

Using Irreducible Group Representations for Invariant 3D Shape Description

Marco Reisert and Hans Burkhardt

University of Freiburg, Computer Science Department,
79110 Freiburg i.Br., Germany
`{reisert}@informatik.uni-freiburg.de`

Abstract. Invariant feature representations for 3D objects are one of the basic needs in 3D object retrieval and classification. One tool to obtain rotation invariance are Spherical Harmonics, which are an orthogonal basis for the functions defined on the 2-sphere. We show that the irreducible representations of the 3D rotation group, which acts on the Spherical Harmonic representation, can give more information about the considered object than the Spherical Harmonic expansion itself. We embed our new feature extraction methods in the group integration framework and show experiments for 3D-surface data (Princeton Shape Benchmark).

1 Introduction

In many fields researchers deal with a huge amount of three dimensional data. In medical and biological applications one usually has to do with volumetric scans of various types. There is a need for fast and reliable feature extraction methods to handle and classify such huge amounts of data. The development of 3D modeling software has increased the number of freely available 3D-surface models, fast retrieval systems are necessary to browse and search for 3D-models in a user-friendly way. As the representation of 3D objects is not canonical and objects often occur at different spatial position and in different rotational poses, the question arises how to compare and classify the objects. One way is to use invariant features.

There are basically two ways to obtain invariance: Group integration and Normalization techniques. Normalization techniques obtain invariance by computing features relative to a global reference frame. The determination of the reference frame makes Normalization techniques extremely sensitive to noise. Whereas Group Integration (GI) is known to be very robust to many kinds of noise. In [1] a detailed overview over GI-techniques is given.

In this work we want to concentrate on 3D surface models. The Princeton Shape Benchmark (PSB) [9] offers a possibility to evaluate 3D-feature extraction techniques. It consists of approx. 1800 surface models collected from the Web. There is already a huge amount of work concerning feature extraction for 3D surface models by the use of Spherical Harmonics (SH). Vranic et al. [10] compute

a so-called spherical extent function of the model-surface and make a spherical harmonic transform of this function, but the rotational invariance is obtained by normalization. Kazhdan et al. [4] eliminate the rotational dependency by taking the magnitude of the invariant subspaces of the Spherical Harmonic transform. They show, that in most cases this is the better alternative than a normalizing approach. In [8] Reiser et al. enhanced the Shape Distribution introduced by Osada [6] by SH-expansion. The currently best performing methods on the PSB are the so-called Light Field Descriptors (LFD) [2]. LFD is an appearance based methods. The shape is rendered from several views and features for each view are computed. Two models are compared by searching the best matching pair of views.

The work is organized as follows: in Section 2 we introduce the basic algebra concerning rotations in 3D and introduce the so called *Wigner* D-matrices, the irreducible representation of the 3D rotation group. Further we give a relation for fast computation. In Section 3 we shortly review the group integration framework and show how the Wigner D-matrices can be used to enhance invariant features and keep more discriminative power. Then in Section 4 we show how this can be applied to extract features from 3D-surface models and show in Section 5 experiments on the Princeton Shape Benchmark. Finally we give a conclusion and an outlook for future work.

2 Life in $SO(3)$

First, some preliminaries about the notation. We always assume complex-valued vector spaces. Finite dimensional vectors \mathbf{x} are printed bold face, their i th component x_i in normal face. Transposition is denoted by \mathbf{x}^T , so scalar products are written by $\mathbf{x}^T \mathbf{x}'$, complex conjugation by x_i^* . Functions X (infinite dimensional vectors) are printed in capital letters and scalar products between functions are denoted by $\langle X|X' \rangle$.

