3D Object Detection using a Fast Voxel-Wise Local Spherical Fourier Tensor Transformation

Henrik Skibbe[1,4], Marco Reisert[2], Thorsten Schmidt[1,4], Klaus Palme[3,4]
, Olaf Ronneberger[1,4], and Hans Burkhardt[1,4]

[1]Department of Computer Science, University of Freiburg, Germany
[2]Dept. of Diagnostic Radiology, Medical Physics, University Medical Center, Freiburg
[3]Institute of Biology II, Freiburg Institute for Advanced Studies (FRIAS)
[4]Center for Biological Signalling Studies (BIOSS), University of Freiburg
skibbe@informatik.uni-freiburg.de,marco.reisert@uniklinik-freiburg.de

Abstract. In this paper we present a novel approach for expanding spherical 3D-tensor fields of arbitrary order in terms of a tensor valued local Fourier basis. For an efficient implementation, a two step approach is suggested combined with the use of spherical derivatives. Based on this new transformation we conduct two experiments utilizing the spherical tensor algebra for computing and using rotation invariant features for object detection and classification. The first experiment covers the successful detection of non-spherical root cap cells of Arabidopsis root tips presented in volumetric microscopical recordings. The second experiment shows how to use these features for successfully detecting $\alpha-$helices in cryo-EM density maps of secondary protein structures, leading to very promising results.

1 Introduction

With the increasing performance of modern computers and the rapid development of new 3D image recording techniques, the amount of volumetric image data has drastically increased during the last years. Due to this fact, there is a need for adapting existing techniques from 2D image analysis to the third dimension. One major problem in image analysis is the extraction of information which is reduced in size as much as possible while still containing all characteristics necessary to describe, analyze, detect, compare or classify different objects. Many methods widely used for extracting features from 2D images make use of the gradient direction to get rid of the rotation. These approaches can often be directly adapted into the 2D+time domain (e.g. SIFT [1]). However, when working with volumetric images the gradient direction gives us only information about two rotation angles leaving the third angle undetermined.

In this work, we introduce a new method for realizing a fast voxel-wise local spherical Fourier transformation of spherical tensor-valued 3D images. The expansion coefficients are used to compute rotation invariant features in an analytical way. From a practical point of view, our method does exactly this: For each Gaussian windowed surrounding of each voxel of a tensor-valued volume, we

simultaneously compute the expansion coefficients of a spherical Fourier tensor expansion up to a certain band. Using spherical tensor algebra (the interested reader is referred to [2,3]), the resulting expansion coefficients can be combined to analytically form rotation invariant features. These voxel-wise features can be used for e.g. object detection, segmentation or object classification. The method proposed here needs a small number of image convolutions followed by successively applying point-wise operations that can run in parallel in a memory efficient way. Outperforming all existing methods we are aware of doing voxel-wise spherical harmonic expansions realized by a huge number of convolutions (e.g. [4]).

This paper is organized as follows: In section 2 we recapitulate the basics and requirements necessary for our mathematical framework. In section 3 we introduce the fast spherical Fourier tensor transformation. Finally, in section 4, two applications are introduced where local rotation invariant features are computed and used for successfully detecting and classifying objects.

2   Preliminaries

We denote scalars in unbold latin face and vectors in latin bold face. Typically, vectors are elements of $\mathbb{C}^{2\ell+1}$ whose basis is written as $\{\mathbf{e}_m^\ell\}_{m=-\ell\ldots\ell}$. Depending on the context we will express the coordinate vector $\mathbf{r} = (x,y,z)^T \in \mathbb{R}^3$ in spherical coordinates (θ, ϕ, r), where $\theta = \arccos(z/\|\mathbf{r}\|)$, $\phi = \operatorname{atan2}(y,x)$ and $r = \|\mathbf{r}\|$. By $Y_m^\ell(\theta, \phi)$ we denote the usual spherical harmonics [2] in Schmidt semi-normalized form. All harmonics for a fixed ℓ are arranged in a vector $\mathbf{Y}^\ell \in \mathbb{C}^{2\ell+1}$. The functions Y_m^ℓ build a complete orthogonal basis for representing functions on the 2-sphere, with $\langle Y_m^\ell, Y_{m'}^{\ell'} \rangle = \frac{4\pi}{2\ell+1}\delta_{\ell\ell'}\delta_{mm'}$. Furthermore, we denote by $\mathbf{R}^\ell : \mathbb{R}^3 \to \mathbb{C}^{2\ell+1}$ the commonly known solid harmonics [2], whose $(2\ell + 1)$ components R_m^ℓ are defined by $R_m^\ell(\mathbf{r}) = r^\ell Y_m^\ell(\theta, \phi)$.

