SPHERICAL BESSEL FILTER FOR 3D OBJECT DETECTION

Henrik Skibbe[1,3]*, Marco Reisert*[2]*, Olaf Ronneberger*[1,3] *and Hans Burkhardt*[1,3]

[1]Department of Computer Science, Albert-Ludwigs-Universität Freiburg, Germany
[2]Dept. of Diagnostic Radiology, Medical Physics , University Medical Center, Freiburg
[3]BIOSS Centre for Biological Signalling Studies, University of Freiburg

skibbe@informatik.uni-freiburg.de, marco.reisert@uniklinik-freiburg.de

ABSTRACT

The detection of 3D objects and landmarks in arbitrary orientations is one of the most challenging tasks in biomedical 3D image analysis. In this paper we introduce the spherical Bessel Filter (BF) for rotation invariant 3D object detection tasks. The BF is based on the Harmonic Filter (HF) and thus inherits all the gentle properties of the HF, in particular the data driven adaptability and the processing speed. In contrast to the HF the BF benefits from a better object representation based on local spherical Fourier basis functions leading to noticeably better object detections and localizations.

1. INTRODUCTION

The rotation invariant object detection in 3D images is a challenging task, due to the huge data size and the non-trivial 3D rotation. Facing these kind of problems, the Harmonic Filter (HF) [5, 8] is a state-of-the-art approach that automatically adapts to certain detection tasks based on training images. Furthermore, it is fast and able to detect 3D objects, feature or landmarks in arbitrary orientations. What we introduce here is an extension of the HF by using spherical Fourier basis functions to describe and locate the desired objects more precisely.

The paper is structured as follows: In section 2 we give a rough overview on the mathematical notations and definitions we use in this paper. In section 3 we introduce the Bessel Filter (BF) itself by pointing out the extensions we make with respect to the HF and we show how to design the filter in an efficient way. Finally, the experiment in section 4 demonstrates the performance of the BF outperforming the HF significantly.

2. PRELIMINARIES

Volumetric images, usually represented in Cartesian coordinates $\mathbf{r} = (x, y, z)^T$, can equivalently be considered to be represented in spherical coordinates $\mathbf{r} = (r, \theta, \phi)$, where r represents the distance to the origin, and θ, ϕ denote the direction as point on the 2-sphere. The Bessel Filter (BF) is mainly based on the theory of spherical Fourier analysis. In order to do Fourier-analysis in spherical coordinates we must know, that the spherical expansion of the plane wave $e^{i\mathbf{k}^T\mathbf{r}}$ is given by $e^{i\mathbf{k}^T\mathbf{r}} = \sum_\ell (i)^\ell (2\ell+1) j_\ell(kr) \mathbf{Y}^\ell(\mathbf{r}) \bullet_0 \mathbf{Y}^\ell(\mathbf{k})$. The functions $\mathbf{Y}^\ell : \mathbb{R}^3 \rightarrow \mathbb{C}^{2\ell+1}$ are called *spherical harmonics* and are widely known from angular momentum theory [3] describing the rotational state of physical systems. The single components $Y^\ell_{m=\{-\ell\dots\ell\}}$ with $\ell \in \mathbb{N}_0$ build an orthogonal basis system for functions on the 2-sphere (see figure 1). The functions $j_\ell(kx)$ with $k \in \mathbb{R}$ (the frequency of the wave) are the spherical Bessel functions [1] which build an orthogonal basis ($k \neq k' \Rightarrow j_\ell(kx) \perp j_\ell(k'x)$) for one dimensional signals with respect to the weighting function r^2 (see plot depicted in figure 1). In [10] it is shown how to expand a function $f : \mathbb{R}^3 \rightarrow \mathbb{R}$ in terms of Fourier basis functions $\mathbf{B}^\ell_k(\mathbf{r}) := \mathbf{Y}^\ell(\theta, \phi) j_\ell(kr)$ separately in angular direction using $\mathbf{Y}^\ell$ of order ℓ and in radial direction using j_ℓ with frequency k. Hence we can represent a function f in terms of $\mathbf{B}^\ell_k$ and the object center $\mathbf{c}$ by expansion coefficients $\mathbf{a}^\ell_k(\mathbf{r}) = \langle f_\mathbf{c}, \mathbf{B}^\ell_k \rangle$, where $f_\mathbf{c}(\mathbf{r}) := f(\mathbf{r} + \mathbf{c})$.

