

[1] Corresponding author: Georges-Köhler-Allee 52, 79110 Freiburg i.Br., Germany, Tel. +49-(0)761-203-8267, Fax. +49-(0)761-203-8262, email: reisert@informatik.uni-freiburg.de

Euclidean motions of the objects. To capture this demand, there are three main direction:

(1) **Matched Filter.** The invariance is obtained while comparing two objects in all possible positions and orientations. The relative pose of the two objects with the highest similarity has to be determined. The amount of similarity for this maximal match serves as the measure.

(2) **Invariance.** The feature representation of the objects is already invariant.
   (a) **Normalization.** The objects are normalized with respect to rotation and translation. Mostly the translation is normalized by shifting the Center of Gravity (COG) to the origin. To eliminate the rotation dependency the second order moment matrix (inertia tensor) is diagonalized. The dissimilarity is determined by an appropriate distance measure.
   (b) **Group Integration.** In contrast to the first technique, where the invariance of the representation is obtained by normalization, it is also possible to obtain the invariance intrinsically. Those features can be easily obtained by averaging arbitrary kernel functions over the considered transformation group (in our case: the Euclidean motion)

Group Integration is a well known approach to obtain invariant features for gray-scale images (1; 2; 3). Group Integration is a constructive approach, i.e. for a given kernel function $f(\mathbf{x})$ of the object we construct an invariant feature by integrating the kernel over the considered group:

$$F = \int_{\mathcal{G}} f(g\mathbf{x})\ dg \tag{1}$$

Recently it was shown that there is also a close connection of such group integration features to the world of discrete structures (4). In fact, it is possible to embed Osada's Shape Distribution (5) in the theory of group integration (6).

In practice, the three methods mentioned above often intersect and are used together. As already mentioned the best performing descriptors are often based on matching schemes. But those have also the highest computational demands and the related similarity measures are not well suited for advanced search techniques and fast indexing schemes. Approaches where the translational and rotational invariance are obtained by group integration often show bad performance due to the putative loss of information. But comparing the objects can easily be done. The corresponding feature vector elements are known a priori and no matching need to be performed. One can simply use e.g. the Euclidean distance for measuring the dissimilarity. The Euclidean distance is often the basis for fast nearest neighbor searches and indexing schemes (7; 8).

Also partial queries have proper meanings. For example in the case of Osada's D2 Shape Distributions (5) a query for longish objects is easily done by asking the database for objects with a high probability for low distances. Recently Ohbuchi et al. (9) presented an improvement of D2 named AD, which outperformed D2 with only a low amount of additional computational costs. For the AD-features partial queries can also have meaningful interpretations.

In this paper we present various further improvements of D2 and AD and show that also small and compact representations obtained by group integration can lead to reliable and informative descriptions of the objects.

The article is organized as follows: in the following subsection we give a rough overview of existing techniques and previous work. In Section 2 we propose the examined features and corresponding similarity measures and Section 4 presents intensive experiments based on the Princeton Shape Benchmark (10).

1.1 Previous Work

A lot of researchers are working on the problem of content based 3D Shape retrieval. The work has led to numerous shape retrieval systems (11; 12; 13; 14; 15), which are mostly based on the feature vector paradigm. Survey papers are provided by Tangelder and Veltkamp (16; 17) and Iyer et al. (18), they give good overviews. Bustos et al. give a comparative study in (19). In the following we shortly summarize the approaches, which are closely related to ours.

We already mentioned Osada's Shape Distributions (5) and Ohbuchi's extension (9) of it. Both are based on measuring the frequency of some translation and rotation invariant properties (e.g. distance) of two random points on the surface of the model. Ohbuchi also extended the approach to a multi-resolution descriptor, which outperforms his simple AD-descriptor, but by the cost of high computation time. Ip et al. (20) used a modified version of the Shape Distributions for categorization of CAD related solid models. Vranic et al. (12) compute a so-called spherical extent function of the model-surface and make a spherical harmonic transform of this function. The rotational-invariance is obtained by normalization. In another work (21) Vranic generalized the spherical extent function by incorporating also the surface orientation. Kazhdan et al. (22) eliminate the rotational dependency by taking the magnitude of the invariant subspaces of the Spherical Harmonic Transform (SHT). They show, that in most cases this is the better alternative than a normalizing approach. In their experiments the SHT-magnitudes of the Gaussian Euclidean Distance Transform (GEDT) perform best. The GEDT is a volume-function, which gives the Gaussian of the Euclidean distance to the nearest point on the surface of the

model. Novotni et al. (23) used 3D Zernike descriptors to describe voxelized shape models. Recently Chen et al. (13) proposed the so-called Light Field Descriptors (LFD), which are based on a set of views of the object, uniformly distributed on a sphere. The LFDs performed best on the Princeton Shape Benchmark (PSB) (10). The PSB forms the basis of our experiments.

There has been a lot of work concerning group integration features. Ronneberger et al. (2) used group integration to classify pollen grains. Schael (24) used group averaging for texture classification. In (25) the connection of group integration and histograms was established and applied to content-based image retrieval. Haasdonk et al. (26) used partial group integration to obtain robustness to slight transformation.

