

SHOG - Spherical HOG Descriptors for Rotation Invariant 3D Object Detection

Henrik Skibbe^{1,3}, Marco Reisert² and Hans Burkhardt^{1,3}

¹Department of Computer Science, University of Freiburg, Germany

²Dept. of Diagnostic Radiology, Medical Physics, University Medical Center, Freiburg

³Center for Biological Signalling Studies (BIOSS), University of Freiburg

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

Abstract. We present a method for densely computing local spherical histograms of oriented gradients (SHOG) in volumetric images. The descriptors are based on the continuous representation of the orientation histograms in the harmonic domain, which we compute very efficiently via spherical tensor products and the fast Fourier transformation. Building upon these local spherical histogram representations, we utilize the Harmonic Filter to create a generic rotation invariant object detection system that benefits from both the highly discriminative representation of local image patches in terms of histograms of oriented gradients and an adaptable trainable voting scheme that forms the filter. We exemplarily demonstrate the effectiveness of such dense spherical 3D descriptors in a detection task on biological 3D images. In a direct comparison to existing approaches, our new filter reveals superior performance.

1 Introduction

The rapid development of imaging techniques has led to a dramatic increase in the amount of volumetric image data that need to be processed. Especially in the field of biomedical imaging the third dimension becomes more and more important as it enables studying organisms in their natural constellation. Objects and organisms are sought to be located and analyzed in any number, at every position, and in every orientation. This means, volumetric data yields not only more demanding constraints regarding computational efficiency, also the interrelationship of neighboring intensity values becomes more complex. One of the most relevant issues is to cope with 3D rotation.

In this paper, we aim at creating filters based on the information of local gradient histograms that offer a robust, dense and rotation invariant object detection in volumetric images. For this we transfer the widely used HOG [2] features to the third dimension and show how to represent them in terms of so-called spherical tensors. Upon this representation we are capable to benefit from the simple rotation behavior of spherical tensors which enables the usage of Harmonic Filters [4,8]. This leads to a trainable 3D object and landmark detection system (figure 1) that benefits from both highly characteristic gradient orientation histograms and a memory and computational efficient trainable filter framework.

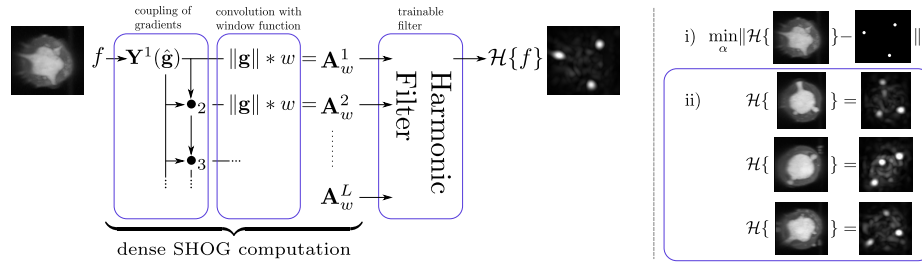


Fig. 1. Aiming at 3D landmark-detection: Above the flow diagram of the trainable SHOG-filter. We first compute the spherical gradient image of an input image and split it into its spherical orientation field $\mathbf{Y}^1(\hat{\mathbf{g}}) : \mathbb{R}^3 \rightarrow S^2$ and its magnitude image $\|\mathbf{g}\| : \mathbb{R}^3 \rightarrow \mathbb{R}$. Then the continuous spherical histograms of oriented gradients (SHOG) are computed densely in the whole image in a recursive manner. We get an expansion of the histograms in terms of vector-valued coefficients \mathbf{A}_w^ℓ . We finally use a trainable filter framework (Harmonic Filter) that learns a non-linear combination of the SHOG coefficients \mathbf{A}_j^ℓ such that a filter response is only given at the desired landmark positions. Moreover, all responses on the remaining positions are suppressed. Thanks to the continuous representation and the design of the coefficients \mathbf{A}_w^ℓ in combination with the Harmonic Filter, the filter response rotates smoothly with respect to the orientation of the landmarks without additional computational costs.