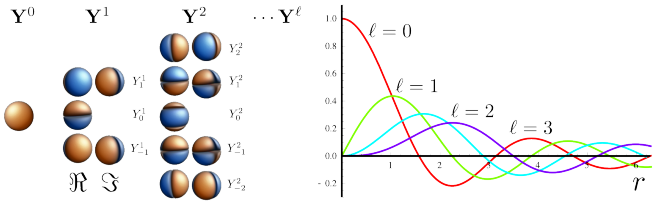

Fig. 1. Left: Spherical harmonics $\mathbf{Y}^0$,$\mathbf{Y}^1$,$\mathbf{Y}^2$ evaluated on the 2-sphere. Right: Spherical Bessel functions $j_\ell(kr)$ for $\ell = 0$,1,2,3 and frequency $k = 2$.

In order to explain the meaning of the symbol $\bullet$ we have to introduce the concept of *spherical tensor* fields [3], first. Cartesian tensor fields are widely used for 3D image analysis, e.g. the gradient vector field or a field based on structure tensors. They all have in common that a rotation acting on the whole field is acting on the field's elements, too. For each

This study was supported by the Excellence Initiative of the German Federal and State Governments (EXC 294)

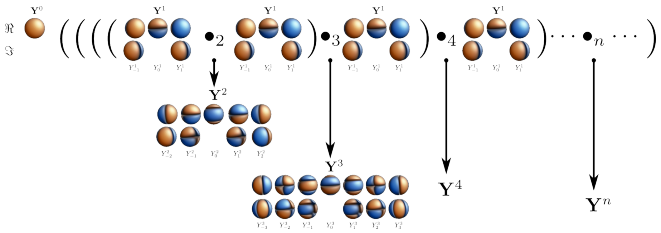

Fig. 2. One of the most important properties of $\bullet_\ell$ and $\mathbf{Y}$: Recursively deriving higher order spherical harmonics $\mathbf{Y}^{\ell+1}$ by element-wise coupling the previous spherical tensor fields $\mathbf{Y}^\ell$ and $\mathbf{Y}^1$.

Cartesian tensor exists an equivalent irreducible representation consisting of vector-valued spherical tensors $\mathbf{v} \in \mathbb{C}^{2\ell+1}$, where ℓ denotes the tensor rank. A scalar field can directly be considered as a spherical tensor field of order 0. The symbol $\bullet_\ell$ represents the *spherical tensor product* $\bullet_\ell : \mathbb{C}^{2\ell_1+1} \times \mathbb{C}^{2\ell_2+1} \to \mathbb{C}^{2\ell+1}$ (see [6] for details) which can be used to efficiently couple two spherical tensors of rank ℓ_1 and ℓ_2 to obtain a new tensor of rank ℓ (A Cartesian counterpart of $\bullet_0$ is the dot product reducing the tensor rank to 0. A Cartesian counterpart of $\bullet_2$ is the traceless outer product of two gradient vectors increasing the rank to 2.). The most important reason for using tensor algebra for 3D image processing is the fact that combining two tensors always implicitly preserves their rotation state. Further when combining tensors using $\bullet_0$ we obtain a tensor of order 0. This 0-ordered tensor is scalar valued and thus is **rotation invariant**.

The spherical harmonic $\mathbf{Y}^\ell$ can be considered as a spherical tensor field of order ℓ. One very important property of $\mathbf{Y}^\ell$ is that we can utilize $\bullet_\ell$ to recursively obtain an $\mathbf{Y}^{\ell+1} = \mathbf{Y}^\ell \bullet_{\ell+1} \mathbf{Y}^1$ as depicted in figure 2. It further follows that if the spherical harmonics are spherical tensors, then $\mathbf{B}_k^\ell(\mathbf{r})$ and the coefficients $\mathbf{a}_k^\ell$ are spherical tensors, too.