2.1 Spherical Harmonics

The Spherical Harmonic expansion is a basic tool for 3D shape representation. We first want to give a short review about its basic properties. The complex functions defined on the two-sphere S^2 form a Hilbert-space with the inner-product $\langle X|X' \rangle = \int_{S^2} X^*(\mathbf{s}) X'(\mathbf{s}) d\mathbf{s}$, where $d\mathbf{s}$ denotes the natural measure on S^2 and $\mathbf{s} \in S^2$ is some unit-vector on the two-sphere. The Spherical Harmonics form an orthonormal basis in this space. They are commonly denoted by $Y_m^l(\mathbf{s})$, where $l \geq 0$ is a spectral index and $-l \leq m \leq l$. Basically the Y_m^l are polynomials of degree l in the normalized coordinates $\mathbf{s} = (x, y, z)^T$. We can expand any function X in our Hilbert space in terms of Spherical Harmonics where the expansion coefficients are simply the projections on the basis functions $\mathbf{a}^l = \langle \mathbf{Y}^l|X \rangle$, where this abbreviates $a_m^l = \langle Y_m^l|X \rangle$. The main property of a Spherical Harmonic expansion is its behavior under rotations. Let $g \in SO(3)$ an element of the rotation group acting on functions X by $(gX)(\mathbf{s}) \mapsto X(\mathbf{R}^T \mathbf{s})$, where

$\mathbf{R} \in \mathbb{R}^{3 \times 3}$ is the corresponding rotation matrix, then the expansion coefficients have the following property

$$\sum_{m'=-l}^{m'=l} D_{mm'}^l(g) a_{m'}^l = \langle Y_m^l | gX \rangle,$$

or shortly $\mathbf{D}^l(g)\mathbf{a}^l = \langle \mathbf{Y}^l | gX \rangle$. Note the close relation $D_{m0}^l = Y_m^l$. The $\mathbf{D}^l(g)$ are unitary transformation matrices depending on the rotation g . Note that only coefficients a_m^l with the same index l mix with each other, i.e. the subspaces with fixed index l stay invariant. This property is often used to obtain invariance against rotations. Due to the unitarity of $\mathbf{D}^l(g)$ the energy within a subspace is preserved. One can easily obtain invariant features of the function X by taking the magnitudes $\|\mathbf{a}^l\| = \sqrt{\sum_{m=-l}^{m=l} |a_m^l|^2}$, which are sometimes called SH-descriptors, analog to the Fourier-Descriptors in 2D.

2.2 Wigner D-Matrix

Let us have a closer look on the $\mathbf{D}^l(g)$ itself. For $l = 1$ there is a close relation to the real-valued ordinary rotation matrix \mathbf{R} by a special linear unitary transformation \mathbf{U} , i.e. $\mathbf{D}^1(g) = \mathbf{U}^T \mathbf{R} \mathbf{U}$. The general $\mathbf{D}^l(g)$ are called the *Wigner D-matrices* and they are the irreducible representations of the three dimensional rotation group. Irreducibility means that there is no further linear decomposition of the a_m^l such that the corresponding subspaces do not mix up during rotations (for references concerning the related group theory see e.g. [5] and references therein) The irreducibility has several important consequences: in fact, the $\mathbf{D}^l(g)$ are an orthogonal basis for the functions defined on the rotation group itself. Before going into detail let us introduce some basics. We obtain a Hilbert-space whose elements are the functions defined on $SO(3)$ by introducing an inner product via the group integral. Let $Z, Z' : SO(3) \mapsto \mathbb{C}$ two functions defined on the rotation group, then

$$\langle Z | Z' \rangle = \int_{SO(3)} Z^*(g) Z'(g) dg$$

defines a regular inner product, where dg is the natural group measure on $SO(3)$, which is left- and right-invariant (in Euler angles $dg = d\psi d\varphi \sin \theta d\theta$). Now we are able to state the orthogonality relation for the irreducible representations

$$\langle D_{m'_1 m_1}^{l_1} | D_{m'_2 m_2}^{l_2} \rangle = \delta_{l_1 l_2} \delta_{m'_1 m'_2} \delta_{m_1 m_2} \frac{8\pi^2}{2l_1 + 1},$$

i.e. any components of the representation matrices $\mathbf{D}^l(g)$ are orthogonal with respect to the given inner product. Now, given a function Z we can expand this function in terms of Wigner D-matrices as follows

$$Z(g) = \sum_{l=0}^{\infty} \sum_{m=-l}^l \sum_{m'=-l}^l b_{m'm}^l D_{m'm}^l(g),$$

where the \mathbf{b}^l are expansion-'matrices' obtained by the projections on the basis-functions $b_{m'm}^l = \frac{2l+1}{8\pi^2} \langle D_{m'm}^l | Z \rangle$.