SHOG-Filter in action: i) Optimizing the filter parameter α . For this a binary label image is required. ii) The filter can now be applied to further objects.

Mathematical Notation: We write vectors $\mathbf{v} \in \mathbb{C}^n$ in bold letters. We denote the complex conjugate of \mathbf{v} by $\bar{\mathbf{v}}$ and the transpose of \mathbf{v} by \mathbf{v}^T . We consider unit-length vectors $\mathbf{n} = (x, y, z)^T \in \mathbb{R}^3$, $\|\mathbf{n}\| = 1$ w.l.o.g as points on the unit-sphere which we denote by $\mathbf{n} \in S_2$. We equivalently can represent \mathbf{n} in spherical coordinates (θ, ϕ) , where $\theta = \arccos(z)$ and $\phi = \text{atan2}(y, x)$ (see figure 2). We denote complex numbers by i , with $i^2 = -1$ and denote the convolution by $*$.

2 SHOG - Spherical Histograms of Oriented Gradients

Local descriptors based on orientation histograms, such as SIFT [3] and HOG [2], have revolutionized detection and matching in natural 2D images. Recently in particular HOG found its way in many applications because it can be computed efficiently and shows excellent performance. One step toward the third dimension HOG based features have been used for describing 3D mesh models [7,1]. What we propose here is a direct extension to volumetric images, where we aim at densely computing HOG at any image position. In contrast to 2D where a histogram is build upon gradient directions in a local neighborhood with respect to one angle (figure 2 a)), we must consider two angles for the 3D case (figure 2 b)). Hence the resulting histogram can be considered to be a histogram on a the 2-sphere (unit-sphere in 3D). We call the 3D representation of a HOG spherical HOG, or shortly SHOG. It is worth noting that the literature differs between R-HOG (rectangular spatial window) and C-HOG (circular,

isotropic window). Since the rotation of objects plays an important role in our framework, we only consider the latter one. Given an image $f : \mathbb{R}^3 \rightarrow \mathbb{R}$. We denote a dense field of SHOG descriptors defined over the whole image domain as $\text{SHOG}\{f\} : \mathbb{R}^3 \times S_2 \rightarrow \mathbb{R}$, where S_2 denotes the unit-sphere. For capturing




Fig. 2. In 3D the gradient direction is described by two angles thus a histogram of oriented gradients (HOG) can be considered as function on the sphere.

only the structure in a voxel's surrounding a window function $w : \mathbb{R}^3 \rightarrow \mathbb{R}$ is required. Such a window function is e.g. the 3D Gaussian function. We compute a local SHOG at position \mathbf{x} by collecting all magnitudes of gradients within the window function w contributing to orientation \mathbf{n} according to the continuous distribution function

$$\text{SHOG}\{f\}_w(\mathbf{x}, \mathbf{n}) = \int_{\mathbf{r} \in \mathbb{R}^3} \|\mathbf{g}(\mathbf{r})\| \delta_{\mathbf{n}}(\hat{\mathbf{g}}(\mathbf{r})) w(\mathbf{x} - \mathbf{r}) d\mathbf{r} \quad , \quad (1)$$

where $\mathbf{g} : \mathbb{R}^3 \rightarrow \mathbb{R}^3$, $\mathbf{g} = \nabla f$ is the gradient field of the volumetric image f , $\hat{\mathbf{g}} := \mathbf{g}/\|\mathbf{g}\|$, $\hat{\mathbf{g}} : \mathbb{R}^3 \rightarrow S_2$ the gradient orientation field and $\mathbf{n} \in S_2$ is the current histogram entry (the direction) taken into account. $\delta_{\mathbf{n}} : S_2 \rightarrow \mathbb{R}$ denotes the Dirac delta function on the unit sphere (see figure 5) that selects those gradients out of \mathbf{g} with orientation \mathbf{n} . A direct extension to the 2D HOG descriptor would require discrete sampling of the orientation space which is trivial in 2D, but in general a non-trivial task in 3D. A discretization would require an equidistant sampling of the sphere which in general can not be solved explicitly (known as Thomson problem [10]). To overcome this problem we propose to keep the histogram continuous and realize the "binning" in frequency domain instead. Due to this reason we gain the following advantages: First, no interpolation is required because our SHOG descriptor is based on the true continuous distribution function. Furthermore, if the window function is isotropic, the descriptor rotates with respect to rotation of its underlying data without leading to any discretization artifacts in the histogram. This plays a very important role when aiming at detecting objects in volumetric images at any position and in any orientation using the Harmonic Filter framework.