The symbol $\widetilde{\bullet}_\ell$ (see [6]) denotes a convolution combining the ordinary convolution and the tensor product $\bullet_\ell$. The resulting tensor field has order ℓ. The operators $\boldsymbol{\nabla}^\ell$ and $\boldsymbol{\nabla}_\ell$ [6] are called the spherical up derivatives and spherical down derivatives, respectively. $\boldsymbol{\nabla}^\ell$ is recursively defined by $\boldsymbol{\nabla}^\ell := \nabla \bullet_\ell \boldsymbol{\nabla}^{\ell-1} f$, where $\nabla = \left(\frac{1}{\sqrt{2}}(\partial_x - i\partial_y), \partial_z, -\frac{1}{\sqrt{2}}(\partial_x + i\partial_y)\right)^T$ is the spherical gradient and $\partial_x, \partial_y, \partial_z$ are the (ordinary Cartesian) partial derivatives. $\boldsymbol{\nabla}_\ell$ is defined similarly but decreases the tensor rank by ℓ. $\overline{\boldsymbol{\nabla}}^\ell$ and $\overline{\boldsymbol{\nabla}}_\ell$ denote the complex conjugate operators of $\boldsymbol{\nabla}^\ell$ and $\boldsymbol{\nabla}_\ell$, respectively.

As mentioned before, a volumetric image can be represented in terms of spherical tensor valued expansion coefficients $\mathbf{a}_k^\ell$. These coefficients can be computed very efficiently when computing them for all voxels simultaneously. This is done by making use of a differential interrelation of spherical Fourier functions (as shown in [9]). The differential formulation is given by

$$\boldsymbol{\nabla}^\ell \mathbf{B}_k^0 = (-1)^\ell k^\ell \mathbf{B}_k^\ell \quad , \tag{1}$$

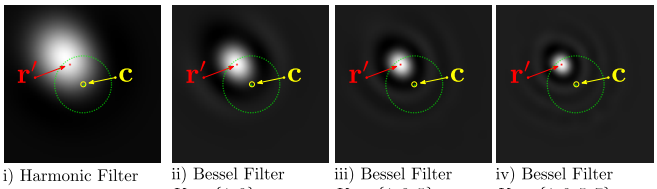

i) Harmonic Filter ii) Bessel Filter $K_i = \{1,3\}$ iii) Bessel Filter $K_i = \{1,3,5\}$ iv) Bessel Filter $K_i = \{1,3,5,7\}$

Fig. 3. Voting functions of the Harmonic filter and the Bessel filter: Centered Z-slices showing the 3D voting functions $V_{\mathbf{c}}(\mathbf{r}-\mathbf{c})$ voting for a supposed object center $\mathbf{r}'$. i) Voting of the Harmonic Filter of order 10. ii) Bessel Filter of order 10 with 2 radial Bessel frequencies. iii-iv) Accurate localization with 3 Bessel frequencies and slightly better location with 4 frequencies, respectively. The green dashed circle indicates the size of the scale of the kernel.

leading to

$$\mathbf{a}_\ell^k(\mathbf{r}) \propto \langle f_{\mathbf{r}}, \boldsymbol{\nabla}^\ell \mathbf{B}_k^0 \rangle = \overline{\boldsymbol{\nabla}}^\ell (f * \mathbf{B}_k^0)(\mathbf{r}) \quad . \tag{2}$$

The differential formulation of spherical Bessel function is crucial for realizing the trainable voting scheme of the Bessel Filter in a memory and runtime efficient way.

3. THE BESSEL FILTER

The Bessel Filter (BF) is an extension to the Harmonic Filter (HF) [5]. The HF as well as the BF can be regarded as a kind of a context depending voting scheme. Each voxel in the volume casts votes to its surrounding, depending on the appearance of its local neighborhood. The function that maps the local appearance of a voxel onto the vote distribution is learned automatically in a parametric way. For the representation of both the local appearance descriptors and voting distribution the HF uses a particular kind of orthogonal functions. We propose to use a different and complete set of functions, which enables the BF to detect and localize objects much more precisely and, additionally, preserves the high efficiency in computation via spherical derivatives.