2.1   Spherical Tensor Fields

In the following we give a short introduction on spherical tensor algebra based on the definitions and notation used in [5]. The spherical tensor algebra is necessary for expanding higher order tensor fields (e.g. vector fields) in terms of tensor-valued Fourier basis functions. We also use spherical tensor algebra for computing rotation invariant features in an analytical way [6].

The central role of spherical tensor algebra play the Wigner D-matrices $\mathbf{D}_g^\ell \in \mathbb{C}^{(2\ell+1)\times(2\ell+1)}$ which are the unitary irreducible representations of the 3D rotation group. Each D-matrix is associated with an element g of the rotation group. They behave like ordinary rotation matrices in the sense that $\mathbf{D}_g^\ell \mathbf{D}_h^\ell = \mathbf{D}_{gh}^\ell$, but act on the high-dimensional complex Hilbert space $\mathbb{C}^{2\ell+1}$. A fundamental property of the Wigner D-matrices is their behavior with respect to spherical harmonic expansion coeffcients. Suppose you have expanded some function f in spherical harmonics $\mathbf{a}^\ell = \langle \mathbf{Y}^\ell, f \rangle$. Then the expansion coefficients of the rotated function gf are related to the $\mathbf{a}^\ell$ just by the Wigner D-matrices

$\mathbf{a}^{\ell'} = \mathbf{D}_g^\ell \mathbf{a}^\ell$. Spherical Tensor algebra utilizes this behavior in a much more general way. We call a function $\mathbf{f} : \mathbb{R}^3 \to \mathbb{C}^{2\ell+1}$ a spherical tensor field of rank ℓ if it transforms with respect to rotations as

$$\forall g \in SO(3): \quad (g\mathbf{f})(\mathbf{r}) := \mathbf{D}_g^\ell \mathbf{f}(\mathbf{U}_g^T \mathbf{r}) \tag{1}$$

where $\mathbf{U}_g \in \mathbb{R}^{3\times3}$ is the corresponding real valued ordinary rotation matrix. The space of all spherical tensor fields of rank ℓ is denoted by $\mathcal{T}_\ell$. Note that for $\ell = 1$ a spherical tensor field is just an ordinary vector field. Interpreting solid harmonics (or spherical harmonics) as spherical tensor fields shows their importance, namely (due to eq. (1)) $g\mathbf{R}^\ell = \mathbf{R}^\ell$, i.e. they are 'fix' with respect to rotations.

Similar to Cartesian tensor fields, where Kronecker products connect tensor fields of different rank, there exist spherical products [5,2] that connect spherical tensor fields of different rank. In fact, for two given spherical tensor fields $\mathbf{v} \in \mathcal{T}_{\ell_1}$ and $\mathbf{w} \in \mathcal{T}_{\ell_2}$, there exists a whole set of different products $\circ_\ell$ to build new spherical tensor fields. More precisely, for every $\ell \geq 0$ that obeys the triangle inequality $|\ell_1 - \ell_2| \leq \ell \leq \ell_1 + \ell_2$ there is a bilinear form $\circ_\ell : \mathbb{C}^{2\ell_1+1} \times \mathbb{C}^{2\ell_2+1} \to \mathbb{C}^{2\ell+1}$ that takes two spherical tensors and gives a new one, i.e

$$(\mathbf{D}_g^{\ell_1}\mathbf{v}) \circ_\ell (\mathbf{D}_g^{\ell_2}\mathbf{w}) = \mathbf{D}_g^\ell(\mathbf{v} \circ_\ell \mathbf{w}) \tag{2}$$

holds for any $\mathbf{v} \in \mathcal{T}_{\ell_1}$ and $\mathbf{w} \in \mathcal{T}_{\ell_2}$. Again the spherical harmonics show a special behavior. In fact, multiplying two spherical harmonics results in another harmonic, i.e. $\mathbf{Y}^{\ell_1} \circ_\ell \mathbf{Y}^{\ell_2} = c_{\ell,\ell_1,\ell_2}\mathbf{Y}^\ell$, where c_{ℓ,ℓ_1,ℓ_2} is a constant related to the Clebsch Gordan coefficients (for details see [2]).