SHOG Decomposition: The Racah normalized spherical harmonic functions[6] $Y_m^\ell : S_2 \rightarrow \mathbb{C}$ build a complete orthogonal basis for functions on the unit sphere $f : S_2 \rightarrow \mathbb{C}$. Similar to the Cartesian Fourier basis, spherical harmonics represent




Fig. 3. Key property of \bullet_ℓ and \mathbf{Y} : higher order spherical harmonics $\mathbf{Y}^{\ell+1}$ can be obtained by element-wise coupling of spherical tensor fields \mathbf{Y}^ℓ with \mathbf{Y}^1 .

the different frequency components of spherical functions. We always have $2\ell + 1$ functions $Y_{m=\{-\ell, \dots, \ell\}}^\ell : S_2 \rightarrow \mathbb{C}$ representing a basis function of frequency ℓ , which can be arranged in a vector-valued function $\mathbf{Y}^\ell : S_2 \rightarrow \mathbb{C}^{2\ell+1}$.

Since a SHOG is a function on the sphere we can represent $\text{SHOG}_w(\mathbf{x})$ in terms of the orthogonal basis functions \mathbf{Y}^ℓ , namely

$$\text{SHOG}\{f\}_w(\mathbf{x}, \mathbf{n}) = \sum_{\ell=0}^{\infty} (\mathbf{A}_w^\ell(\mathbf{x}))^T \mathbf{Y}^\ell(\mathbf{n}) \quad , \quad (2)$$

where $\mathbf{A}_w^\ell(\mathbf{x}) \in \mathbb{C}^{2\ell+1}$ are the vector valued expansion coefficients completely representing the SHOG at image position \mathbf{x} in the spherical harmonic domain.

How to compute the coefficients $\mathbf{A}_w^\ell(\mathbf{x})$ densely for all image positions in an efficient manner? We identify the coefficients $\mathbf{A}_w^\ell(\mathbf{x})$ by plugging the spherical expansion of the Dirac delta function (figure 5) into eq. (2):

$$\begin{aligned} \text{SHOG}\{f\}_w(\mathbf{x}, \mathbf{n}) &:= \int_{\mathbf{r} \in \mathbb{R}^3} \|\mathbf{g}(\mathbf{r})\| \delta_{\mathbf{n}}(\hat{\mathbf{g}}(\mathbf{r})) w(\mathbf{x} - \mathbf{r}) d\mathbf{r} \\ &= \sum_{\ell=0}^{\infty} \int (2\ell + 1) \|\mathbf{g}(\mathbf{r})\| (\overline{\mathbf{Y}^\ell}(\hat{\mathbf{g}}(\mathbf{r})))^T w(\mathbf{x} - \mathbf{r}) d\mathbf{r} \mathbf{Y}^\ell(\mathbf{n}) \\ &= \sum_{\ell=0}^{\infty} (2\ell + 1) \underbrace{\left(\left(\|\mathbf{g}\| (\overline{\mathbf{Y}^\ell}(\hat{\mathbf{g}}))^T \right) * w \right) (\mathbf{x})}_{=\mathbf{A}_w^\ell(\mathbf{x}) \in \mathbb{C}^{2\ell+1}} \mathbf{Y}^\ell(\mathbf{n}) = \sum_{\ell=0}^{\infty} (\mathbf{A}_w^\ell(\mathbf{x}))^T \mathbf{Y}^\ell(\mathbf{n}) \quad . \quad (3) \end{aligned}$$




Fig. 4. The most left image shows a quantized SHOG. A band limited expansion in terms of spherical harmonics offers a smooth rotation with the underlying data and a memory efficient representation (here for $\ell \leq 5$).