2 The examined features

In this section we introduce the examined second-order features. Let us first explain, what *second-order* exactly means. Let a surface be a measurable subset $\mathbf{S} \subset \mathbb{R}^3$, which is represented by its indicator-function $x : \mathbb{R}^3 \mapsto \mathbb{R}$, i.e. x is given by $x(p) = \int_{\mathbf{S}} g(p-s)\ ds$, where g is some isotropic 'fill'-function, for example a Gaussian or the Dirac-Distribution. A second-order monomial Group Integration feature is written as

$$F = \int_{\mathcal{G}} x(Ru_1 + p)\ x(Ru_2 + p)\ dR\ dp, \tag{2}$$

where dR and dp are appropriate group measures for the rotation and translation group. Such features are successfully used e.g. in (2) for pollen recognition. In (6) it was shown that such features can be expressed by

$$F = \int_{(p_1,p_2)\in\mathbf{S}\times\mathbf{S}} h(p_1,p_2)\ dp_1 dp_2. \tag{3}$$

Now the integral ranges over the cartesian product of the shape $\mathbf{S}$ with itself, instead of the group as in (2). The function $h : \mathbb{R}^3 \times \mathbb{R}^3 \mapsto \mathbb{R}$ is some real-valued function depending on the choice of the 'fill'-function g. In the case of the monomial features (2) the function h can be interpreted as some kind of histogram-binning function. And it is actually the case that we arrive at Osada's Shape Distributions, if we choose $h(p_1,p_2) = \delta_d(||p_1 - p_2||)$, where δ_d is the usual histogram-binning-function. There is also a illustrative explanation of the equivalence of both approaches. Move two points with some fix distance d in all orientation and all translational poses over the shape, multiply the values of the indicator function x at those points and sum everything up. This is nothing else than integrating a second order monom over the Euclidean group. And it is also obvious that this integral exactly counts the frequency

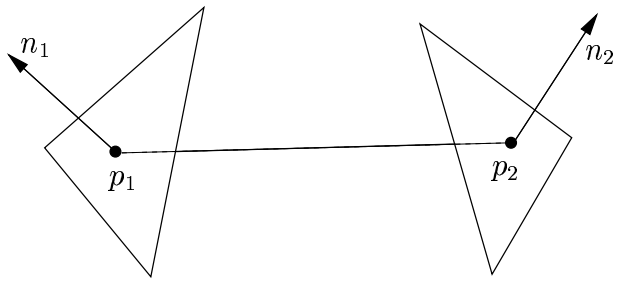

Fig. 1. Two points on the surface of a model. The configuration is uniquely determined up to translation and rotation by: the distance $||p_1 - p_2||$, the 3 cosines $n_1 \cdot n_2$ and $n_1 \cdot \frac{p_1 - p_2}{||p_1 - p_2||}$ and $n_2 \cdot \frac{p_1 - p_2}{||p_1 - p_2||}$, and the orientation of the three vectors n_1,n_2 and $(p_1 - p_2)$.

that two points on the surface of the model have distance d, which is nothing else than Osada's D2 Shape Distribution.

We do not want to restrict to such simple binning-functions like $\delta_d(||p_1 - p_2||)$. Since we are working with oriented surfaces, any $p \in \mathbf{S}$ is attached with a surface-normal n, describing the pose of the surface-plane (in Figure 2 we give an illustration of the configuration), so we are also able to include the surface-information in our binning-function, which gives us the opportunity to extract more information.

In the following we will mention Osada's D2 Shape Distributions first, they are some kind of basis of all upcoming methods. The pure Distance Histogram D2 is extended by a more discriminative Spherical Harmonic description as Kazhdan (27) has proposed in his dissertation, but we show how more information one can save than just the norms of the spherical harmonic descriptors. The enhanced version of D2 by Ohbuchi, named AD, also incorporating angles between surface normals is also considered. Further we introduce a new descriptor relying on the joint distribution of the distances, the angles between the distance vectors and the surface normals. Also the combination of those features with Ohbuchi's is considered. And finally we extend those with a spherical harmonic description to enhance the discriminativity. In the following we introduce the considered shape functions:

- **Distance Histograms (D)**
  The well known D2 Shape Distributions by Osada et al. can be written as

$$D(d) = \int_{(p_1,p_2)\in\mathbf{S}\times\mathbf{S}} \delta_d(||p_1 - p_2||)$$

  (in the following we omit the differential $dp_1 dp_2$ for simplicity). In practice the integral is approximated by Monte-Carlo integration, i.e. we adopted the method proposed in (5). The binning function δ_d is defined as usual, i.e. $\delta_d(d')$ gives a contribution if d' falls into an interval around d, where the the intervals are all non-overlapping and equal sized.