Finally, we present the third important ingredient of spherical tensor calculus: the spherical derivative. In ordinary vector calculus differential operators like the gradient, divergence or the Hessian connect Cartesian tensor fields of different rank. There are also spherical counterparts. In the following we need just one type of spherical derivative, the spherical up-derivative $\boldsymbol{\nabla}^1$, which increases the rank of the spherical field (see [7] for further details and proofs). If $\mathbf{f} \in \mathcal{T}_\ell$ is a tensor-field of order ℓ, then the spherical up-derivative $\boldsymbol{\nabla}^1 : \mathcal{T}_\ell \to \mathcal{T}_{\ell+1}$ maps $\mathbf{f}$ onto a field of rank $\ell + 1$. For multiple application of the spherical derivative we write $\boldsymbol{\nabla}^\ell : \mathcal{T}_{\ell_0} \to \mathcal{T}_{\ell_0+\ell}$. For example, applying $\boldsymbol{\nabla}^1$ on a scalar field gives just the spherical counterpart of an ordinary gradient of the field. The result of applying $\boldsymbol{\nabla}^1$ twice is linearly related to the traceless Hessian of the scalar field.

Relation between Cartesian and Spherical Tensors. The theories of Cartesian and spherical tensors are basically equivalent. Up to rank 2 the relations connecting both worlds are well known and reported e.g. in [2] or [8]. To get an impression; a general real Cartesian tensor of rank 2 (basically a 3×3 matrix) can be decomposed into a spherical tensor of rank 0 (the trace), of rank 1 (the antisymmetric part) and of rank 2 (the traceless symmetric part).

3   Spherical Fourier Tensor Transformation

In this section, we introduce the Fourier basis which we use for spherical tensor field expansion. This basis can be seen as an extension to [9], replacing the spherical harmonics by tensorial harmonics. With this we are able to decompose spherical tensor fields of any order into basic Fourier patterns.We further show how to expand fast and efficiently by utilizing the spherical derivatives.

3.1   Tensorial Bessel Harmonics

In order to obtain a representation of spherical tensor fields with tensor valued Fourier basis functions having the same convenient rotation properties known from the spherical harmonics, we perform a spherical tensor field expansion based on *tensorial harmonics* [5]. In addition to the tensorial harmonic expansion given in [5] for representing spherical tensor fields on the 2-sphere, we use the spherical Bessel function $j_\ell(r)$ for representing the radial part (see [10] for definition). This directly extends the spherical Fourier basis $\mathbf{B}_k^\ell(\mathbf{r}) := \mathbf{Y}^\ell(\mathbf{r})j_\ell(kr)$ presented in [9] to higher order tensor fields, where $k \in \mathbb{R}_{>0}$ represents the frequency in radial direction. The spherical Fourier tensor field expansion of $\mathbf{f} \in \mathcal{T}_J$ in terms of $\circ$ and the Fourier basis $\mathbf{B}_k^\ell$ is given by

$$\mathbf{f}(\mathbf{r}) = \int_0^\infty \sum_{\ell=0}^\infty \sum_{j=-J}^{j=J} \mathbf{c}_{jk}^\ell \circ_J (\alpha_{\ell jk}^{\frac{1}{2}} \mathbf{B}_k^\ell(\mathbf{r}))dk \ , \tag{3}$$

where $\alpha_{\ell jk} = \frac{2k^2}{\pi} \frac{2(\ell+j)+1}{2J+1} \frac{2\ell+1}{4\pi}$ are scalar valued normalization factors, and $\mathbf{c}_{jk}^\ell \in \mathbb{C}^{2(\ell+j)+1}$ are the spherical tensor valued expansion coefficients of $\mathbf{f}$.