Now the question arises how the $\mathbf{D}^l(g)$ look like explicitly. There is an expression involving the Jacobi-polynomials depending on the Euler-angles corresponding to g . The direct computation of this expression is a little bit cumbersome and of high computational complexity. Moreover the parameterization in Euler-angles is not advantageous for our purposes. We need a formulation in terms of the corresponding three-dimensional rotation matrix itself. In the following we want to point out the alternative way.

2.3 Product relations and Clebsch-Gordan coefficients

As we know that any function on the sphere can be expanded in terms of Spherical Harmonics, then also products of two Spherical Harmonics must have such a representation. The corresponding expansion coefficients are the so called Clebsch-Gordan coefficients.

$$Y_{m_1}^{l_1}(\mathbf{s})Y_{m_2}^{l_2}(\mathbf{s}) = \sum_{l=0}^{l_1+l_2} \sum_{m=-l}^{m=l} \langle lm|l_1m_1l_2m_2\rangle Y_m^l(\mathbf{s}) \quad (1)$$

The Clebsch-Gordan coefficients $\langle lm|l_1m_1l_2m_2\rangle$ fulfill two selection rules. They only give a contribution when $|l_1-l_2| \leq l \leq l_1+l_2$ and $m = m_1+m_2$. Additionally the Clebsch-Gordan coefficients themselves fulfill several orthogonality relation by what we can reformulate equation (1) as follows

$$Y_m^l(\mathbf{s}) = \sum_{m_1=-l_1}^{m_1=l_1} \sum_{m_2=-l_2}^{m_2=l_2} \langle lm|l_1m_1l_2m_2\rangle Y_{m_1}^{l_1}(\mathbf{s})Y_{m_2}^{l_2}(\mathbf{s}),$$

where l_1 and l_2 have to be chosen such that $l = l_1 + l_2$ due to the selection rules of the Clebsch-Gordan coefficients. By choosing $l_1 = l - 1$ and $l_2 = 1$ we have an iterative way to compute the Spherical Harmonics. The computation of the $\mathbf{Y}^l(\mathbf{s})$ only involves $\mathbf{Y}^{l-1}(\mathbf{s})$ and $\mathbf{Y}^1(\mathbf{s})$. For a fast implementation the Clebsch-Gordan coefficients can be precomputed and stored in a lookup-table. Considering the selection rules the actual algorithm to compute the Spherical Harmonics is a convolution-like method. In fact, the overall computational complexity is linear in the number of coefficients to be computed.

The computation of the Wigner D-matrices is very much the same as above. Products of Wigner D-matrix elements show nearly the same behavior, but needing products of Clebsch-Gordan coefficients. The basis for the iteration is now

$$D_{m'm}^l(g) = \sum_{\substack{m'_1=-l_1 \\ m_1=-l_1}}^{\substack{m'_1=l_1 \\ m_1=l_1}} \sum_{\substack{m'_2=-l_2 \\ m_2=-l_2}}^{\substack{m'_2=l_2 \\ m_2=l_2}} \langle m'|l_1m'_1l_2m'_2\rangle \langle lm|l_1m_1l_2m_2\rangle D_{m'_1m_1}^{l_1}(g)D_{m'_2m_2}^{l_2}(g),$$

and we again choose $l_1 = l - 1$ and $l_2 = 1$. The computational complexity is again linear in the number of computed matrix elements.