For a fast computation of SHOG we utilize so-called spherical tensor products $\bullet_\ell : \mathbb{C}^{2\ell_1+1} \times \mathbb{C}^{2\ell_2+1} \rightarrow \mathbb{C}^{2\ell+1}$ [4] which can be used for coupling *spherical tensors* associated with different *orders* ℓ_1, ℓ_2 to form new tensors of higher or lower order ℓ i.e. in our scenario here for (point-wise) coupling different functions \mathbf{Y}^ℓ or for coupling the expansion coefficients \mathbf{A}_w^ℓ of SHOG_w . Most important, we can use \bullet_ℓ for recursively deriving spherical harmonics of order $\ell + 1$ by coupling two spherical harmonics of order ℓ and 1 with $\mathbf{Y}^{\ell+1} = \mathbf{Y}^\ell \bullet_{\ell+1} \mathbf{Y}^1$ for $\ell \geq 1$. In Fig. 3 we illustrate how higher order spherical harmonics can be computed recursively. Utilizing this property we gain a recursive rule with which we avoid an explicit, expensive computation of $\mathbf{Y}^\ell(\hat{\mathbf{g}})$, namely

$$\mathbf{Y}^{\ell+1}(\hat{\mathbf{g}}) = \mathbf{Y}^\ell(\hat{\mathbf{g}}) \bullet_{\ell+1} \mathbf{Y}^1(\hat{\mathbf{g}}) \quad . \quad (4)$$

Moreover, it turns out that $\|\mathbf{g}\| \mathbf{Y}^1(\hat{\mathbf{g}}) = \left(\frac{1}{\sqrt{2}} \left(\frac{\partial f}{\partial x} - i \frac{\partial f}{\partial y} \right), \frac{\partial f}{\partial z}, -\frac{1}{\sqrt{2}} \left(\frac{\partial f}{\partial x} + i \frac{\partial f}{\partial y} \right) \right)^T$ is just the spherical gradient of f which we compute in an initial step. The remaining computations are just the convolutions with the window function w that can be realized very efficiently by utilizing the Fast Fourier Transform.




Fig. 5. Band limited expansion of the Dirac delta $\delta_{\mathbf{n}} : S_2 \rightarrow \mathbb{R}$ on the unit-sphere: $\delta_{\mathbf{n}}(\mathbf{n}') := \sum_{\ell=0}^{\infty} (2\ell + 1) (\overline{\mathbf{Y}^\ell(\mathbf{n}')})^T \mathbf{Y}^\ell(\mathbf{n})$. For our experiments we use $\ell \leq 5$.

Object Detection in 3D - SHOG Features for Harmonic Filters: The Harmonic Filter [4,8] is a nonlinear polynomial filter that is designed for detecting arbitrary structures in volumetric images. The most important characteristic of this filter is a trainable voting scheme. The scheme comprises local image features to train a voting function such that the filter responses only to certain structures while responses to all remaining structures in the image are suppressed. This is achieved in an initial training step where the voting scheme is learned by providing a reference image together with a binary-valued label image (see our introductory example in figure 1). The local features of the original Harmonic Filter are the spherical derivatives of the 3D Gaussian encoding the intensity values of a voxel's surrounding in some kind of Taylor expansion coefficients. These features are then combined in a weighted, non-linear way. These weights are the free parameters that are optimized during the training step. Because of the spherical representation of the derivatives the features show a special, very simple rotation behavior depending on the rotation state of the underlying data. The filter comprises the rotation state of the features to steer

the voting function wherefore the filter response itself rotates smoothly with respect to the underlying data. Hence structures like objects or landmarks can be detected in any orientation. Since the spherical expansion coefficients \mathbf{A}_w^ℓ of the SHOG obey the same rotation behavior like the expansion coefficients of the spherical Gaussian derivatives in the original filter we propose to simply replace the Gaussian derivatives by SHOG in the Harmonic Filter framework. In addition to a non-linear combination of all expansion coefficients \mathbf{A}_w^ℓ we propose to additionally compute and combine coefficients derived from different window functions w_n (an angular cross-correlation of different local SHOG). Accordingly the voting function of the harmonic filter (eq. (6) in [4]) changes to

$$\mathbf{V}^\ell(\mathbf{x}) := \sum_{\substack{|\ell_1 - \ell_2| \leq \ell \leq \ell_1 + \ell_2 \\ \ell_1 + \ell_2 + \ell \text{ even} \\ \ell_1, \ell_2, \ell \leq N \\ n, m}} \alpha_{\ell_1, \ell_2, \ell}^{n, m} \underbrace{(\mathbf{A}_{w_n}^{\ell_1}(\mathbf{x}) \bullet_\ell \mathbf{A}_{w_m}^{\ell_2}(\mathbf{x}))}_{\text{non-linear combination of coefficients}} ; \quad (5)$$