   The expansion coefficients $\mathbf{c}_{jk}^\ell$ can be computed by directly projecting onto tensorial Bessel harmonics, with $c_{jkm}^\ell = \langle \mathbf{f}, \mathbf{e}_m^{\ell+j} \circ_J \mathbf{B}_k^\ell \rangle$. However, it would be quite expensive to do this voxel by voxel in a large volume. According to [8], we suggest to compute the expansion coefficients $\mathbf{c}_{jk}^\ell$ of $\mathbf{f} \in \mathcal{T}_\ell$ in two steps. First, we separately transform all $2J+1$ components of $\mathbf{f}$ into the harmonic domain, i.e. we express each component of f_M in terms of $\mathbf{B}_k^\ell$, with

$$f_M(\mathbf{r}) = \int_0^\infty \sum_{\ell=0}^\infty \mathbf{a}_k^{\ell M T} \mathbf{B}_k^\ell(\mathbf{r})dk \ , \tag{4}$$

where $\mathbf{a}_k^{\ell M}$ are the expansion coefficients representing the M-th component of $\mathbf{f}$. Given the coefficients $\mathbf{a}_k^{\ell M}$ we then compute the expansion coefficients $\mathbf{c}_{jk}^\ell$:

$$c_{jkm}^\ell = \frac{2(\ell+j)+1}{2J+1} \sum_M a_{kn}^{\ell M} \langle (\ell+j)m, \ell n \,|\, JM \rangle \ , \tag{5}$$

where $m = -2(\ell+j), \ldots, 2(\ell+j)$ and $n = M-m$. Until now we have not reached any computational benefit, but we show in the following section that we can use the spherical derivatives for the computation of the expansion coefficients $a_{kn}^{\ell M}$. This avoids explicit convolutions and pre-computation of convolution kernels, which makes the computation practical in terms of speed and memory consumption.

3.2   Differential Formulation of the Tensorial Bessel Expansion

In this section we show how to compute the expansion coefficients $\mathbf{a}_k^\ell$ (eq. (4)) with respect to all positions $\mathbf{x} \in \mathbb{R}^3$ of a given 3D image simultaneously by utilizing the spherical derivatives. As a consequence of having a fast transformation for computing the coefficients $\mathbf{a}_k^\ell$ we directly get a fast method for computing the tensorial Bessel harmonic coefficients $\mathbf{c}_{jk}^\ell$ (see eq. (5)).

One possible solution to transform efficiently is separately convolving a function $f \in \mathcal{T}_0$ with all $(2\ell+1)$ scalar valued components $B_{kn}^\ell \in \mathcal{T}_0$ of all $\mathbf{B}_k^\ell \in \mathcal{T}_\ell$. In our scenario, doing a transformation using K different values for k and an angular expansion up to the L-th band, we would need $K(L+1)^2$ convolutions! In contrast, the method proposed here makes use of an iterative differential formulation of the $\mathbf{B}_k^\ell$ for realizing the transformation in angular direction. Hence we only need K convolutions followed by L times applying the spherical derivative operator $\boldsymbol{\nabla}$. In this case, the number of convolutions does not depend on the parameter L. Furthermore, the spherical derivative operator can be implemented efficiently using finite differences, widely used for fast computing derivatives from scalar valued fields which can be executed in parallel in a memory efficient way. This enables us to process huge volumetric images in seconds rather than minutes (e.g. given an image of size 256^3 and doing the voxel-wise expansion up to order 15 (double precision) takes 1.3 minutes by convolutions using the multi-threaded fft [11] with planning flag FFTW_MEASURE and 14.1 seconds using our approach. The experiments are run on a 6×quad-core system, each core with 2,7 GHz.). The relation between the convolution based approach and our differential based approach is given by:

$$\mathbf{a}_k^\ell(\mathbf{x}) = \langle f_\mathbf{x}, \mathbf{B}_k^\ell \rangle = \underbrace{(f * \overline{\mathbf{B}}_k^\ell)(\mathbf{x})}_{\substack{(2\ell+1)\text{ scalar}\\ \text{valued convolutions}}} = \frac{(-1)^\ell}{k^\ell}\langle f_\mathbf{x}, \boldsymbol{\nabla}^\ell \mathbf{B}_k^0 \rangle = \left[\frac{\overline{\boldsymbol{\nabla}}^1}{(-k)}\right]^\ell \underbrace{(f * \overline{\mathbf{B}}_k^0)(\mathbf{x})}_{\substack{1\text{ scalar}\\ \text{valued convolution}}},$$

where $f_\mathbf{x}(\mathbf{r}) := f(\mathbf{r} + \mathbf{x})$ and $\overline{\boldsymbol{\nabla}}^\ell$ is the complex conjugate of $\boldsymbol{\nabla}^\ell$ (for proof see appendix A). As a result the expansion coefficients $\mathbf{a}_k^0, \ldots \mathbf{a}_k^\ell$ can be computed iteratively by ℓ times applying the spherical derivative operator:

$$\mathbf{a}_k^\ell(\mathbf{x}) = \underbrace{(\tfrac{-1}{k}\overline{\boldsymbol{\nabla}}^1 \ldots (\tfrac{-1}{k}\overline{\boldsymbol{\nabla}}^1 (\underbrace{\tfrac{-1}{k}\overline{\boldsymbol{\nabla}}^1 \underbrace{(f * \overline{\mathbf{B}}_k^0}_{= \mathbf{a}_k^0})}_{= \mathbf{a}_k^1})))(\mathbf{x})}_{= \mathbf{a}_k^\ell}, \tag{6}$$