- **SHT Distance Histograms (SD)**
  The simple D2 Shape Distributions can be interpreted as the lowest order mode of a more general spherical harmonic description of the pair-wise distances of the model. Let $\mathbf{a}_l(d) \in \mathbb{C}^{2l+1}$ be the spherical descriptor obtained by

$$\mathbf{a}_l(d) = \int_{(p_1,p_2)\in \mathbf{S}\times\mathbf{S}} \mathbf{Y}_l(\tfrac{p_1-p_2}{||p_1-p_2||})\ \delta_d(||p_1-p_2||),$$

  where $\mathbf{Y}_l = (Y^l_{-l}, ..., Y^l_l)^T$ are the usual spherical harmonics. It is well known, that the magnitude of the $\mathbf{a}_l(d)$ are invariant to rotations of the underlying object. Thus we choose

$$SD_l(d) = ||\mathbf{a}_l(d)||$$

  for our rotation and translation-invariant descriptor. Since $Y_0 \propto 1$, it is obvious that $SD_0(d) \propto D(d)$, i.e. $SD_l(d)$ is a generalization of $D(d)$. Besides, only SD_l with even l is actually unequal to zero, which is caused by the point-symmetry of the distance distribution.

- **Extended SHT Distance Histograms (SDE)**
  Let us shortly examine why magnitudes of the $\mathbf{a}^l$ are invariant to rotations. The subspaces where the $\mathbf{a}_l$ live in are invariant to rotations, i.e. a rotation of the objects acts by an unitary transformation on those subspaces. The description of the rotated object maps by $\mathbf{a}_l \mapsto \mathbf{D}_l\mathbf{a}_l$, where the $\mathbf{D}_l$ are the so-called Wigner D-matrices, which are unitarian and have a direct connection to the applied rotation. Now it is clear why the magnitudes are invariant,

$$||\mathbf{D}_l\mathbf{a}_l||^2 = (\mathbf{D}_l\mathbf{a}_l)^\dagger\mathbf{D}_l\mathbf{a}_l = \mathbf{a}_l^\dagger\mathbf{D_l}^\dagger\mathbf{D}_l\mathbf{a}_l = \mathbf{a}_l^\dagger\mathbf{a}_l = ||\mathbf{a}_l||^2,$$

  where we use $\mathbf{D}_l^\dagger\mathbf{D}_l = \mathbf{I}$. This statement also holds by taking two different descriptors, namely $\mathbf{a}^l(d_1)$ and $\mathbf{a}^l(d_2)$, the only requirement is, that they have the same transformation behavior under rotations. Thus, $(\mathbf{D}_l\mathbf{a}_l(d_1))^\dagger\mathbf{D}_l\mathbf{a}_l(d_2) = \mathbf{a}_l(d_1)^\dagger\mathbf{a}_l(d_2)$ also holds and we define our extended features by

$$SDE_l(d_1,d_2) = \sqrt{|\mathbf{a}_l(d_1)^\dagger\mathbf{a}_l(d_2)|}\ \text{sign}\ (\mathbf{a}_l(d_1)^\dagger\mathbf{a}_l(d_2)).$$

  The product $\mathbf{a}_l(d_1)^\dagger\mathbf{a}_l(d_2)$ is actually real and thus we only have to recover the sign. Indeed SDE equals SD for equal distance arguments.

- **Alpha/Distance Histograms (AD)**
  Like Ohbuchi et al. already discovered it is possible to incorporate more information than only the distance between two selected points of the shape. Every point on the surface of the shape is associated with an orientation of the surface plane, represented by its surface normal. The AD-histogram simply measures the joint distribution of the distance of two points and the

cosine of the angle of the two associated surface normals, formally,

$$AD(d,\alpha) = \int_{(p_1,p_2)\in \mathbf{S}\times\mathbf{S}} \delta_d(||p_1 - p_2||)\delta_\alpha(|n_1 \cdot n_2|),$$

where n_1 and n_2 are the surface normals and $\cdot$ the dot-product in $\mathbb{R}^3$. We neglect the sign of the cosine, because in most cases the sign of the normals are not consistently chosen.

- **Beta/Distance Histograms (BD)**
  Another possibility is to measure the angle between the distance vector $p_1 - p_2$ and one of the surface normals n_1 or n_2. It is not necessary to measure both angles, because the resulting histogram would be symmetric with respect to the exchange of n_1 and n_2. We call the joint distribution of the distances and one of the angles the BD-histogram, which is defined by

  $$AD(d,\beta) = \int_{(p_1,p_2)\in \mathbf{S}\times\mathbf{S}} \delta_d(||p_1 - p_2||)\delta_\beta(\tfrac{|n_1\cdot(p_1-p_2)|}{||p_1-p_2||}).$$

  Thereby the sign of the cosine was again neglected due to the reason mentioned above.

- **Alpha/Beta Histograms (AB)**
  It is also possible to forget about the distance between the two points and letting the feature only rely on the angles. An important property of an AB-histogram is the scale-invariance, i.e. an isotropic scaling do not have any effect to the features.

- **Alpha/Beta/Distance Histograms (ABD)**
  Of course, it is not a bad idea to combine all three informations, the distances, the angles between the normals and the angles between the distance vector and one of the normals in a three dimensional histogram, which we call the ABD-histogram, which looks as follows

  $$ABD(d,\alpha,\beta) = \int_{(p_1,p_2)\in \mathbf{S}\times\mathbf{S}} \delta_d(||p_1 - p_2||)\delta_\alpha(|n_1 \cdot n_2|)\delta_\beta(\tfrac{|n_1\cdot(p_1-p_2)|}{||p_1-p_2||}).$$

  It is worthwhile to mention that this type should keep the most information back and is expected to perform best among the pure histograms.