First, the HF utilizes the Gaussian windowed solid harmonics (GSH) to represent the voting distribution. For the Bessel Filter we use the Fourier basis $\mathbf{B}_k^\ell$. Consider Figure 3: assume we are at position $\mathbf{c}$ and want to cast a vote to position $\mathbf{r}'$. Ideally we just want to have a sharp peak at position $\mathbf{r}'$, but the HF is restricted to use GSH as a basis. Figure 3i) shows the best approximation of this peak with respect to the GFH basis, which is rather blurry and inaccurate. By changing the basis to the proposed Fourier basis $\mathbf{B}_k^\ell$ we can see that we are able to represent the peak much better (Figure ii-iv). In contrast to eq. (4) in [5] the contribution to $\mathbf{r}$ originating from position $\mathbf{c}$ is now determined by

$$V_{\mathbf{c}}(\mathbf{r}) = \sum_i \sum_{\ell=0}^{\infty} \mathbf{V}_{K_i}^\ell(\mathbf{c}) \bullet_0 \mathbf{B}_{K_i}^\ell(\mathbf{r}-\mathbf{c}) \quad , \tag{3}$$

where $\mathbf{V}_K^\ell \in \mathbb{C}^{2\ell+1}$ are the expansion coefficients of the filter and $K_i \in \mathbb{R}$ is a finite number of radial frequencies considered for the filter expansion. According to the definition of the HF the BF of order N can then be defined by

$$\begin{aligned}\mathcal{H}\{f\}(\mathbf{r}) &:= \int_{\mathbb{R}^3} V_{\mathbf{c}}(\mathbf{r}) d\mathbf{c} \\ &= \sum_{\ell=0}^{N} \sum_i (\mathbf{B}_{K_i}^\ell \widetilde{\bullet}_0 \mathbf{V}_{K_i}^\ell)(\mathbf{r}) \quad . \end{aligned} \tag{4}$$

In a similar way to [5] we can determine the differential formulation of $\mathcal{H}$ in terms of $\boldsymbol{\nabla}$ using the properties of $\boldsymbol{\nabla}$,$\widetilde{\bullet}_0$ [6] and eq. (1), highly reducing the computation time by making the number of scalar valued convolutions independent from ℓ, namely

$$\mathcal{H}\{f\} := \sum_i \mathbf{B}_{K_i}^0 * \sum_\ell^N \boldsymbol{\nabla}_\ell \mathbf{V}_{K_i}^\ell \quad . \tag{5}$$

The key idea is that differentiating using $\boldsymbol{\nabla}$ can be realized in linear time in number of voxel $O(N)$ using finite differences while the convolution $\widetilde{\bullet}_0$ has complexity $O(N \log N)$ (where N is the number of voxel).

Secondly, the way of describing the local appearance of a voxel differs, too. Instead of Gaussian windowed solid harmonics we also use the Fourier basis $\mathbf{B}_k^\ell$ to find appropriate coefficients $\mathbf{V}_{K_i}^\ell$ in an image content dependent way. This is done by selecting a finite set of radial frequencies $\{k_1, k_2, \dots\}$ with which we directly select a subset of Fourier expansion coefficients $\mathbf{a}_{k_j}^\ell$ describing each voxel surrounding as local image descriptors, with (according to eq. (2))

$$\begin{aligned}\mathbf{V}_{K_i}^\ell &:= \textstyle\sum_j \mathbf{V}_{K_i k_j}^\ell [\mathbf{a}_{k_j}^0, \dots, \mathbf{a}_{k_j}^N] \quad , \\ &= \textstyle\sum_j \mathbf{V}_{K_i k_j}^\ell [\overline{\boldsymbol{\nabla}}^0 \mathbf{a}_{k_j}^0, \dots, \overline{\boldsymbol{\nabla}}^N \mathbf{a}_{k_j}^0] \quad , \end{aligned} \tag{6}$$

where

$$\mathbf{V}_{K_i k_j}^\ell [\mathbf{a}_{k_j}^0, \dots, \mathbf{a}_{k_j}^N] := \sum_{\substack{|\ell_1-\ell_2| \le \ell \le \ell_1+\ell_2 \\ \ell_1+\ell_2+\ell \text{ even} \\ \ell_1,\ell_2,\ell \le N}} \alpha_{\ell_1,\ell_2,\ell}^{K_i k_j} (\mathbf{a}_{k_j}^{\ell_1} \bullet_\ell \mathbf{a}_{k_j}^{\ell_2}). \tag{7}$$

$\alpha_{\ell_1,\ell_2,\ell}^{K_i k_j} \in \mathbb{R}$ are scalar valued expansion coefficients and $k_j \in \mathbb{R}$ denotes the radial frequency used for the local image descriptors. The scalar valued, image content dependent expansion coefficients $\alpha_{\ell_1,\ell_2,\ell}^{K_i k_j}$ are determined in a training phase by a least square fit of the filter response to a binary-valued label image (similar to [5]).