We finally obtain the expansion coefficients $\mathbf{c}_{jk}^\ell(\mathbf{x})$ of the spherical tensor field by first doing a component-wise transformation of the tensor field (eq. (4)) utilizing the spherical derivatives (eq. (6)) followed by coupling the expansion coefficients according to eq. (5).
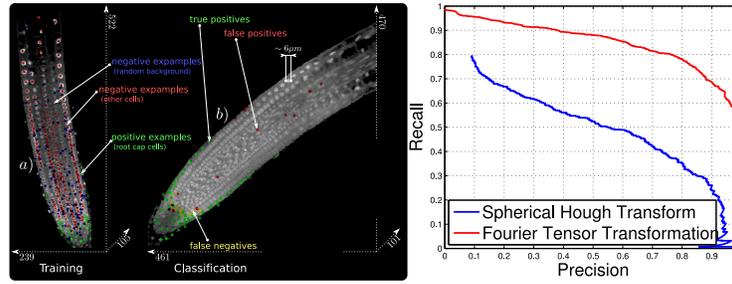
Fig. 1. Left image: a) Training dataset. b) Test dataset. Right PR-graph: Comparing the performance of our detection and classification to the detection rate of the spherical hough transformation.

4   Applications & Experiments

We conduct two experiments demonstrating the performance of our method. In the first experiment we detect nuclei of root cap cells of Arabidopsis root tips. In the second experiment we detect α-helices in secondary protein structures. Since we aim to describe local image structures and further need a finite convolution kernel, we use a Gaussian windowed Bessel function in both experiments. The Gaussian windowed convolution kernel is given by $(g_\sigma(r)j_\ell(kr))\mathbf{Y}^\ell(\theta,\phi) = \mathbf{B}_k^0(\mathbf{r})g_\sigma(r)$ where σ determines the width of the Gaussian window function $g_\sigma(r) = e^{-\frac{r^2}{\sigma}}$. Scaling is done by assuming different voxel sizes. Considering the Fourier transform of this function, the parameter k determines the distance of a spherical harmonic from the origin, while the parameter σ determines the size of the Gaussian with which the spherical harmonic is convolved.

There are several ways for obtaining rotation invariant features based on spherical tensor fields (see e.g. [3,9,6]). Similar to [6] we are utilizing the spherical tensor product (eq. (2)) for coupling expansion coefficients of equal rank. By coupling coefficients with themselves we obtain the *power-spectrum* known from ordinary Fourier analysis, with $(\frac{1}{\sqrt{2\ell+1}}(\mathbf{a}_\ell^k \circ_0 \mathbf{a}_\ell^k))^{\frac{1}{2}} = \langle \mathbf{a}_\ell^k, \mathbf{a}_\ell^k \rangle^{\frac{1}{2}} = \|\mathbf{a}_\ell^k\|$. We further can couple expansion coefficients corresponding to Bessel functions $j_\ell(rk_1)$ and $j_\ell(rk_2)$ of different frequencies k_1 and k_2, with $(\frac{1}{\sqrt{2\ell+1}}(\mathbf{a}_\ell^{k_1} \circ_0 \mathbf{a}_\ell^{k_2}))^{\frac{1}{2}} = \langle \mathbf{a}_\ell^{k_1}, \mathbf{a}_\ell^{k_2} \rangle^{\frac{1}{2}}$. In the experiments this feature is called the *Phase*-feature.