3 Group Integration Features

Group Integration is a well known tool to gain invariance for object representations. Suppose a given 3D-object X which has to be represented invariant against rotations. The easiest way to get an invariant feature is to extract some non-invariant non-linear 'kernel'-feature $f(X)$ and sum up this feature for all rotational poses

$$I_f(X) = \int_{SO(3)} f(gX) dg$$

Due to the integration lots of information about the original object is getting lost. In previous work [7] we have shown how to use Spherical Harmonics to keep more discriminative power. We will shortly review how this was done. The integral over a 3D rotation can be decomposed into an integration over the sphere S^2 and the circle:

$$I_f(X) = \int_{S^2} \int_0^{2\pi} f(g_{\mathbf{s},\varphi}X) d\varphi d\mathbf{s},$$

where $g_{\mathbf{s},\varphi}$ is a rotation around the axis \mathbf{s} with angle φ . Now, instead of just integrating over the sphere we projected the inner integral $F_X(\mathbf{s}) = \int_0^{2\pi} f(g_{\mathbf{s},\varphi}X) d\varphi$ on the $\mathbf{Y}^l(\mathbf{s})$ by $\mathbf{I}_f^l(X) = \langle \mathbf{Y}^l | F_X \rangle$ and obtained invariance by taking the band-wise magnitudes $\|\mathbf{I}_f^l(X)\|$. Doing this we already retained much information, but we still lose information about the ϕ -angle.

The Wigner D-matrices offer a much more natural way to extend the discriminative power of the features. Instead of an artificial decomposition of the integral we can simply project the function $F_X(g) := f(gX)$ on the irreducible group representation itself

$$\mathbf{I}_f^l(X) = \langle F_X | \mathbf{D}^l \rangle = \int_{SO(3)} f^*(gX) \mathbf{D}^l(g) dg \quad (2)$$

But what effect does have a rotation of the object $X \mapsto g'X$ on the $\mathbf{I}_f^l(X)$? Since the $\mathbf{D}^l(g)$ are unitary representations of the rotation group, i.e. $\mathbf{D}^l(gg') = \mathbf{D}^l(g)\mathbf{D}^l(g')$, we can show that

$$\mathbf{I}_f^l(gX) = \mathbf{I}_f^l(X) \mathbf{D}^l(g^{-1}),$$

i.e. a rotation of the object X leads to right-multiplication of the features matrices \mathbf{I}_f^l with \mathbf{D}^l . This means that the magnitudes of the columns are preserved during rotations and hence form an invariant feature. The final invariant features we use look as follows

$$\mathbf{I}_f^{l,m}(X) = \sqrt{\sum_{m'=-l}^{m'=l} |I_{f,mm'}^l(X)|^2},$$

where $I_{f,mm'}^l(X)$ denote the components of the feature matrices $\mathbf{I}_f^l(X)$.

4 Application to Surface Models

In [8] the group integration framework was already applied to surface models. The surface model is represented by a function $\mathbf{X} : \mathbb{R}^3 \mapsto \mathbb{R}^3$ giving only contribution on the surface \mathcal{S} of the shape, where the function value on the surface is given by the surface-normal at this position.

4.1 First-Order

To obtain translation invariance we shift the origin of the coordinate system into the center of gravity of the shape \mathbf{X} , i.e. we first use a simple normalization approach. For the group integration we will use the following kernel-function

$$f_{\mathbf{h},\mathbf{r}}(\mathbf{X}) = \delta(\mathbf{h}^T \mathbf{X}(\mathbf{r}) - 1), \quad (3)$$

where $\mathbf{r}, \mathbf{h} \in \mathbb{R}^3$ are parameters with $\|\mathbf{h}\| = 1$ and δ is the Delta-Distribution. The kernel-function gives contribution whenever at position \mathbf{r} the surface is present and its normal is parallel to \mathbf{h} . Inserting the kernel into (2) gives

$$\begin{aligned} \mathbf{I}_{\mathbf{h},\mathbf{r}}^l(\mathbf{X}) &= \int_{SO(3)} \delta(\mathbf{h}^T \mathbf{R} \mathbf{X}(\mathbf{R}^T \mathbf{r}) - 1) \mathbf{D}^l(g) dg \\ &= \int_{SO(3)} \int_{\mathbb{R}^3} \delta(\mathbf{h}^T \mathbf{R} \mathbf{X}(\mathbf{r}') - 1) \delta(\mathbf{R}^T \mathbf{r} - \mathbf{r}') \mathbf{D}^l(g) dg d\mathbf{r}'. \end{aligned}$$