- **Alpha/Beta/Distance SHT Histograms (ABSD)**
  Similarly to the SD-features, we are able to use the idea of the Spherical Harmonic Transform for the ABD-features to enhance the information content. For every bin we compute the spherical harmonic description of the pairwise distances up to a certain order and take the norm of the invariant subspaces as rotation-invariant features.

- **Alpha/Beta/Distance Extended SHT Histograms (ABSDE)**
  The same as for the SDE-features. But here we have to be careful, because the feature size grows tremendously, i.e. the feature size is the squared size of the features above.

2.1 Implementation

We adopted the method proposed in (5) to generate uniformly distributed points on the surface of the considered models. Our histograms always consist of equally sized bins. The number of bins are denoted by N_D, N_A and N_B, hence a ABD-histogram is of size $\#bin = N_D N_A N_B$. The featuresize for ABSD is given by $\#bin \cdot l_{max}/2$, for ABSDE by $\#bin(\#bin+1)/2 \cdot l_{max}/2$.

Every quantity, i.e. the measured distances and cosines, is recorded in the interval $[0,1]$ (the objects we consider are scaled such that their extents are about 0.5). For the Spherical Harmonics implementation we used the C++ library *Matpack*. All experiments were done on *Intel P4 2.8 Ghz* with 512MB on a *Linux*-system.

3 Similarity Measures

We only consider symmetric distance measures. As we want to compare some kind of probability distribution, we are tended to use distance measures known from information theory and statistics. In the following we concentrate on four measures.

- L_1 **- Norm** $d_1(F_1,F_2) = \sum_{i=1}^{n} |F_1(i) - F_2(i)|$. The L_1 - Norm is closely related to the Histogram-Intersection Kernel, which predestines it for our purposes.

- L_2 **- Norm** $d_2(F_1,F_2) = \sum_{i=1}^{n} (F_1(i) - F_2(i))^2$. The Euclidean distance is probably the most common distance measure.

- **Symmetric Kullback-Leibler divergence** The Kullback-Leibler divergence can be considered as a kind of a distance between the two probability densities. Originally it is not symmetric, we consider the symmetrizied KL-divergence:
$$d_{KL}(F_1,F_2) = \sum_{i=1}^{n} (F_1(i) - F_2(i)) \log(\tfrac{F_1(i)}{F_2(i)}).$$
Since our features only can be considered in some cases as probability distributions, we adapted the above distance measure to signed features by

Descriptor	size	comp.time (s)	$N_A/N_B/N_D$	l_{max}
D	32	1.0	-/-/32	-
SD	40	1.6	-/-/10	8
SDE	220	1.6	-/-/10	8
AD	256	1.4	16/-/16	-
BD	256	1.4	-/16/16	-
AB	256	1.4	16/16/-	-
ABD	256	1.4	4/4/16	-
ABSD	768	1.7	4/8/8	6
ABSDE	6240	2.0	4/4/4	6

Table 1
Parameters for feature extraction.

taking the absolute values for the arguments of the logarithm and adding some small epsilon for numerical stability.

- **Symmetric χ^2 distance measure** The χ^2-distance is an approximation of the Kullback-Leibler distance.

$$d_{\chi^2}(F_1, F_2) = \sum_{i=1}^{n} \frac{(F_1(i) - F_2(i))^2}{F_1(i) + F_2(i)},$$

where we use a symmetrizied version again. The χ^2-distance is much less computional expensive than d_{KL}.

4 Experiments

In this section we present intensive experiments evaluating the proposed features. First we give a short overview over the used data and give a description of the considered performance measures. For a detailed description see (10). Finally we present the results and discuss them.

4.1 The Database and the Performance Measures

The basis of our experimental examinations is the *Princeton Shape Benchmark* (PSB) recently proposed by Shilane et al. (10). The PSB provides a set of 1814 classified 3D models collected from the Web. The models are divided in 907 training and test-models. The classification is given in a hierarchical