4. EXPERIMENT

It has already been shown in [5] that the Harmonic Filter (HF) outperforms a general Hough transform [2] based on local 3D SIFT descriptors in rotation invariant 3D object detection tasks in precision and speed. It is further shown [5] that the HF outperforms spherical harmonic based descriptors (see. e.g. [4, 9]) in combination with a nearest neighbor classifier or a support vector machine when only the detection and not the classification of objects is required. However, when using Fourier basis functions instead of Gaussian windowed solid harmonics for both the computation of local image descriptors and for expanding the filter we expect a lower false positive rate due to a better object description and a much more accurate object localization due to the improved voting scheme (see figure 3). We conduct an experiment where we compare the performance of the HF to the performance of the Bessel Filter (BF) using the same database as used in [5], where we detect porates of airborne pollen. Due to the fact that the Bessel functions are infinite, we use the scalar-valued Gaussian windowed convolution kernel

$$\mathbf{B}_k^0(\mathbf{r}/t) e^{-r^2/(2t^2)} \tag{8}$$

for implementation, where $t \in \mathbb{R}$ defines the scale of the kernel.

In addition to the murgwort pollen dataset used in [5] (15 3D images, 45 porates) we use a birch pollen dataset (15 3D images, 45 porates) and an alder pollen dataset (15 3D images, 60 porates). Each dataset is recorded by a confocal laser scanning microscope (see [7]) and is used to compare the performance of the BF to the HF in a landmark detection task in biological images. Each image size is about 80^3 voxels. The voxel size corresponds to $0.408948 \mu m$. The experimental setup is identical to the experiment in [7] using only one training image containing one single pollen having 3 porates (mugwort and birch) or 4 porates in case of the alder pollen dataset. The three training images are depicted in figure 5. The porates are labeled manually by an expert and these labels are used for training and evaluating the performance.

We first conduct experiments for finding the most suitable parameters for the BF. The following parameters must be set: A small set of frequencies $\{K_1, .., K_i, k_1, .., k_j\}$ used for designing the voting function and computing the local image descriptors. Secondly, we have the parameters already known from the HF, namely, the filter degree N determining the granularity of the voting function (see e.g. figure 8) and of

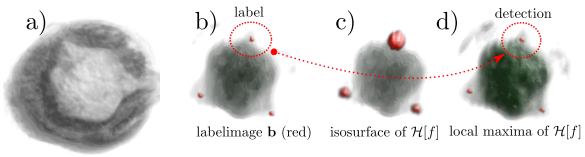

Fig. 4. a) 3D rendering of an airborne pollen. b) Labels (here: 3 voxel have been set to 1, all others to 0) are indicating the center of the porates and are used for training and evaluation. c) Isosurface showing the response of the Bessel Filter. d) Local maxima of the responses. The porates are localized very precisely.

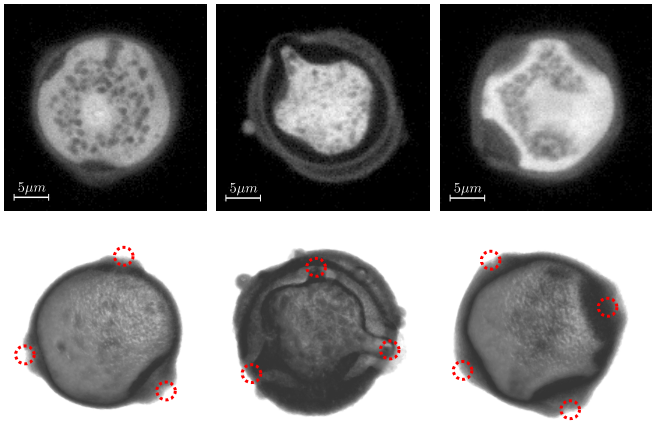

Fig. 5. Birch pollen (3 train, 42 test), murgwort pollen (3 train, 42 test), alder pollen (4 train, 56 test)