We see that the integral only gives contribution, whenever $\|\mathbf{r}\| = \|\mathbf{r}'\|$ and $\mathbf{r}' \parallel \mathbf{R}^T \mathbf{r}$ and $\mathbf{h} \parallel \mathbf{R} \mathbf{X}(\mathbf{r}')$. The last two conditions are only satisfiable, if $\mathbf{r}^T \mathbf{h} = \mathbf{r}'^T \mathbf{X}(\mathbf{r}')$. If they are satisfied they determine the rotation matrix \mathbf{R} uniquely. So the group integral disappears and only the \mathbf{r}' -integral is left over. As the function \mathbf{X} gives contribution on the surface of the shape, the volume-integral is actually a surface-integral and we arrive at

$$\mathbf{I}_{\mathbf{h},\mathbf{r}}^l(\mathbf{X}) = \int_{\mathbf{r}' \in \mathcal{S}} \delta(\|\mathbf{r}'\| - r) \delta\left(\frac{\mathbf{r}'^T \mathbf{X}(\mathbf{r}')}{\|\mathbf{r}'\|} - \alpha\right) \mathbf{D}^l(g^*) d\mathbf{r}', \quad (4)$$

where $r = \|\mathbf{r}\|$ and $\alpha = \frac{\mathbf{r}^T \mathbf{h}}{\|\mathbf{r}\|}$ (for illustration see Figure 1). The g^* denotes the rotation which turns $\mathbf{r}', \mathbf{X}(\mathbf{r}')$ into \mathbf{r}, \mathbf{h} . One can see that g^* is the only part that depends on the actual values of \mathbf{r} and \mathbf{h} , the rest on their relative directions and the length of \mathbf{r} . A joint rotation of the two parameters leads to a left-multiplication of $\mathbf{I}_{\mathbf{r},\mathbf{h}}^l$ with \mathbf{D}^l . As this is only a unitary transformation of the features, we can restrict us for computation to one standard pose of \mathbf{h} and \mathbf{r} and hence the features depend on r and α only. But the actual invariant features depend on the actual choice of the standard pose. Following expression (4) for the integral the algorithm looks as follows:

Start with result array $\mathbf{I}_{r,\alpha}^l$ initialized with zeros.

For all points \mathbf{r}' on the surface of the object

 Compute $\alpha = \frac{\mathbf{r}'^T \mathbf{X}(\mathbf{r}')}{\|\mathbf{r}'\|}$, $r = \|\mathbf{r}'\|$

 Compute g^* , which turns $\mathbf{r}', \mathbf{X}(\mathbf{r}')$ into $\mathbf{r}_{\text{norm}}, \mathbf{h}_{\text{norm}}$

 Update $\mathbf{I}_{r,\alpha}^l = \mathbf{I}_{r,\alpha}^l + \mathbf{D}^l(g^*)$ for all $l \leq l_{\text{max}}$

The vectors $\mathbf{r}_{\text{norm}}, \mathbf{h}_{\text{norm}}$ are the parameter-vectors in normalized pose such that $\alpha = \frac{\mathbf{r}_{\text{norm}}^T \mathbf{h}_{\text{norm}}}{\|\mathbf{r}_{\text{norm}}\|}$ and $\|\mathbf{r}_{\text{norm}}\| = r$. Actually we compute for $l = 0$ something like a histogram. We count how often a point on the surface of the shape occur within a distance r to the COG and with angle $\arccos(\alpha)$ between the surface-normal and the vector connecting the point with the COG.