First Experiment: Detection and Classification of Cells. In this experiment we aim to detect DAPI[1]-stained nuclei of root cap cells represented in volumetric images of Arabidopsis root tips. The data was recorded using a confocal laser-scanning microscope. Experiments have shown, that, in contrast to inner cells, root cap cells can hardly be detected by strategies suitable for detecting roundish structures, e.g. using the spherical Hough transform [12]. For this

---

[1] 4'-6-Diamidino-2-phenylindole

scenario a voxel-wise feature computation and classification approach is suitable for first learning the shape and structure of root cap cells using rotation invariant features, followed by a detection and classification of cells in unclassified data sets. The experiment is organized as follows: We select one image for training (size $461 \times 470 \times 101$) and one further image for testing (size $239 \times 522 \times 105$) (depicted in the left image of figure 1 a) and b)). The voxel size of each images is $1\mu m$. Each image contains several hundred cells. The center of the nuclei of the cells were labeled manually and divided into two classes: Root Cap Cells and Inner Cells. We compute rotation invariant voxel-wise features for $k = 1, 3, 5, 7, 9$ and a band-with limit $\ell \leq 5$. σ is set to 2π. The kernel is scaled by a factor of 6 approximately covering a whole cell. We separately normalize each feature dimension with respect to the mean and the variance of the whole set of features computed for a single dataset. For training we select all features representing cells from the training image. We further randomly select features describing locations not belonging to cells representing the background. We train a two-class cSVM [13] using an RBF kernel with $\gamma=1$ and cost=1. We first conduct experiments comparing features based on the power-spectrum to the phase-features. For the voxel-wise classification we use the local maxima of the decision values of the SVM. We have a true positive detection if correctly classifying a root cap cell in a $3\mu m$ surrounding of a positive label and a false negative for each root cap cell, which is not detected. All remaining voxels, wrongly classified as root cap cells count as false positives. Results are depicted in the upper left PR-graph shown in figure 2. Surprisingly the phase-feature performs much better than the power-spectrum feature. We belief that this is caused by the textural information of cells in radial direction which is better preserved in the features when coupling coefficients representing different radial frequency components. However, as expected, coupling all possible coefficients leads to the best results (Powerspectrum+Phase). Although we often successfully use the spherical
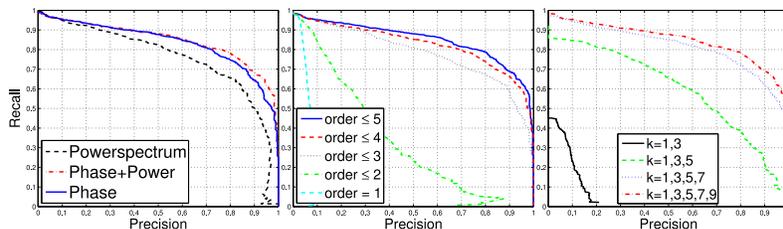

Fig. 2. Left: Comparing the features obtained by only coupling coefficients with themself (power-spectrum), coupling all possible coefficients (phase+power) and finally only coupling coefficients with different values of k (phase). Middle: Performance for different band-width limits ℓ. Right: Performance for different numbers of radial functions.

Hough transformation for cell detection, it is not possible using it to detect more than 60% of the cells due to their non-roundish shape. In the right PR-graph

of figure 1 we show the performance of our method outperforming the detection rate of the spherical Hough transformation significantly.

We further conduct experiments for different numbers of radial and angular frequency components. As expected the performance increases when considering further higher frequency components. This is true for both increasing the number of spherical Bessel functions and increasing the order of the spherical harmonics. The results are depicted in the middle and right graph of figure 2.

Second Experiment: α-Helices in Secondary Structures of Proteins. Electron cryomicroscopy is a powerful technique for analyzing the dynamics and functions of large flexible macromolecule assemblies. One major challenge in analyzing such density maps is the detection of subunits and their conformations. One important step in this procedure is the detection of secondary structure elements, mainly the α-helices.