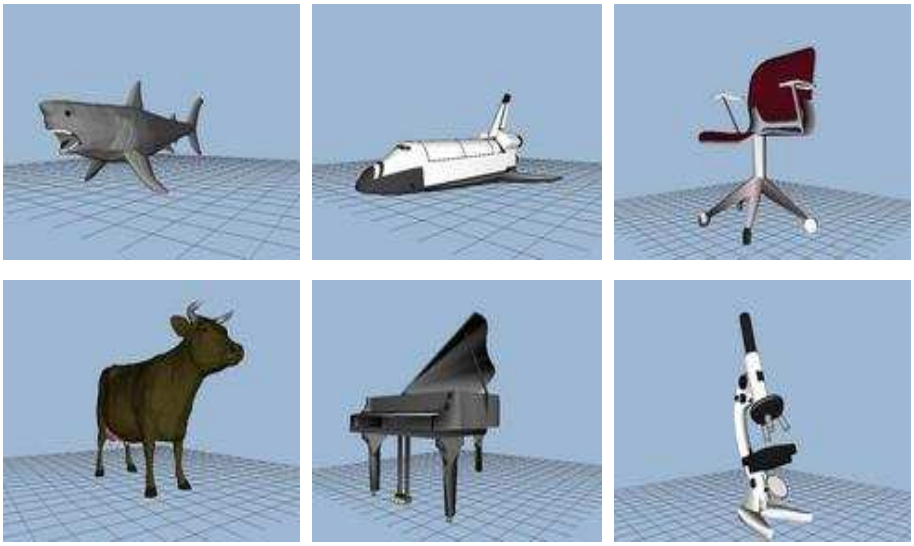

Fig. 2. Examples (28) from the Princeton Shape Benchmark (10)

manner. There are four different granularities. The finest level groups the objects by function and form, continuing to coarser levels, where finally the coarsest discriminates between natural and artificial objects only.

The *Princeton Shape Benchmark* also provides several tools for measuring the performance of the considered features. In the following we give a short description of the quality measures.

- **Nearest Neighbor** The leave-one-out error of the nearest neighbor classifier. The percentage that the closest matches belong to the same class as the query object.

- **N-tier** The percentage of models belonging to the query class that appear within the top $N*(|C|-1)$ matches, where $|C|$ is the size of the query-class. N is chosen 1 or 2.

- **E-Measure** The E-Measure incorporates only the first 32 retrieved models and calculates Precision (P) and Recall (R) for those results. It is defined by $E = \frac{2PR}{P+R}$. Higher values for E indicate better results.

- **Discounted Cumulative Gain (DCG)** An user of a search engine is usually more interested in results of the first result page, thus those should be more weighted than the later ones. Considering a query result for object k from class C, we call

$$DCG_k = \frac{\sum_{i \in C} \mathrm{ld}^{-1}(rank_k(i))}{\sum_{j=2}^{|C|} \mathrm{ld}^{-1}(j)}$$

the Discounted Cumulative Gain of the query, where $rank_k(i)$ is the rank

of object i in the query for object k. The final score is the average of DCG_k over all models in the database.

4.2 Results and Discussion

The experiments were only performed on the Test-Set of the database to keep the results comparable to those presented in (10). We computed the performance measures given above for the four distance measures given in Section 3 and for the finest granularity of the classification hierarchy. In Table 1 we specify the parameters of the proposed features we used in the experiments. We have chosen the size of the bins such that two seconds of computation time are enough to compute stable features. With an increasing number of bins it is obvious to increase the number of random events as well to get stable features.

In Table 2 the results for the finest level of classification are given. The corresponding Precision/Recall-Graphs are given in Figure 3. Let us analyze the results. Looking at the different similarity measures, the χ^2-distance is the overall winner. The KL-divergence and L_1-norm give very similar results, but recall that for the KL-divergence we need to evaluate the logarithm, which needs a lot of time in comparison to the simple computation of L_1-norm. The L_2-norm looses the competition in all disciplines.

Now let us have a look at the different features competing against each other. The pure distance histogram obviously performs worst. The spherical harmonic descriptors SD and SDE improve the performance of D. The SDE-descriptors perform better than SD, which one could have expected. One can suppose that the idea of the SDE-features could improve any SHT-based features, which get rid of the rotation by taking the norm of the invariant subspaces, but of course the number of features grows. Having a look at the angle-based histograms, it is actually the better choice to use our new BD-features instead of the AD-features proposed by Ohbuchi, i.e. it seems that the angle between surface-normal and distance-vector gives us more information for the retrieval task than the angle between the surface-normals. The pure angle histograms AB do not show such a performance like AD or BD, which may be caused by the additional scale-invariance. But for applications, where scale-invariance is explicitly demanded, AB could do a good job and one need not to bother about a scale normalization. As already expected the ABD-histograms are performing very well, they beat all except its own extensions. The SHT extensions of the ABD-features improve the results of ABD, but not in such a rigorous manner like it was the case for the D-features. Unfortunately the feature size grows for ABSDE so tremendously, that one is only able to choose a 4/4/4 bin-distribution to get an acceptable feature size, while the improvement is much less than one would expect (except for the