$\alpha_{\ell_1, \ell_2, \ell}^{n, m} \in \mathbb{R}$ are the new weighting parameters that are learned in a training step.

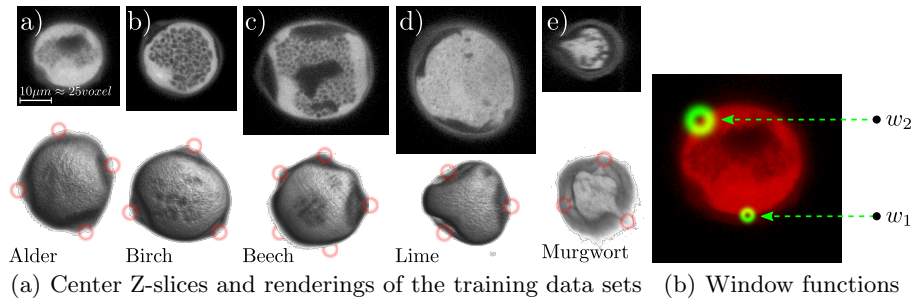


Fig. 6. Our Database: Alder pollen (4 porates for training, 56 for testing), Birch pollen (3 train, 42 test), Beech pollen (4 train, 65 test), Lime pollen (3 train, 42 test), Murgwort pollen (3 train, 42 test). **Figure b)** illustrates the size and shape of the two window functions that we use in our experiments (two nested smoothed spheres).

3 Experiments

For evaluating the performance of the SHOG-filter we aim at detecting landmarks in volumetric confocal recordings of airborne pollen. In particular we aim at detecting porates in 5 different kinds of pollen species [5], namely (see figure 6) Alder, Birch, Beech, Lime and Murgwort pollen. Each dataset consists of 15 volumetric images where the porates have been manually labeled by an expert for training and evaluation. Note that the number of porates varies between but also within the different species. The image sizes for the Alder, Birch and Murgwort pollen are about 80^3 voxels. For the Beech and Lime pollen we have about 110^3 and 120^3 , respectively. One voxel corresponds to $0.4\mu\text{m}$.

Table 1. a) Filter parameters in our Experiments. The filter parameters are given in voxel size (voxel size $\approx 0.4\mu\text{m}$). **b)** Performance on the Birch dataset when using different normalization strategies.

(a) Filter Parameters					(b) Different Normalizations	
data	Filter	L	Filter Parameter	Feature Parameters	Recall	Precision
a) Alder	SHOG	5	$\eta = 5$	$\{d, \sigma\} = \{2, 1\}, \{4, 2\}$	0.9	0.2
b) Birch	SHOG	5	$\eta = 5$	$\{d, \sigma\} = \{2, 1\}, \{4, 2\}$	0.8	0.4
c) Beech	SHOG	5	$\eta = 6$	$\{d, \sigma\} = \{2, 1\}, \{4, 2\}$	0.7	0.6
d) Lime	SHOG	5	$\eta = 7$	$\{d, \sigma\} = \{4, 2\}, \{6, 2\}$	0.6	0.8
e) Mugwort	SHOG	5	$\eta = 4$	$\{d, \sigma\} = \{2, 1\}, \{4, 2\}$	0.5	1.0

Apart from our SHOG-Filter, we consider two other trainable filters, namely the original Harmonic Filter [4] and the Bessel Filter [9]. For all filters the experimental setup is as follows: We conduct 5 different experiments based on the different pollen datasets. For each experiment we use one single dataset for training. The training sets are depicted in figure 6 (the labels for training are marked by a red circle). In this step the filters optimize their parameter (least square fit) such that the filter response is most similar to the labeling. The filters are then applied to the remaining datasets for evaluation.

