

FAST AND ROBUST SEGMENTATION OF SPHERICAL PARTICLES IN VOLUMETRIC DATA SETS FROM BRIGHTFIELD MICROSCOPY

Olaf Ronneberger, Qing Wang*, and Hans Burkhardt*

Universität Freiburg, Institut für Informatik, Lehrstuhl für Mustererkennung und Bildverarbeitung,
Georges-Köhler-Allee Geb. 052, 79110 Freiburg, Deutschland
{ronneber,qwang,burkhardt}@informatik.uni-freiburg.de

ABSTRACT

In this article we present an approach for a precise segmentation of spherical particles in transmitted light image stacks. A main goal was its fast operation and a high robustness to occlusions and agglomerations of the particles. The system is based on a voting procedure that finds the centers and radii of the particles and a subsequent precise segmentation with an active contour approach. To meet the demands of an online pollenmonitor for high speed and low memory consumption a multi-scale approach was applied. The proposed techniques successfully segmented the pollen grains in a vast amount of different air samples (about 2.7TB of raw data). The results on one of the most cluttered samples are presented in this paper.

Index Terms— Spherical, segmentation, detection, Hough transform, snakes, pollen, pollenmonitor

1. INTRODUCTION

Motivation. Reliable detection and segmentation of touching and partially occluded objects in 3D volumetric data sets is a difficult problem, even if the data is recorded with high-end microscopes under well controllable laboratory conditions and the computation time does not matter.

In this paper we present a detection and segmentation technique for spherical particles, which was developed for the first prototype of an online pollenmonitor. The pollenmonitor is a fully automated machine for measuring real-time pollen concentration in the air. It was developed within the BMBF-funded project OMNIBUSS. It can collect and prepare air samples and record them with an integrated microscope. In online operation the hourly air samples have to be evaluated within one hour. For a reliable estimation of the pollen concentration about 2 m^3 of air have to be analyzed, which corresponds to an area of about 26 mm^2 on the air sample. Due to large size variances of pollen grains, and their important fine details, a 3D volumetric data set of the air sample is recorded. This data set consists of about 200 transmitted light image stacks, each containing $1392 \times 1040 \times 70$ voxels of size $0.3225 \times 0.3225 \times 1.5\mu\text{m}^3$. Therefore, in the online operation of the pollenmonitor each image stack of about 100 megavoxels needs to be completely processed in about 18 sec, or less. As segmentation is only one step in the recognition pipe, it should be finished in even shorter time.

A further challenge in this setup is the automated unattended operation on ambient air samples. It must be robust to occlusions with any type of particles that might be found in the ambient air. Moreover it was not possible to completely isolate the microscopy

part from the rest of the machine. Therefore vibrations from outside caused some jitter in the recorded image stacks.

Our work with this dataset was first presented in [1], where the emphasis was on the derivation of invariant features. Here we will concentrate on the detection and segmentation of pollen grains.

State of the art. A reliable segmentation of pollen in air samples is difficult and only a few publications describe approaches for the solution. Bonton et al. [2] use the color information after a special staining of the air samples, but such staining procedures are impractical for our automated system. Rodriguez-Damian et al. [3] tried different standard segmentation techniques, among them also a circular Hough transform, followed by a snake approach. Due to the use of manual prepared pure pollen samples without occlusions or agglomerations, the standard circular Hough transform was sufficient there. The standard snake failed for many pollen grains, because it was attracted to the wrong borders. Ranzato et al. [4] use the DOG (difference of Gaussian) interest point detector for the detection of pollen grains in ambient air samples. According to these authors, “a segmentation based approach has never given good results”. Therefore they only determine a coarse bounding box and use all pixels therein for feature extraction, such that the features are disturbed by neighboring dust. The obtained recognition rate of 64.9% recognition rate at a precision of 30% for 8 different pollen taxa are significantly lower than the best segmentation based approaches of 84.3% recognition rate at a precision of 96.7% for 33 different taxa [1].

2. APPROACH

The proposed approach is based on an extension of the generalized Hough transform [5] for the detection of spherical objects and a subsequent precise segmentation with an active contour approach.

There are basically three parts that need to be considered, when a hough transform shall be used: First, the selection of the voting pixels (usually done by an edge detector or thresholding), second, the contribution of this voter to the accumulator (usually the positions are computed from the local gradient direction, and the contribution is either constant or the gradient magnitude), and third, the way how the different votes are combined to the final vote (usually a simple summation). Due to speed and robustness considerations we had to replace each of these standard algorithms with faster or more robust ones.