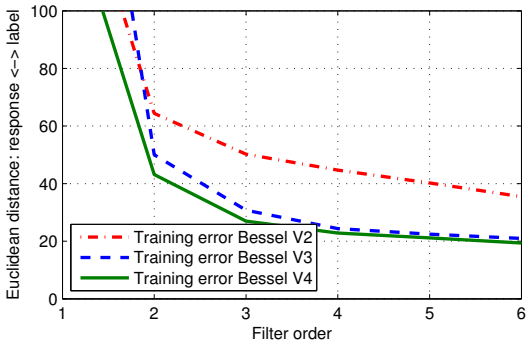

Fig. 6. Bessel Filters of order up to 6 have been trained using the training dataset of the birch-pollen database. The graph shows the Euclidean distance between the filter response and the label image (The label images consists of three one-valued voxels marking the positions of the porates. All other voxels are set to 0). The performance of three different kind of filters with two (V2), three (V3) and four (V4) radial frequencies are shown. Considering the results, a filter of order N=5 with 3 radial frequencies should be close to the best results that we can expect when given one training sample and using the Bessel Filter.

the local descriptors. Furthermore, the size of the surrounding considered for computing the local features (we call this σ, which means we set $t := \sigma$ in eq. (8) when computing the features) and the scale of the voting function (setting $t := \eta$ when computing the voting function). Consider, that the parameters might vary strongly for different kind of data. However, we experienced that the procedure we introduce here helps to find a good parameter set reliably.

The Bessel Filter is able to steer its voting function very precisely, hence choosing η to be half the diameter of the structure we aim to detect already leads to a very good performance. σ should be about half the size of η. To find a good combination of frequencies and the filter order we use the following procedure: As exemplary shown in figure 6 for varying the parameters of the voting function we train each filter for different orders N (here we use the training image of the birch-pollen dataset). In figure 6 the Euclidean distance between the label image used for training (zero-valued image but 3 one-valued voxel at the porate positions) and the filter response on the training image is depicted as a function of increasing filter order N. Considering the different curves we can observe that no significant improvement can be expected when using more than three radial frequencies (here $K_i = \{1, 3, 5\}$) and a filter of order 5. The training and detection procedure is illustrated in figure 4. Images demonstrating the adaptability of the corresponding voting functions and the filter responses on the training image are depicted in figure 8. With this method we find the following parameters for the birch pollen dataset: $(N = 5, K_i = \{1, 3, 5\}, k_j = \{1, 3\}, \eta = 2.86\mu m, \sigma = 1.4\mu m)$. Evaluating the performance on the corresponding test dataset justifies our decision (see figure 9 exemplary showing the performance for different orders N, and figure 10 comparing the performance of different kind of filters of order 5).

The parameters used for the three databases are summarized in table 1. The parameters for the Harmonic Filter are determined in the same way. However, due to both the simpler voting scheme and the features the HF is much more sensitive to small changes in η and σ. Hence it takes a lot more effort to find good parameters for the HF compared to the BF.

For being comparable to the experiments conducted in [5] we only use one image for training for each dataset (figure 5 shows the images used for training and determining the parameters). For each dataset we considered two scenarios according to the experiments in [5] comparing the performance of the BF to the performance of the HF. For evaluation we consider the local maxima of the filter response. The first scenario counts a detection as true-positive when it lies in a 4-voxel surround of the ground-truth label, else we count a false-positive. We count a false-negative for each label that is not detected. We use the magnitude of the local filter response for thresholding when generating the PR-graphs. The second scenario tolerates detections in an 8-voxel surrounding. The results for the birch-dataset are depicted in figure 7 a), the results for the murgwort pollen are shown in figure 7 b) and for the alder pollen in figure 7 c), respectively.

In all our experiments the Bessel Filter outperforms the Harmonic Filter. Especially for birch and alder pollen (both inducing hypersensitivity) the BF outperforms the Harmonic Filter significantly (figure 7 a) and b)).