For all filters some parameters must be set manually. For finding the optimal parameters we follow the way proposed in [9]. The optimization is done by varying the parameters during several training steps until the Euclidean distance of the filter responses to the training label-images cannot be further reduced significantly. For the SHOG-Filter we must determine the following parameters: A filter degree $L \in \mathbb{N}$ that limits the number of expansion coefficients of the SHOG filter \mathbf{A}_w^ℓ , $\ell \leq L$. Furthermore, for the Harmonic Filter framework we need to set the parameter η that steers the size of a Gaussian window that restricts the SHOG features that can contribute to a local filter response. We finally must define one or more window function w for the SHOG itself. We observed that for the given data two nested Gaussian windowed spheres $w(\mathbf{r}, d, \sigma) := e^{-\frac{(\|\mathbf{r}\| - d)^2}{2\sigma^2}}$ lead to the best performance. We exemplarily illustrate the size and shape of the window functions we use for Alder datasets in figure 6(b). The parameters for all filters are summarized in table 1(a).

The gradient magnitude highly varies over a wide range due to variations in illuminations and in particular in volumetric biomedical images due to absorption and occlusion effects. Similar to [2] we observed that unnormalized gradients lead to poor performance. See figure 1(b) for results on the Birch dataset. Normalizing SHOG with respect to the l_2 -norm or with respect to the standard deviation of the local intensity values [4] increases the performance significantly. However, we achieve the best performance when almost neglecting the gradient magnitude and only considering the gradient orientation by performing a gamma correction of the gradient field, whereas $\mathbf{g}_\gamma = \|\mathbf{g}\|^\gamma \hat{\mathbf{g}}$. For our experiments we use $\gamma = 0.1$.

Figure 7 lists the PR graphs showing the performance of the filters in all 5 experiments for both tolerating a 8 voxel ($\approx 3.27\mu m$) displacement to the ground truth and tolerating only a more strict 4 voxel ($\approx 1.64\mu m$) displacement. For a better comparison we list the equal error rate (EER) for all experiments in table 1(a). We additionally show qualitative results of the SHOG filter in figure 8. The SHOG-Filter produces only clear responses at the correct porate positions. All remaining regions of the pollen are successfully suppressed. Moreover, thanks to the SHOG-Filter we detected a pollen belonging to a related but different species that accidentally found its way into the database. The structure of the porates differ strongly from the training set and thus the filter didn't respond at all (figure 8 d)).

In figure 9 we exemplarily show detections on the Beech dataset corresponding to the Harmonic Filter. Here we can observe that the harmonic filter clearly can detect the porates but produces a lot of false positive detections within the pollen. Similar for the remaining pollen species having high variations within the pollen. We observed similar problems for the Bessel Filter. We were not able to suppress responses on the inner structures of the Beech and Lime pollen while still getting clear responses at porate positions. The main difference of the SHOG-Filter is that the gamma normalized SHOG features are mainly comprising the gradient orientations. Thus SHOG is very robust against non-linear, local illumination and contrast changes. In contrast, the Bessel and Harmonic Filters are both indirectly encoding the gradient magnitudes in their features and thus are sensitive to non-linear illumination changes and noise.