The extraction of invariant features is based on the boundary of a pollen grain, which makes precise segmentation necessary. Active contour method is employed to find a smooth and complete boundary. But the standard method suffers from a severe problem: the snake is often attracted to undesired edges. The problem was solved

*First two authors contributed equally to this work

by using a modified Canny edge detection and by careful edge analysis.

3. DETECTION OF SPHERICAL OBJECTS

Detection of spherical objects is carried out in the following steps:

1. Calculating voting vector fields. The voting vector fields have to be chosen according to the characteristics of the dataset. The pollen data is recorded on image stacks by transmitted light microscopy, for which the recording properties in z-direction differ significantly from those in xy-direction. The borders in z-direction of the pollen grain are hidden by the diffraction patterns (see Fig. 1). It makes sense to omit the z-component of the gradient vector.




Fig. 1. Orthogonal slices of a transmitted light stack from a *Betula* pollen grain. Only the gradients in x- and y- direction should be used to find the object center. In z-direction the gradients are dominated by the diffraction patterns.

Furthermore, as can be found from Fig. 1, the borders of pollen grains are lines of different thickness. It motivated us to use line information instead of edge information for detection. We do it by replacing gradient magnitude with the largest eigenvalue of the Hessian, and by replacing gradient direction with the corresponding eigenvector.

The Hessian of a function $f(x, y)$ is defined as

$$Hf(x, y) = \begin{pmatrix} \partial^2 f / \partial x^2 & \partial^2 f / \partial x \partial y \\ \partial^2 f / \partial y \partial x & \partial^2 f / \partial y^2 \end{pmatrix}. \quad (1)$$

If the Hessian matrix for a data point is

$$\begin{pmatrix} a_{11} & a_{12} \\ a_{21} & a_{22} \end{pmatrix}, \quad (2)$$

its eigenvalues are

$$\lambda_{\pm} = \frac{1}{2} \left[(a_{11} + a_{22}) \pm \sqrt{\Delta} \right], \text{ where } \Delta = (a_{11} - a_{22})^2 + 4a_{12}a_{21}. \quad (3)$$

Theoretically Δ should be nonnegative as $a_{12} = a_{21}$. Numerically we can avoid $\Delta < 0$ by replacing a_{12} and a_{21} with their average.

If $\Delta = 0$, we assign the voting vector to be 0 as there is no preferred direction for voting. When $\Delta > 0$, λ_+ is the largest eigenvalue, we assign the magnitude of the voting vector \mathbf{v} to be $v = \lambda_+$ and its direction to be

$$\mathbf{e}_v = \frac{1}{\sqrt{a_{12}^2 + (a_{11} - v)^2}} \begin{pmatrix} a_{12} \\ a_{11} - v \end{pmatrix}, \quad (4)$$

which is the eigenvector corresponding to λ_+ .

2. Selecting voting points. The selection of voting points is made by short line detection in the xy-plane and requiring a local maximum in z-direction. This is done in two runs: 1.) In xy plane,

two neighboring points are found which lie on each side of the candidate point along the direction perpendicular to the vector direction (These two points usually do not lie on the grid and the vector value on them are interpolated bilinearly). The candidate point is kept only when it and the two neighboring points are all local maxima along the direction of the voting vector (see Fig. 2a). 2.) Compare each selected point with the selected points in the layers above and below, if it is not the maximum, it is also discarded.




Fig. 2. Selection of the voting points in two steps. a) The voting point must be located on a ridge of the eigenvalues. b) The ridge must be higher than the corresponding ridges in the layers above and below.

3. Voting. The votes are collected in an accumulator of five dimensions (x, y, z, r, d) . The first four parameterize the sphere and the fifth dimension is the voting direction and is quantized into n_{dir} bins. Each identified voter contributes to the accumulator. A direct rendering of the voting kernel into the accumulator is too slow. So a radially symmetric voting kernel is taken, such that each vote can be approximated by adding a delta peak into the accumulator and a subsequent convolution of the accumulator with the respective kernel after the voting is finished [6]. Here this convolution is approximated with a recursively computed iir-filter (infinite impulse response), which can be computed in linear time. The width of the kernel was adapted according to the expected deviation of the true shape from a perfect circle. The combination of the votes from the n_{dir} different directions at accumulator-position (x, y, z, r) is done by sorting the votes and taking the n th element, where n can be interpreted as the number of *required* voting directions.