5. CONCLUSION

We have shown how to utilize the basis functions of the spherical Fourier transformation to build an image filter for detecting 3D objects in arbitrary orientations. The Bessel Filter is

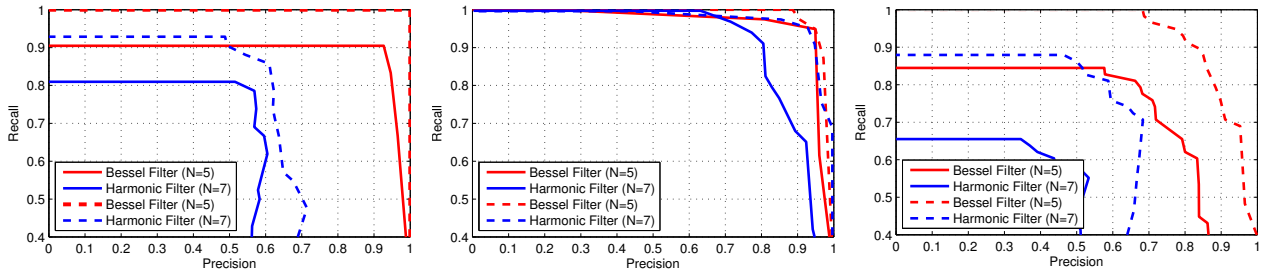

(a) Birch (b) Murgwort (c) Alder

Fig. 7. Detection of airborne-pollen porates. The dashed line represents the performance when tolerating a 8 voxel ($\approx 3.27\mu m$) displacement to the ground truth. The continuous line the results when tolerating a 4 voxel ($\approx 1.64\mu m$) displacement. The Bessel Filter outperforms the Harmonic Filter significantly.

based on the Harmonic Filter and thus inherits all the gentle properties of the Harmonic Filter, in particular the data driven adaptability and the processing speed outperforming many common techniques like SIFT based voting approaches or steerable filters (see [5, 8]). We conducted experiments demonstrating that due to the much more accurate representation of local surroundings in radial directions when using the spherical Bessel functions the Bessel Filter outperforms the Harmonic Filter significantly. Due to the data driven adaptability of the Bessel Filter and the low number of required training samples there are a lot of further applications in the biomedical field that will highly benefit from this framework (e.g. automated landmark detection for landmark based registrations, or for pre-aligning deformable models to the data).

6. REFERENCES

[1] M. Abramowitz and I. A. Stegun. *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*. Dover books on mathematics. Dover Publications, 1 edition, June 1970.

[2] D. H. Ballard. Generalizing the hough transform to detect arbitrary shapes. *Pattern Recognition*, 13(2):111–122, 1981.

[3] D. M. Brink and G. Satchler. *Angular Momentum*. Oxford Science Publications, 1993.

[4] M. Kazhdan, T. Funkhouser, and S. Rusinkiewicz. Rotation invariant spherical harmonic representation of 3d shape descriptors. In *SGP '03: Proc. of the 2003 Eurographics/ACM SIG-*

Table 1. Filter Parameters for our Experiments. The parameters η and σ are given in μm.

data	Filter	N	η	σ	$\{K_i\}$	$\{k_i\}$
Birch	HF	8	2.86	1.0		
	BF	5	2.86	1.4	$\{1,3,5\}$	$\{1,3\}$
Murgwort	HF	8	3.27	1.0		
	BF	5	2.86	1.4	$\{1,3,5\}$	$\{1,3\}$
Alder	HF	8	2.86	1.0		
	BF	5	2.86	1.4	$\{1,3,5\}$	$\{1,3\}$

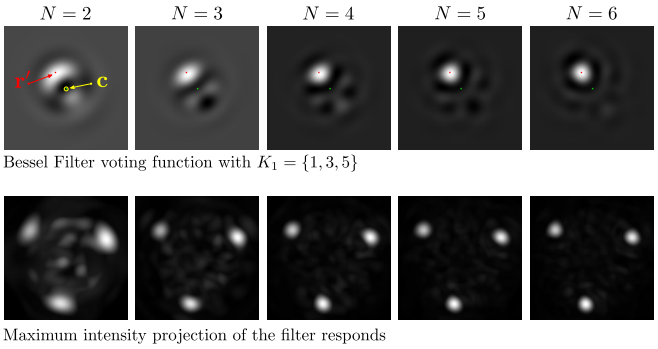

Fig. 8. *Upper row*: Voting function of a Bessel filter with 3 radial frequencies computed for an increasing order N. (Order 10 and further details are given in image iii) of figure 3). *Lower row*: Maximum intensity projection of the response of filters with increasing order N applied to the training sample of the birch-pollen dataset itself (Most left image in figure 5). The filter should respond to the three porates, only. Beginning with order 4 we only have little improvement. In our experiment an over-fitting takes place when using filters of order $\gg 4$ and the performance decreases (figure 9)

