2017 IEEE International Conference on Pattern Recognition (ICPR)

Pollen Recognition in Optical Microscopy by Matching Multifocal Image Sequences

Roman Filipovych*, Amar Daood†, Eraldo Ribeiro* and Mark Bush‡
*Computer Vision Laboratory Department of Computer Sciences
†Department of Electrical and Computer Engineering
‡Department of Biological Sciences
Florida Institute of Technology, Melbourne, FL 32951, U.S.A.

***Abstract*—We present a novel technique for pollen identification from sets of multifocal image sequences obtained from optical microscopy. Our algorithm analyzes the visual texture of pollen grains for each focal image, and performs identification using a fast sequence-matching algorithm. Although we develop a pollen-recognition protocol, the method is applicable to other microscopy object-recognition tasks. The proposed method requires little manual interaction, and does not rely on specialized imaging procedures such as florescence and deconvolution. We test our method on images of tropical fossil pollen.**

I. Introduction

The identification and counting of pollen grain are both central and time-consuming tasks in palynology research. Despite recent developments in modern computing and automation, automatic pollen-grain identification remains a largely unsolved problem. Representing the rich visual variability for hundreds of pollen types is major challenge. The problem's complexity is further increased by the pollen's 3-D nature and the small depth of field available in optical microscopes. The limited depth of field means that only a small part of the pollen surface is in focus during the analysis. Thus, distinguishing grain features are only visible when imaged at the correct focal plane (Figure 1, top left). Traditionally, pollen recognition by visual inspection is done by measuring morphological surface attributes such as the presence of pores and salient geometric features [1] (e.g., ridges and spikes). This identification methodology is suitable for human inspection but it does not directly leads to computational algorithms.

In the last two decades, most automatic pollen recognition methods have focused on measurements of visual texture, contour, and in some cases, three-dimensional shape. For example, Carrion et al. [2] described a method for classifying honeybee pollen using multiscale texture signatures. Similar approaches include the works by Li et al. [3] and by Zhang et al. [4]. These algorithms used texture descriptors calculated at small fixed-sized image subregions. However, image subregions may not capture the full complexity of the texture across a whole grain. In addition, some pollen grains cannot be distinguished by texture measurements only.

Some recent methods analyze the pollen's three-dimensional shape, which is reconstructed from a sequence of confocal-microscopy images. Ronneberger et al. [5] classified aerial pollen by reconstructing grains using fluorescence confocal microscopy. Although 3-D reconstruction provides information about the whole grain, the method requires expensive confocal-microscopy hardware. Furthermore, fluorescence microscopy is not suitable for fossil pollen that are the basis of most paleoecology studies. A similar approach was developed by Boucher et al. [6] who used expert palynological knowledge for feature identification by measuring morphological features such as size and number of pores, ridges, and grain's diameter. Of course, measurements of morphological features require both accurate 3-D reconstruction and feature detection.

Pollen grains can be also distinguished by the shape of their outer contour [7], [8]. These methods used a pollen's contour descriptor to classify certain types of pollen that have distinguishable shapes. However, the contour of a 3-D object in images changes with the imaging viewpoint. This apparent deformation, along with occasional collapsing of the grain's surface limit the practical use of contour shape as a main feature for automatic pollen identification.

We present a novel pollen-identification technique that uses multifocal image sequences from optical microscopy. Figure 1 shows a graphical overview of the method. Our algorithm analyzes the pollen's surface texture across varying focal planes. The method consists of three main steps. First, we use an active-contour segmentation method to extract the entire subvolume containing the pollen grain regions from the multifocal image sequence. Image regions in these subvolumes contain relevant visual information (e.g., texture, blob features, and contour fragments) while discarding unwanted neighboring information (e.g., background clutter and debris). Secondly, we represent the texture appearance inside the segmented multifocal subvolume using a histogram of responses from a set of representative filters. This appearance model carries information about both the scale and orientation of pollen surface features for each focal plane. Finally, pollen identification is performed by matching sequences of multifocal appearance descriptors using a robust sequence-alignment method. We test our method on images acquired at Florida Tech's Paleoecology Laboratory using an optical microscope. Experiments show the potential of our method for pollen recognition (Section VI).