4.2 Second-Order

Now we assume an unnormalized model and obtain translation invariance by group-integration. We first perform a group integration over the three dimensional translation group \mathcal{T} and treat the result as above by projecting it on the irreducible representations. Equation (2) becomes

$$\mathbf{I}_l(\mathbf{X}) = \int_{SO(3)} \left(\int_{\mathcal{T}} f^*(\tau g \mathbf{X}) d\tau \right) \mathbf{D}^l(g) dg.$$

In this case the kernel-function (3) is too simple, the results would be indiscriminative. We need a more complex kernel-function. A simple generalization of (3) is

$$f_{\mathbf{h}, \mathbf{h}', \mathbf{d}}(\mathbf{X}) = \delta(\mathbf{h}^T \mathbf{X}(0) - 1) \delta(\mathbf{h}'^T \mathbf{X}(\mathbf{d}) - 1).$$

We do not want to give the complete derivation again since it is very much the same like in the first case, so we outline the resulting algorithm directly.

Start with result array $\mathbf{I}_{d, \alpha, \beta, \gamma}^l$ initialized with zeros.

For all pairs of points $\mathbf{r}_1, \mathbf{r}_2$ on the surface of the object

Let $\mathbf{d} = \mathbf{r}_1 - \mathbf{r}_2$ and $d = \|\mathbf{d}\|$

Compute $\alpha = \mathbf{d}^T \mathbf{X}(\mathbf{r}_1)/d$, $\beta = \mathbf{d}^T \mathbf{X}(\mathbf{r}_2)/d$

Compute $\gamma = \frac{\mathbf{P}_d \mathbf{X}(\mathbf{r}_1)^T \mathbf{P}_d \mathbf{X}(\mathbf{r}_2)}{\|\mathbf{P}_d \mathbf{X}(\mathbf{r}_1)\| \|\mathbf{P}_d \mathbf{X}(\mathbf{r}_2)\|}$

Compute g^* , which turns $\mathbf{d}, \mathbf{X}(\mathbf{r}_1), \mathbf{X}(\mathbf{r}_2)$ into $\mathbf{d}_{\text{norm}}, \mathbf{h}_{\text{norm}}, \mathbf{h}'_{\text{norm}}$

Update $\mathbf{I}_{d, \alpha, \beta, \gamma}^l = \mathbf{I}_{d, \alpha, \beta, \gamma}^l + \mathbf{D}^l(g^*)$ for all $l \leq l_{\text{max}}$

The matrix \mathbf{P}_d denotes a projection matrix, projecting on the plane given by the vector \mathbf{d} . For illustration of the variables see Figure 1. Again $\mathbf{d}_{\text{norm}}, \mathbf{h}_{\text{norm}}, \mathbf{h}'_{\text{norm}}$ are the parameters in standard pose. For $l = 0$ the feature can also be interpreted as a histogram, the frequency of occurrences of two points within a specific distance d and a specific surface-normal configuration determined by the parameters α, β, γ .

4.3 Implementation details

For evaluation of our features we use the Princeton Shape Benchmark [9]. It consists of approx. 1800 polygonal mesh models. We represent the models in 256^3 voxel grid. Additionally, each voxel gets a reference to the triangle it is stemming from to enable us to incorporate the original surface normals in our calculations. To exclude non visible constructional artefacts inside the closed surface we use a floodfill operation.

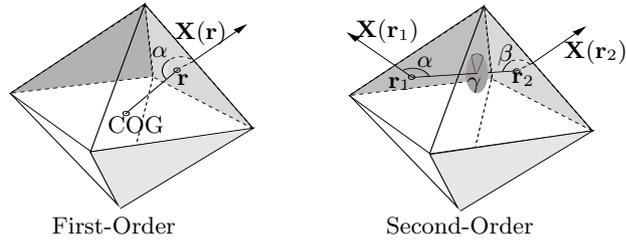


Fig. 1. Explanation of variables. For first-order just the length of \mathbf{r} and the angle between \mathbf{r} and $\mathbf{X}(\mathbf{r})$ are relevant. For second-order we have four parameters, α, β, γ and distance between \mathbf{r}_1 and \mathbf{r}_2 . The γ -angle is the angle between $\mathbf{X}(\mathbf{r}_1)$ and $\mathbf{X}(\mathbf{r}_2)$ after projection onto the plane orthogonal to $\mathbf{r}_1 - \mathbf{r}_2$.