Descriptor	NN	1-T	2-T	EM	DCG	NN	FT	ST	EM	DCG
Distance	L_1					L_2				
D	33.3	16.0	24.4	14.7	43.8	32.9	15.9	24.1	14.6	43.7
SD	37.3	18.9	27.4	16.1	46.4	36.9	18.4	26.5	15.8	45.7
SDE	43.0	21.1	30.4	17.6	49.0	42.8	21.2	29.7	17.4	48.5
AD	42.9	21.6	31.3	18.8	49.9	38.8	20.1	28.9	17.3	48.0
BD	50.1	25.0	34.8	20.3	53.1	44.1	22.9	31.7	18.5	50.5
AB	41.0	21.6	31.5	18.9	50.2	37.3	18.7	27.5	16.4	47.0
ABD	54.5	27.7	37.9	22.6	55.9	50.4	25.1	34.4	20.5	53.0
ABSD	56.3	28.8	39.2	23.4	57.0	52.3	24.8	34.1	20.3	52.9
ABSDE	55.7	30.3	41.4	24.5	57.9	55.5	29.1	40.2	23.5	56.9
Distance	KL					χ^2				
D	33.8	16.6	25.0	15.0	44.3	34.2	16.9	25.2	15.2	44.5
SD	37.8	18.9	27.2	16.1	45.2	38.6	19.2	27.2	16.3	46.5
SDE	37.2	21.3	30.4	18.0	48.5	41.7	21.7	30.6	18.2	49.1
AD	42.3	21.2	30.7	18.2	49.6	43.7	21.6	31.1	18.5	49.9
BD	48.1	24.6	34.1	19.9	52.6	48.6	24.9	34.7	20.2	53.1
AB	41.1	21.4	30.9	18.5	49.9	41.2	21.6	31.1	18.8	50.2
ABD	53.9	27.0	37.1	22.3	55.3	54.7	27.4	37.5	22.5	55.8
ABSD	57.6	28.8	38.9	23.1	56.8	58.0	29.2	39.7	23.5	57.3
ABSDE	52.6	27.3	37.9	22.5	55.2	57.1	30.5	41.5	24.4	58.2

Table 2
Performance for Base-Level classifications

L_2-norm, here one gets nice improvements to ABSD).

In the following we compare our methods to the state-of-the-art methods. All the reference results in Table 3 are obtained from (10). Additionally we give the feature size and how the methods obtain their invariances in Table 3. The methods below are used for comparison:

(1) **Gaussian Euclidean Distance Transform (GEDT).** A volumetric function, where the value at each voxel is the Gaussian of the Euclidean distance to the nearest point on the surface of the model. Before computing the similarity the translation of the model is normalized by its COG, and the rotation is normalized by diagonalizing its inertia tensor and sorting by the directional inertia. (22)

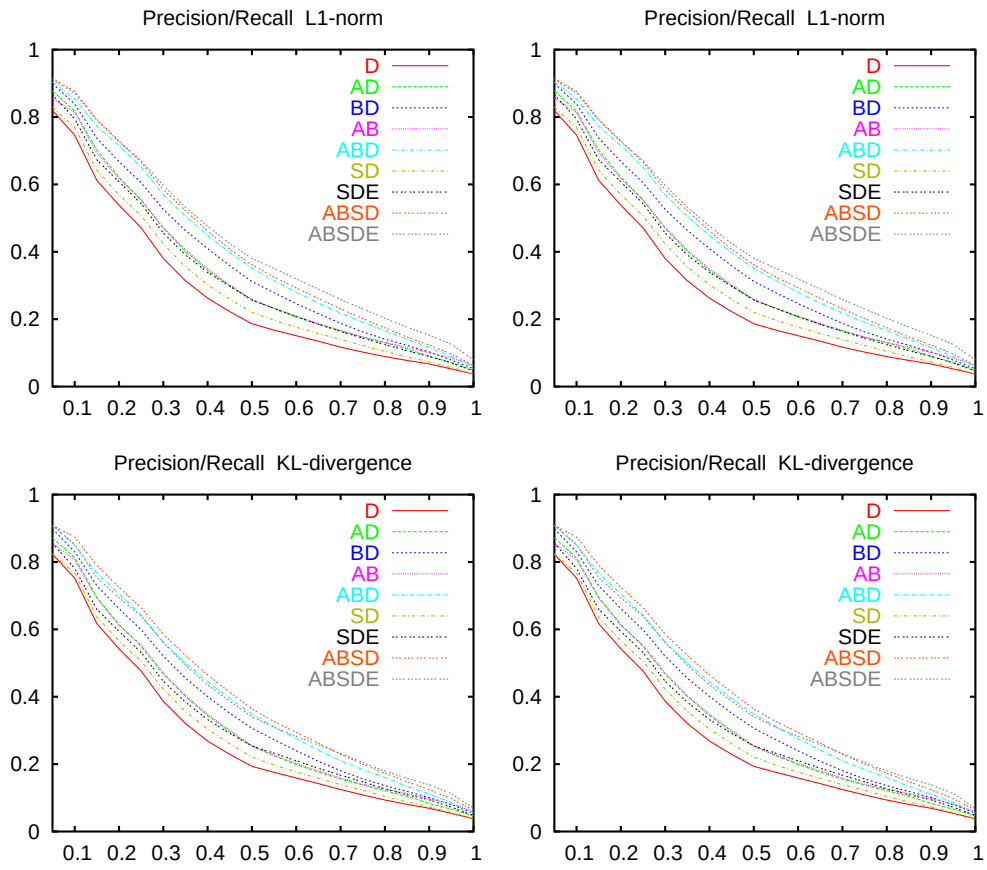

Fig. 3. PR-graphs corresponding to Table 2.