The database for our experiments consists of simulated electron microscopic volumes of 56 polymers with an EM-resolution of 10Å and 1Å per voxel [14]. The data is divided into a training set (4 files) and a test set (52 files). We first try to detect helices using Helixhunter [15], mainly based on an eigen-analysis of the second moment tensor of local structures. We further use a harmonic filter of order 5 [7]. For pre- and post-smoothing we use a Gaussian with $\sigma = 1.5$. Finally we perform experiments using our own rotation invariant power-spectrum-features, based on the coefficients of the Fourier tensor transformation. We compute features directly based on the intensity values. We further compute a second order tensor field by computing the structure tensor at each voxel position. We use eq. 2.1 for representing the traceless, symmetric parts of the resulting Cartesian tensor field in terms of a spherical tensor field. We compute features for both the intensity values and the structure-tensor field with $k = 1, 3, 5$ and $\ell \leq 5$. We scale the kernel by a factor of 2 and use $\sigma = 2$ for the Gaussian window function. The features are normalized by weighting the components with respect to their frequency using the weights $k^{\lambda_k} \ell^{\lambda_\ell}$. We obtain the best results when suppressing the lower frequency components by amplifying the higher frequency components using $\lambda_\ell = 3$ and $\lambda_k = 5$. For voxel-wise classification we use a 20KNN-classifier using the $l1$-norm. We only count voxels correctly classified as α-helices as true positives. The results of our experiments are depicted in figure 3. In our experiments Helixhunter has major problems to determine the exact locations of helices. For the experiments based on the harmonic filter, we vary the size of the Gaussian convolution kernels as well as the order of the filter. The order which works best here is 5. Increasing the order to a higher extent does not noticeable increase the performance. This is caused by the low resolution of the data. The same behavior can be observed for our features, too. Similar to the order of the harmonic filter, the bandwidth limit $\ell \leq 5$ restricts the number of tensorial harmonics representing the signal in angular directions. We experienced that in contrast to the intensity features, it is much easier to obtain good results without having much effort for finding good parameters using the structure-tensor based features.
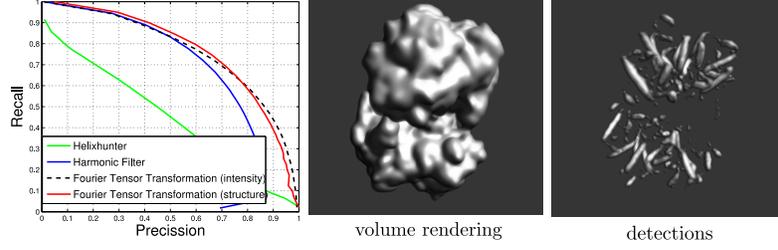
Fig. 3. From left to right: PR-graph showing the performance of the detection procedures used in the experiments. Surface rendering representing the secondary structure with PDB code 1m3z. Corresponding detections of helices using our method.

5   Conclusion

In this paper we extended the spherical Fourier transformation presented in [9] to higher order tensor fields and further presented a new method for performing the transformation in a fast and memory efficient way. We have shown how to utilize the spherical tensor algebra to compute rotation invariant features from the Fourier coefficients in an analytical way. We introduced all details necessary for an implementation, and gave two examples where our rotation invariant local features are used for successfully detecting and classifying objects in volumetric images leading to very promising results.

A   Differential Formulation of Spherical Fourier Functions

We obtain $\mathbf{B}_k^\ell$ by ℓ times applying $\boldsymbol{\nabla}^1$ to $\mathbf{B}_k^0$. Proof:

$$
\begin{aligned}
\langle e^{i\mathbf{k}^T\mathbf{r}}, B_{k'm}^\ell(\mathbf{r})\rangle &= \tfrac{2}{\pi}\int_0^\infty (i)^{\ell'}(2\ell'+1)\ j_{\ell'}(kr)j_\ell(k'r)r^2 \sum_{\ell'} Y_m^{\ell'}(\mathbf{k})\tfrac{1}{(2\ell+1)}\delta_{\ell\ell'}dr \\
&= \tfrac{2}{\pi}(-i)^\ell Y_m^\ell(\mathbf{k})\underbrace{\int_0^\infty\ j_\ell(kr)j_\ell(k'r)r^2 dr}_{\frac{\pi}{2k^2}\delta(k-k')} = (-i)^\ell Y_m^\ell(\mathbf{k})\tfrac{1}{k^2}\delta(k-k') \quad .
\end{aligned}
\tag{7}
$$

This means that a Bessel function $\mathbf{B}_k^\ell$ in the Fourier domain is a spherical harmonic living on a sphere with radius k. Consider the representation of both the Bessel function $\mathbf{B}_k^0$ and the operator $\boldsymbol{\nabla}^\ell$ in the frequency domain (eq. (7) and according to [5] $\widetilde{\boldsymbol{\nabla}^\ell} = i^\ell \mathbf{R}^\ell(\mathbf{k})$). In this scenario we can observe, that $\widetilde{\boldsymbol{\nabla}^\ell \mathbf{B}_{k'}^0} = i^\ell \mathbf{R}^\ell(\mathbf{k})\tfrac{1}{k^2}\delta(k-k')$. Performing the inverse transformation into the spatial domain we obtain

$$
\langle e^{-i\mathbf{k}^T\mathbf{r}}, \widetilde{\boldsymbol{\nabla}^\ell \mathbf{B}_{k'}^0}\rangle = (-1)^\ell k'^\ell j_\ell(rk')\mathbf{Y}^\ell(\mathbf{r}) = (-1)^\ell k'^\ell \mathbf{B}_{k'}^\ell(\mathbf{r}) \quad .
$$