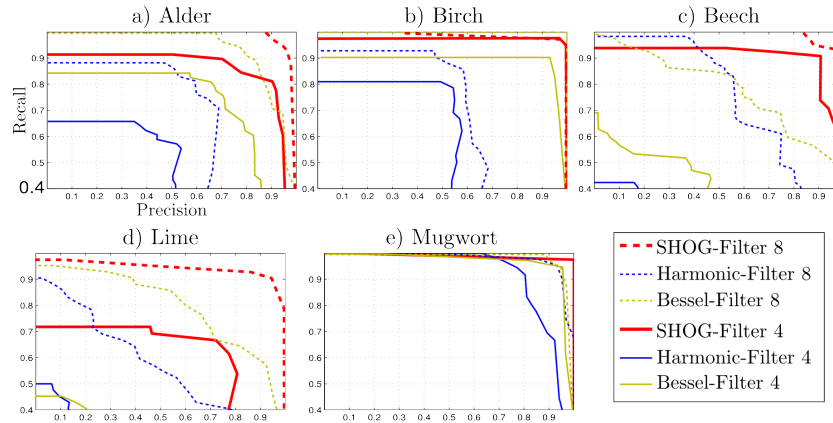


Fig. 7. The PR-curves are showing the performance of our SHOG-Filter compared to two existing state-of-the-art approaches for all 5 datasets. The dashed lines show the performance when tolerating an 8-voxel displacement to the ground-truth. The straight line shows the performance when only tolerating a 4-voxel displacement. We additionally show the maximum intensity projections of the raw filter responses of the SHOG filter in figure 8, clearly emphasizing the superior performance the SHOG Filter.

4 Conclusions

In this paper, we have presented a way to efficiently compute dense spherical HOG (SHOG) descriptors in volumetric images. Upon these descriptors we extended the Harmonic Filter to comprise the SHOG features instead of simple Gaussian derivatives to benefit from both a dense, robust and discriminative description in terms of gradient histograms and the trainable voting scheme of the Harmonic Filter which can be realized very computational and memory efficient.

We have shown the superior detection performance of our filter compared to previous state-of-the-art trainable 3D filters. These results are very promising in connection with the growing importance of volumetric data especially in the life sciences. In order to foster further research and experiments, we will provide public executables for using the proposed filter upon acceptance of this paper.

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

References

1. Buch, N., Orwell, J., Velastin, S.: 3d extended histogram of oriented gradients (3dhog) for classification of road users in urban scenes. In: BMVC. London, UK (2009)
2. Dalal, N., Triggs, B.: Histograms of oriented gradients for human detection. In: Proc. of the CVPR. pp. 886–893 (2005)
3. Lowe, D.G.: Distinctive image features from scale-invariant keypoints. *Int. J. Comput. Vision* 60(2), 91–110 (November 2004)
4. Reiser, M., Burkhardt, H.: Harmonic filters for generic feature detection in 3D. In: Proc. of the DAGM. pp. 131–140. LNCS, Springer, Jena, Germany (2009)
5. Ronneberger, O., Wang, Q., Burkhardt, H.: 3D invariants with high robustness to local deformations for automated pollen recognition. In: Proc. of the DAGM. pp. 455–435. LNCS, Springer, Heidelberg, Germany (2007)
6. Rose, M.: *Elementary Theory of Angular Momentum*. Dover Publications (1995)
7. Scherer, M., Walter, M., Schreck, T.: Histograms of oriented gradients for 3d object retrieval. In: Proceedings of the WSCG. pp. 41–48. Plzen, Czech Republic (2010)
8. Schlachter, M., Reiser, M., Herz, C., Schluermann, F., Lassmann, S., Werner, M., Burkhardt, H., Ronneberger, O.: 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. Skibbe, H., Reiser, M., Ronneberger, O., Burkhardt, H.: Spherical Bessel filter for 3D object detection. In: Proc. of the 8th ISBI. Chicago, Illinois, USA (April, 2011)
10. Thomson, J.J.: On the structure of the atom: an investigation of the stability and periods of oscillation of a number of corpuscles arranged at equal intervals around the circumference of a circle; with application of the results to the theory of atomic structure. *Philosophical Magazine Series 6* 7(39), 237–265 (March 1904)




Fig. 8. Detection of airborne-pollen porates in 5 different datasets (see figure 6). We show the maximum intensity projection (MIP) of the raw filter responses of the SHOG Filter together with the MIP of the detections (colored images) after thresholding (local maxima, threshold selected with respect to the ERR). The SHOG-Filter clearly only response to the porates. Furthermore, the SHOG-Filter didn't respond to a pollen that has accidentally found its way into the database (red mark).




Fig. 9. Detection of a Harmonic Filter for the Beech dataset (compare to figure 8 c)). The Harmonic Filter detects the porates. However, we where not able to avoid responses on inner-pollen structures. Similar for the Bessel Filter on the Lime and Beech dataset.