4. Searching spheres. The objects are not very dense, we search the spheres by a two pass-algorithm, which first extracts a list of all candidates for a sphere by a local maximum criterion, sorts them according to their probability and then extracts the non-overlapping spheres from this list.

Speeding up by multi-scale approach. A naive implementation of the proposed methods needs great amounts of memory and computation time. But it was straightforward to apply a multi-scale approach. In the first step a rough estimate of the circle positions and radii are obtained on a down-scaled image and a coarsely quantized accumulator. In the second step, each found parameter set (radius, x, y, z) is refined by a direct search in its surrounding. In this refinement step a much higher number of voting directions can be used, which significantly improves the accuracy.

4. EXACT SEGMENTATION OF POLLEN GRAINS

It is important to know the exact boundary of a pollen grain to extract features for pollen recognition. The boundary is only searched on the sharpest layer near the center of a detected sphere. As parts of object border are often missing or not clear, we use active contours to find

smooth and complete boundaries. Among many different flavors of active contour methods, the one based on Gradient Vector Flow (will be called GVF snake later) is chosen due to its various advantages [7].

The input for GVF snake is an image which is computed from the image to be segmented and indicates where the snake should be attracted to. Although the border of the pollen grains usually appear as dark lines of various thickness, they should be completely included inside the boundary as their thickness is also characteristic of pollen species. Therefore we use edge images as input for the snakes, which are different from the detection step. The simplest edge image is the gradient magnitude. With it, however, the snake is often attracted to undesired position. The same is also true with edge images acquired with standard edge detection methods like Canny edge detection (see Fig. 4). This problem has its fundamental reason: Edges do not necessarily corresponds to the object boundary. To solve this problem, we take the following steps for segmentation (Illustration of these steps are given in Fig. 4):

1. Applying modified Canny edge detection. The modification is with the gradient vector image, on which Canny edge detection works. As pollen grains from the species we are interested in have a nearly round shape, the edges that are approximately perpendicular to the radial direction (relative to the center obtained from detection step) are more relevant. We replace the gradient vector with its radial component so that the strength of other edges will be reduced. To avoid the influence of the inner contour of the dark lines, the magnitude of the gradient vector is greatly reduced when it points approximately to the center instead away from it. In this way, the inner contours will be detected as very weak edges if they can be detected at all.

2. Analyzing edges. The edges obtained from last step undergo a careful analysis. The curvature of each edge pixel, the location of each edge relative to the center and to other edges as well as the strength of the edges are taken into consideration. Weight is assigned to edge pixel according to carefully-chosen criterion. The edge pixels most probably corresponding to the boundary or those having no much influence to the snake are left with high weights. As a result, a much clearer weighted edge image is obtained.

3. Running GVF-snake on the weighted edge image. The initial position of the snake is chosen to be near the circle found from the detection step.

5. RESULTS

5.1. Detection Results

To obtain quantitative results one of the most cluttered air samples of the pollenmonitor data set (containing lots of dust particles and agglomerated pollen grains) was chosen. It was recorded during a test-operation and contains therefore only 25 image stacks recorded with the integrated microscope of the pollenmonitor in transmitted light mode (20x lens, 0.8 N.A.). Each stack has 1392x1040x70 voxels (≈ 100 megavoxels). An example is depicted in Fig. 3a. The computed eigenvalues (only the higher one) of the Hessian for each position are shown in Fig. 3b. The selected voters are shown in Fig 3c. In the traditional Hough transform the votes are simply added (Fig. 3d). Our approach requires that votes come from different directions: Correct maxima are emphasized while wrong maxima (e.g. in the upper right corner) disappear (Fig. 3e). For validation the obtained circles are compared to the manually labeled correct centers of all pollen (see Fig. 3f).




Fig. 3. Detection of spherical pollen grains in a part of a transmitted light image stack of an air sample

First the direct implementation (without the multi scale approach) was tested for an accumulator with only 11 different radii between $10\ \mu\text{m}$ to $60\ \mu\text{m}$. The computation time for one stack was about 5 min using a single core of a QuadCore Xeon X5365 with 3GHz. The peak memory usage was about 24GB of RAM. Using the multi-scale approach the same volume was processed in about 24 seconds, and needed only about 1GB of RAM. Moreover this approach delivers the circles radius and position with $0.32\ \mu\text{m}$ (one pixel) precision.