The implementation of the first-order features is straight-forward, because each voxel has to be accessed only once. The time for a computation of one feature set is less than a second (*Pentium4* 2.8Ghz), where most of the time is spend on voxelization. For the second-order features the complexity is quadratic in the number of voxels. Typical models consist of several tens of thousand voxels, hence an exhaustive computation is not possible. We use a randomized approach to keep the computation time low. Computation times are varying from one to several seconds, depending on the accuracy of the computation and the number of Wigner D-matrix coefficients.

For the computation of Clebsch-Gordan coefficients we use *Matpack*. Due to the precomputation of the coefficients the running time does not depend on their implementation. We tried several discretizations of the parameters r, α and d, α, β, γ . In the experiments we always give the results for the best quantizations, which always depend on the type of feature. To obtain invariance we take the magnitudes of the rows of the 'matrix'-features \mathbf{I}^l . As already mentioned the features depend on the absolute pose of the parameter-vectors \mathbf{r}_{norm} and \mathbf{h}_{norm} ($\mathbf{d}_{\text{norm}}, \mathbf{h}_{\text{norm}}, \mathbf{h}'_{\text{norm}}$ for second-order). But which pose one should choose? The most simple representation $\mathbf{h}_{\text{norm}} = (\alpha, \sqrt{1-\alpha^2}, 0)^T$ and $\mathbf{r}_{\text{norm}} = (r, 0, 0)^T$ has the disadvantage that the magnitude of the rows of \mathbf{D}^l are the same for m with the same absolute value. We have found that the complex representation $\mathbf{h}_{\text{norm}} = (\sqrt{\frac{1-\alpha^2}{2}}e^{i\frac{\pi}{4}}, \alpha, \sqrt{\frac{1-\alpha^2}{2}}e^{-i\frac{\pi}{4}})^T$ and $\mathbf{r}_{\text{norm}} = (0, r, 0)^T$ keeps the magnitudes of the rows of \mathbf{D}^l more independent. And it actually performs better than the first simple one.

5 Experiments

In order to keep the results comparable to experiments given in [9] we conducted our experiments only on the test set of the PSB at the finest granularity. To show the superiority of the Wigner D-matrices over the Spherical Harmonics we also give the results for the best corresponding SH-feature (see [8]). As distance measure between the features we use the L_1 -norm. In Table 1 the results for the first- and second-order features are shown. For a description of the used

performance measures Nearest-Neighbor/1-Tier/2-Tier/E-Mesure/Discounted-Cumulative-Gain see [9]. Experiments were made for different cutoff indices l_{\max} . We found that higher cutoff indices than 4 show only marginal improvements if at all.

Table 1. Results for finest granularity on the PSB. All results are given in percent. Higher rates mean better performance. (DW) stands for our new approach followed by l_{\max} the cutoff-band. (SH) for the old Spherical Harmonic approach. Best results in bold face.

First-Order						Second-Order					
method	NN	1T	2T	EM	DCG	method	NN	1T	2T	EM	DCG
DW 1	48.3	24.6	34.4	19.8	52.3	DW 1	60.5	31.6	42.2	24.5	59.3
DW 2	53.4	28.2	38.2	21.9	55.4	DW 2	62.5	32.7	43.9	25.5	60.1
DW 3	54.0	28.3	38.0	22.0	55.5	DW 3	63.5	32.7	43.8	25.5	60.2
DW 4	55.3	29.1	38.5	22.5	56.1	DW 4	63.5	32.9	44.0	25.5	60.5
SH 4	50.1	26.1	36.1	21.0	53.6	SH 4	60.6	31.5	42.9	24.7	59.4