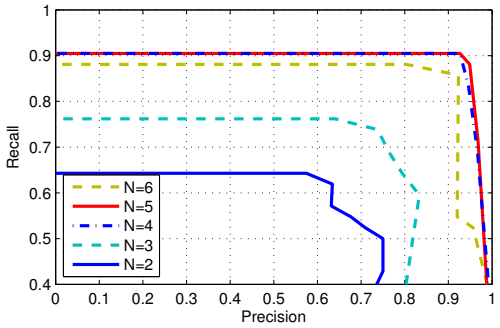

Fig. 9. Bessel Filters of order up to 6 have been trained and evaluated on the birch-pollen database. The filters are based on 2 different radial frequencies ($k = \{1,3\}$) for the features and 3 radial frequencies ($K = \{1,3,5\}$) to fit the voting functions.

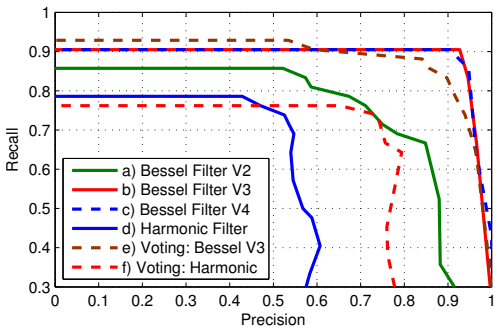

Fig. 10. Comparison of Bessel Filter (BF) and Harmonic Filter (HF) and combination of them. All filters in this experiment are of order N=5. Experiments are based on the birch-pollen dataset. For all BF we use 2 different frequencies for the features. For the filters a),b) and c) we use 2,3 and 4 radial frequencies for the voting functions. Filter d) is an ordinary HF. Filter e) is a filter using features of the HF and the same voting function as filter b). Filter f) uses the features of the BF and the voting function of a HF. These experiments show, that both the consideration of Bessel based features and the new voting scheme are increasing the performance significantly.

GRAPH symposium on Geometry processing, pages 156–164. Eurographics Association, 2003.

[5] M. Reisert and H. Burkhardt. Harmonic filters for generic feature detection in 3D. In *Proc. of the DAGM*, pages 131–140, Jena, Germany, 2009. LNCS, Springer.

[6] M. Reisert and H. Burkhardt. Spherical tensor calculus for local adaptive filtering. In S. Aja-Fernández, R. de Luis García, D. Tao, and X. Li, editors, *Tensors in Image Processing and Computer Vision*. Springer, 2009.

[7] O. Ronneberger, H. Burkhardt, and E. Schultz. General-purpose Object Recognition in 3D Volume Data Sets using Gray-Scale Invariants – Classification of Airborne Pollen-Grains Recorded with a Confocal Laser Scanning Microscope. In *Proc. of the International Conference on Pattern Recognition*, Quebec, Canada, 2002.

[8] M. Schlachter, M. Reisert, C. Herz, F. Schluermann, S. Lassmann, M. Werner, H. Burkhardt, and O. Ronneberger. Harmonic filters for 3d multi-channel data: Rotation invariant detection of mitoses in colorectal cancer. *Medical Imaging, IEEE Transactions on*, 29(8):1485–1495, 2010.

[9] H. Skibbe, M. Reisert, T. Schmidt, K. Palme, O. Ronneberger, and H. Burkhardt. 3D object detection using a fast voxel-wise local spherical Fourier tensor transformation. In *Proc. of the DAGM*, pages 412–421, Darmstadt, Germany, 2010. LNCS, Springer.

[10] Q. Wang, O. Ronneberger, and H. Burkhardt. Rotational invariance based on Fourier analysis in polar and spherical coordinates. *IEEE Trans. on PAMI*, 31:1715–1722, 2009.