The overall recall and precision of the method was evaluated by counting a found circle as a hit, if the center did not differ more than $4\ \mu\text{m}$ (in xy -direction) from the manually labeled experts center. Pollen that are closer than $11\ \mu\text{m}$ to the border were discarded. The overall recall was 98.2%. The missed 14 pollen grains (from the 789) were all highly cluttered and even hard to detect for the expert (similar to the pollen grain, which is nearly completely covered by the black dust particle near the center in Fig. 3f). The number of false alarms (only 148) is very low compared to the traditional approaches, and many of them belong to other spherical airborne particles with are not pollen grains. The details are listed in table 1.

Table 1. Results of the proposed detector

True number of pollen	N_{expert}	789
Found number of pollen	N_{computer}	923
Number of hits	N_{hits}	775
Precision	$N_{\text{hits}}/N_{\text{computer}}$	84.0%
Recall	$N_{\text{hits}}/N_{\text{expert}}$	98.2%

5.2. Results of exact segmentation

Examples of exact segmentation of pollen grains are given in Fig. 4. To get a quantitative impression how well it works, we evaluate it on the same air sample used above for detection evaluation. Here only pollen grains locating completely inside the images are taken in account. That is 689 objects. Among them, only 35 are wrongly segmented (some examples are shown in Fig. 5). The other 654 pollen grains have similar segmentation quality as shown in the last row in Fig. 4. That counts to about 95% correct segmentation. And this is with a very cluttered air sample, where segmentation is very difficult. On many other air samples, we have almost 100% correct segmentation. Furthermore, as mentioned above, correct segmentation is essential for pollen recognition. The good results of the pollen recognition [1] is also an indication of the success of the segmentation.

6. CONCLUSION AND OUTLOOK

The proposed algorithms perform a robust, reliable and very fast segmentation (only 24 sec. for a 100 megavoxel image stack) of pollen grains in a highly cluttered surrounding.

The system can be used for any other spherical particles. First experiments on DIC images of *Drosophila* Schneider cells have shown nearly perfect results. Due to the fact that the voting is completely based on local maxima criteria, the proposed algorithms allow to combine votes from different data channels (e.g. a channel recorded by fluorescence microscopy and a channel recorded by transmitted light microscopy) without the need to define arbitrary factors for the recombination.

7. REFERENCES

- [1] O. Ronneberger, Q. Wang, and H. Burkhardt, "3d invariants with high robustness to local deformations for automated pollen recognition," in *Proc. of the DAGM*. 2007, vol. 4713/2007 of *LNCS*, pp. 425–435, Springer Berlin / Heidelberg.
- [2] P. Bonton, A. Boucher, M. Thonnat, et al., "Colour image in 2d and 3d microscopy for the automation of pollen rate measurement," *Image Analysis & Stereology*, vol. 21, pp. 25–30, march 2002.
- [3] M. Rodriguez-Damian, E. Cernadas, A. Formella, M. Fernandez-Delgado, and Pilar De Sa-Otero, "Automatic detection and classification of grains of pollen based on shape and texture," *Systems, Man and Cybernetics, Part C, IEEE Transactions on*, vol. 36, no. 4, pp. 531–542, 2006.
- [4] M. Ranzato, P. E. Taylor, J. M. House, R. C. Flagan, Y. LeCun, and P. Perona, "Automatic recognition of biological particles in microscopic images," *Pattern Recognition Letters*, vol. 28, no. 1, pp. 31–39, January 2007.
- [5] D. H. Ballard, "Generalizing the hough transform to detect arbitrary shapes," *Pattern Recognition*, vol. 13, no. 2, pp. 111–122, 1981.
- [6] J. Schulz, T. Schmidt, O. Ronneberger, H. Burkhardt, T. Pasternak, A. Dovzhenko, and K. Palme, "Fast scalar and vectorial grayscale based invariant features for 3d cell nuclei localization and classification," in *Proc. of the DAGM*, Berlin, 2006.
- [7] Chenyang Xu and J.L. Prince, "Snakes, shapes, and gradient vector flow," *Image Processing, IEEE Transactions on*, vol. 7, no. 3, pp. 359–369, 1998.




Fig. 4. Exact segmentation of pollen grains. The first row shows the original images. The second and third rows show segmentation results with gradient magnitude and original Canny edges as edge images respectively. The last 6 rows show the important steps of the proposed methods. They are gradient magnitude, magnitude of modified gradient, edge images obtained with modified Canny edge detection, weighted edge images, initial snake position and the final boundaries.




Fig. 5. Examples of bad segmentation