(2) **Spherical Harmonics of GEDT (SHD).** A rotation invariant description of the GEDT. Concentric spheres around the COG of the object are transformed by a Spherical Harmonic Transform and the norms of the SHT serve as a rotation invariant descriptor. (22)

(3) **Light Field Descriptor (LFD).** The model is rendered from different viewpoints and 2D descriptors are computed for each view. The distance between two objects is calculated by taking the minimum of all pairwise distances over all views. (13).

Having a look at Table 3 and Figure 4 one can see that our second-order methods can reach similar performances like SHD or GEDT. It is astonishing that our method, which keeps second-order information only , i.e. relative properties about two points averaged over the whole shape, give us similar results like GEDT, which nearly keeps the whole information about the shape. To the LFD-descriptors there is still a big gap, but LFD is based on a matching scheme, which makes them inappropriate for indexing schemes and partial queries and of course the similarity measure is expensive to compute.

Descriptor	size	NN	1-T	2-T	EM	DCG	Rot.	Trans.
D	32	34.2	16.9	25.2	15.2	44.5	I	I
SDE	220	41.7	21.7	30.6	18.2	49.1	I	I
ABD	256	54.7	27.4	37.5	22.5	55.8	I	I
ABSD	768	58.0	29.2	39.7	23.5	57.3	I	I
ABSDE	6240	57.1	30.5	41.5	24.4	58.2	I	I
GEDT	16384	60.3	31.3	40.7	23.7	58.4	N	N
SHD	1092	55.6	30.9	41.1	24.1	58.4	I	N
LFD	2350	65.7	38.0	48.7	28.0	64.4	M	N

Table 3
Comparison to the best known methods on Base-Level classification. We always used χ^2-distance. In the last two columns it is denoted how rotation/translation invariances are obtained: by Matching (M), Normalization (N) or Group Integration (I).

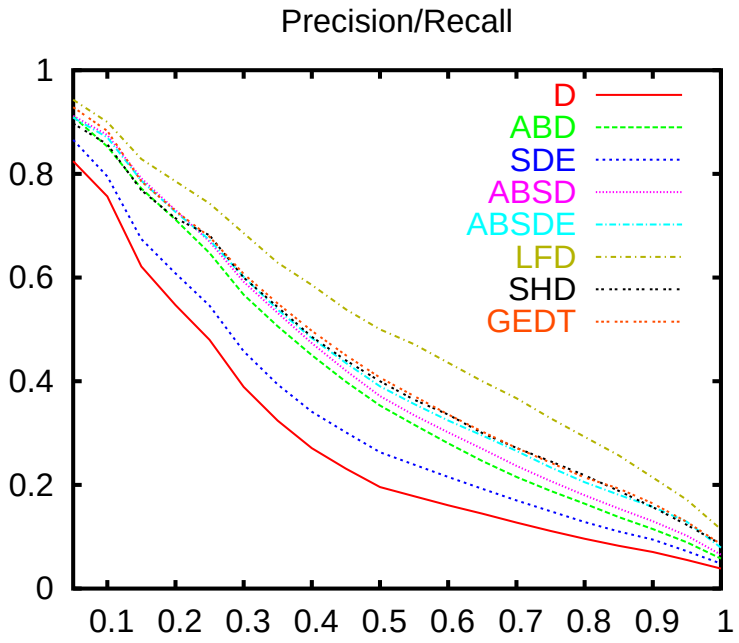

Fig. 4. PR-graph for comparison to the best known methods on Base-Level. The plot is partial copy from (10)

5 Conclusion and Future Work

We presented various extension of already proposed second-order shape features, which are invariant to Euclidean motion. We were able to show that their performance and discrimination ability is very good relative to the small

information content. Especially, the new BD- and ABD-histograms were introduced, which totally outperformed the already known AB-histograms. Further we used the Spherical Harmonic Transform to extend those features and give a general rule to compute rotation invariant features from the SHT-coefficients. The 'summit' of our exhaustive studies, the ABSDE-features, show similar retrieval accuracy like the SHD or GEDT, which both obtain their invariance to translation by normalization.

For future work it would be interesting to examine higher-order features. Indeed Osada has shown that D2 outperforms higher order variants, but there are still many other possible shape function, which are not considered yet. The tremendous growth in the number of features for the extended SHT-features is not satisfactory. Are there possible ways out, e.g. by sparse coding, or more clever similarity measures? Since group integration features do not have to cope with with correspondence problems, it is easy to adopt methods from machine learning to enhance retrieval performance. And finally, preliminary experiments with the ABSD-features on volume-data acquired from microscopy have led to promising results.

Acknowledgements

Many thanks to the anonymous reviewers for their useful suggestions and to the Princeton Shape Retrieval and Analysis Group for providing the PSB.

References

[1] H. Burkhardt, S.Siggelkow, Invariant features in pattern recognition - fundamentals and applications. In C. Kotropoulos and I. Pitas, editors, Nonlinear Model-Based Image/Video Processing and Analysis, John Wiley and Sons, 2001.

[2] O. Ronneberger, H. Burkhardt, E. Schultz, General-purpose Object Recognition in 3D Volume Data Sets using Gray-Scale Invariants, in: Proceedings of the International Conference on Pattern Recognition, Quebec, Canada, 2002.