It follows, that we obtain higher order Bessel functions $\mathbf{B}_k^\ell$ by iteratively applying the spherical derivative operator $\boldsymbol{\nabla}^1$ to $\mathbf{B}_k^0$, namely $\boldsymbol{\nabla}^\ell \mathbf{B}_k^0 = (-1)^\ell k^\ell \mathbf{B}_k^\ell$ $\square$.

References

1. Scovanner, P., Ali, S., Shah, M.: A 3-dimensional sift descriptor and its application to action recognition. In: Proc. of the 15th international conference on Multimedia, ACM Press New York, NY, USA (2007) 357–360
2. Rose, M.: Elementary Theory of Angular Momentum. Dover Publications (1995)
3. Kazhdan, M., Funkhouser, T., Rusinkiewicz, S.: Rotation invariant spherical harmonic representation of 3d shape descriptors. In: SGP '03: Proc. of the 2003 Eurographics/ACM SIGGRAPH symposium on Geometry processing, Eurographics Association (2003) 156–164
4. Ronneberger, O., Fehr, J., Burkhardt, H.: Voxel-wise gray scale invariants for simultaneous segmentation and classification. In: Proc. of the DAGM, Vienna, Austria, LNCS, Springer (2005) 85–92
5. Reisert, M., Burkhardt, H.: Spherical tensor calculus for local adaptive filtering. In Aja-Fernández, S., de Luis García, R., Tao, D., Li, X., eds.: Tensors in Image Processing and Computer Vision. Springer (2009)
6. Skibbe, H., Reisert, M., Wang, Q., Ronneberger, O., Burkhardt., H.: Fast computation of 3D spherical Fourier harmonic descriptors - a complete orthonormal basis for a rotational invariant representation of three-dimensional objects. In: Proc. of the 3DIM, part of the ICCV, Kyoto, Japan (2009)
7. Reisert, M., Burkhardt, H.: Harmonic filters for generic feature detection in 3D. In: Proc. of the DAGM, Jena, Germany, LNCS, Springer (2009) 131–140
8. Skibbe, H., Reisert, M., Ronneberger, O., Burkhardt, H.: Increasing the dimension of creativity in rotation invariant feature design using 3d tensorial harmonics. In: Proc. of the DAGM, Jena, Germany, LNCS, Springer (2009) 141–150
9. Wang, Q., Ronneberger, O., Burkhardt, H.: Rotational invariance based on Fourier analysis in polar and spherical coordinates. IEEE Trans. on PAMI **31** (2009) 1715–1722
10. Abramowitz, M., Stegun, I.A.: Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables. 1 edn. Dover books on mathematics. Dover Publications (June 1970)
11. Frigo, M., Johnson, S.G.: The design and implementation of FFTW3. Proc. of the IEEE **93**(2) (2005) 216–231 Special issue on "Program Generation, Optimization, and Platform Adaptation".
12. Schulz, J., Schmidt, T., Ronneberger, O., Burkhardt, H., Pasternak, T., Dovzhenko, A., Palme, K.: Fast scalar and vectorial grayscale based invariant features for 3d cell nuclei localization and classification. In: Proc. of the DAGM, Berlin, Germany, LNCS, Springer (2006) 182–191
13. Chang, C.C., Lin, C.J.: LIBSVM: a library for support vector machines. (2001) Available at `http://www.csie.ntu.edu.tw/~cjlin/libsvm`.
14. Berman, H.M., Westbrook, J.D., Feng, Z., Gilliland, G., Bhat, T.N., Weissig, H., Shindyalov, I.N., Bourne, P.E.: The protein data bank. Nucleic Acids Research **28**(1) (2000) 235–242
15. Jiang, W., Baker, M.L., Ludtke, S.J., Chiu, W.: Bridging the information gap: computational tools for intermediate resolution structure interpretation. Journal of Molecular Biology **308**(5) (2001) 1033–1044