We see that the Wigner expansion is superior to SH expansion in all cases, which gives evidence that the proposed descriptors can, in fact, carry more information and discriminative power than the pure Spherical Harmonic approach. Of course, in this comparison, the number of features is larger in the Wigner case, because for each l we have $2l + 1$ features instead of only one. The absolute number is also rather high, because the quantization setting for the Wigner case is 32, 2, 8, 8, (for d, γ, α, β) resulting in several thousand of features. So we additionally tried to reduce the number of features by feature selection using *Simba* [3]. In Figure 2 we show the performance results and precision/recall graph with 10% of the second-order features with $l_{\max} = 4$. In fact, feature selection improves the results very much, while reducing the number of features drastically. We also give results of two other shape descriptors, LFD and GEDT (for references see [9]). The PR-graph shows that our approach works with the same precision as the LFD approach, while recall is comparable to the GEDT method. It is astonishing that a group integration approach, which keeps second-order information only, i.e. relative properties about two points averaged over the whole shape, give us similar results like LFD which is basically a registration approach, or GEDT which keeps nearly the whole information about the shape.

6 Conclusion and Future Work

We presented how the irreducible representation of the 3D rotation group can be used for invariant shape representation. The Wigner expansion can be seen as the canonical generalization of the Spherical Harmonic expansion. We applied the Wigner expansion in the group integration framework and were able to show the superiority of them over SH-expansion in a shape retrieval task. The performance of the features is comparable to the currently best performing methods on

method	NN	1T	2T	EM	DCG
10% DW	65.7	34.8	46.0	26.7	62.1
LFD	65.7	38.0	48.7	28.0	64.3
GEDT	60.3	31.3	40.7	23.7	58.1

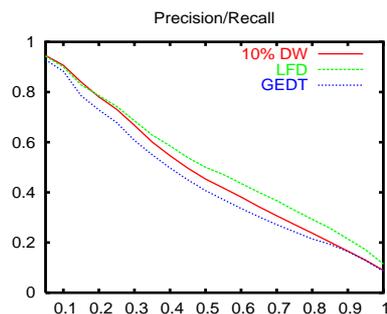


Fig. 2. Results for the second-order features with feature selection in comparison to LFD and GEDT. (For references see [9])

the Princeton Shape Benchmark, but our approach is 'single-feature'-method, while LFD is a 'compositions' of features. For future work we want to apply our methods for volume data. Further we want to examine other methods to obtain invariance instead of just taking the magnitude of the rows of the 'matrix'-features.

References

1. H. Burkhardt and S. Siggelkow. *Invariant features in pattern recognition - fundamentals and applications*. In *Nonlinear Model-Based Image/Video Processing and Analysis*. John Wiley and Sons, 2001.
2. D. Chen, X. Tian, Y. Shen, and M. Ouhyoung. On visual similarity based 3d model retrieval. In *Computer Graphics Forum*, volume 22/3, 2003.
3. R. Gilad-Bachrach, A. Navot, and N. Tishby. Margin based feature selection - theory and algorithms. In *Proc. of ICML04*, 2004.
4. M. Kazhdan, T. Funkhouser, and S. Rusinkiewicz. Rotation invariant spherical harmonic representation of 3d shape descriptors. In *Symposium on Geometry Processing*, 2003.
5. W. Miller. Topics in harmonic analysis with applications to radar and sonar. *IMA Volumes in Mathematics and its Applications*, 1991.
6. R. Osada, T. Funkhouser, and B. Chazelle und D. Dobkin. Matching 3d models with shape distribution. In *Proceedings Shape Modeling International*, 2001.
7. M. Reisert and H. Burkhardt. Invariant features for 3d-data based on group integration using directional information and spherical harmonic expansion. In *the Proceedings of the ICPR06*, 2006.
8. M. Reisert and H. Burkhardt. Second order 3d shape features: An exhaustive study. *C&G, Special Issue on Shape Reasoning and Understanding*, 30(2), 2006.
9. P. Shilane, P. Min, M. Kazhdan, and T. Funkhouser. The princeton shape benchmark. In *Shape Modeling International, Genova, Italy*, 2004.
10. D.V. Vranic, D. Saupe, and J. Richter. Tools for 3d-object retrieval: Karhunen-loeve transform and spherical harmonics. In *Proceedings of the IEEE Workshop Multimedia Signal Processing*, 2001.