[3] H. Schulz-Mirbach, Anwendung von Invarianzprinzipien zur Merkmalgewinnung in der Mustererkennung, Ph.D. thesis, TU Hamburg Harburg, reihe 10, Nr. 372, VDI-Verlag (Feb. 1995).

[4] H. Burkhardt, M.Reisert, H.Li, Invariants for discrete structures - an extension of haar integrals over transformation groups of dirac delta functions, in: Proceedings of the 26th DAGM Symposium, 2004.

[5] R. Osada, T. Funkhouser, B. C. und David Dobkin, Matching 3d models

with shape distribution, in: Proceedings Shape Modeling International, 2001.

[6] M. Reisert, The connection of shape distributions and haar integrals, Tech. Rep. 03/05, Albert-Ludwig University Freiburg *lmb.informatik.uni-freiburg.de/people/reisert* (2005).
URL `lmb.informatik.uni-freiburg.de`

[7] T. Sebastian, P. Klein, B. Kimia, Shock-based indexing into large shape databases, in: Proceedings of ECCV, 2002, pp. 731–746.

[8] J.Vleugels, R.Veltkamp, Efficent image retrieval through vantage objects, Pattern Recognition 35(1) (2002) 69–80.

[9] R. Ohbuchi, T. Minamitani, T. Takei, Shape-similarity search of 3d models by using enhanced shape functions, International Journal of Computer Applications in Technology (IJCAT) 23 (2005) 70–85.

[10] P. Shilane, P. Min, M. Kazhdan, T. Funkhouser, The princeton shape benchmark, in: Shape Modeling International, Genova, Italy, 2004.

[11] P. Min, J. Halderman, T. F. M. Kazhdan, Early experiences with a 3d model search engine, in: Proc. Web3D Symposium, pages 7-18, 2003.

[12] D. V. Vranic, D. Saupe, J. Richter, Tools for 3d-object retrieval: Karhunen-loeve transform and spherical harmonics, in: Proceedings of the IEEE Workshop Multimedia Signal Processing, 2001.

[13] D. Chen, X. Tian, Y. Shen, M. Ouhyoung, On visual similarity based 3d model retrieval, in: Computer Graphics Forum, Vol. 22/3, 2003.

[14] M. Suziki, A search engine for polygonal models to support development of 3d e-learning applications, in: 10th International WWW conference Posert Proceedings, 2001.

[15] J. Tangelder, R. Veltkamp, Polyhedral model retrieval using weighted point sets, International Journal of Image and Graphics 3(1) (2003) 209–229.

[16] J. Tangelder, R. Veltkamp, A survey of content based 3d shape retrieval, in: Proceedings Shape Modeling International, 2004, pp. 145–156.

[17] R. Veltkamp, Shape matching: Similarity measure and algorithms, in: Proceedings of Shape Modelling International, 2001, pp. 188–197.

[18] N. Iyer, S. Jayanti, K. Lou, Y. Kalyanaraman, K. Ramani, Three-dimensional shape searching: state-of-the-art review and future trends, Journal Computer-Aided Design 37(5) (2005) 509–530.

[19] B. Bustos, D. Keim, D. Saupe, T. Schreck, D.Vranic, An experimental comparison of feature-based 3d retrieval methods, in: Second International Symposium on 3D Data Processing, Visualization, and Transmission (3DPVT'04) hessaloniki, Greece, 2004.

[20] C.Ip, W. Lapadat, L. Sieger, W. Regli, Using shape distributions to compare solid models, in: Solid Modelling, 2002.

[21] D. Vranic, D. Saupe, Description of 3d-shape using a complex function on the sphere (2002) 177–180.

[22] M. Kazhdan, T. Funkhouser, S. Rusinkiewicz, Rotation invariant spherical harmonic represenation of 3d shape descriptors, in: Symposium on

Geometry Processing, 2003.
[23] M. Novotni, R.Klein, 3d zernike descriptors for content based shape retrieval, in: Solid Modelling, 2003.
[24] M. Schael, Invariant texture classification using group averaging with relational kernel functions, in: Texture 2002 the 2nd international workshop on texture analysis and synthesis, 2002, pp. 129–134.
[25] S. Siggelkow, H. Burkhardt, Improvement of histogram-based image retrieval and classification, in: Proceedings of the International Conference on Pattern Recognition, Vol. 3, 2002, pp. 367–370.
[26] B. Haasdonk, A. Halawani, H. Burkhardt, Adjustable invariant features by partial haar-integration, in: Proceedings of the 17th International Conference on Pattern Recognition (ICPR), Vol. 2, 2004, pp. 769–774.
[27] M. Kazhdan, Shape represenatations and algorithms for 3d model retrieval, Ph.D. thesis, Princenton University (April 2004).
[28] I. citations, chair http://mowat.virtualave.net/chair4.3ds
shuttle http://www.3dcafe.com/models/shuttle3.zip
piano http://64.198.60.83/baument/Models/grandpiano.zip
shark http://www.avatara.com/avatars2/greatwhite.wrl
microscope http://www.3dcafe.com/models/microscope.zip
cow http://www.sdsu.edu/ boyns/vrml/demo/cow.wrl.