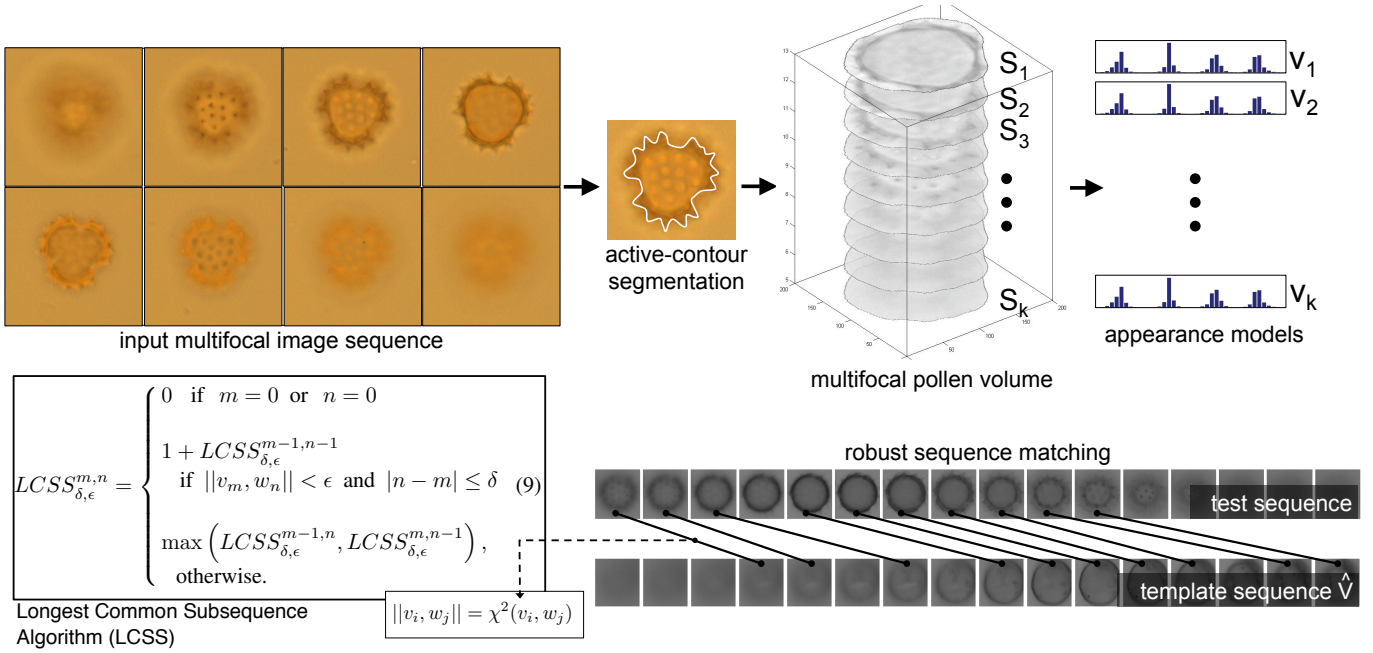

Fig. 1. Overview of our method. A multifocal image sequence containing a pollen grain (*Darwiniothamnus* pollen shown as an example) is given as input. Pollen regions in each image are segmented using an active-contour method, which creates a multifocal pollen volume. Each pollen focal plane is then represented by a frequency histogram of filter responses forming a sequence of appearance models. The appearance-model sequence representing the pollen is matched to the ones in a database using the robust sequence-matching algorithm.

II. Contributions

We propose a method that uses visual information from sequences of focal planes to identify pollen grains in optical microscopy images. Our method does not require images to be acquired at fixed focal-length intervals. We propose a multifocal volume segmentation method that allows us to extract both textured regions inside the pollen boundary and boundary-shape information. The visual information inside this segmented volume is then used for constructing an appearance model of the pollen's surface texture for each focal plane. The identification step considers the multifocal volume as a data sequence (sequence of appearance models). Although we focus on the recognition of fossil pollen grains, an extension to other microscopic objects is straightforward. Finally, domain-specific information such as symmetry, contours, and colors can be also included into the model. An overview of the main steps of our method is given below:

a) Step 1 - Multifocal volume segmentation.: In this step, we automatically extract the region inside the boundary of the pollen grain for the entire multifocal volume using an active-contour segmentation method. Here, rather than processing each focal slice individually, we use information from the entire multifocal volume to support the segmentation task. By treating the entire volume as a single object, our method can improve segmentation by combining useful data from different focal slices while discarding inconsistent noisy pixels. Section III describes the segmentation method.

b) Step 2 - Modeling the texture content.: In our appearance model, each pixel in the polar-mapped images is represented by a vector of Gabor filter responses [9]. Filter response maps are obtained by convolving the original focal slices with a bank of multi-scale and multi-orientation filters. Here, the bank of filters represent a set of basic shapes at varying scales that, when convolved with the pollen images, will produce vectors describing the information in each pixel in terms of filter-bank components. This filter-based representation also encodes, to some extent, information about the pollen contour at each focal slice. Our filter bank is a reduced representation of the one proposed by Leung and Malik [10], [11]. Section IV describes our filter-based appearance model.

c) Step 3 - Classification by sequence matching.: Once the texture descriptor for the segmented multifocal pollen sequence is obtained, identification and detection are accomplished using a sequence-matching method. We use a modified version of the longest-Common Subsequence (LCSS) matching scheme introduced by Vlachos et al. [12]. Details of this step are described in Section V.

III. Multifocal Pollen Volume Segmentation

The first step of our method is to detect the boundary of the multifocal sequence (multifocal volume) containing the target pollen grain. Challenges to be overcome in detecting pollen boundaries in images include: (1) the boundaries can be undefined or blurred due to microscope's small depth of field, (2) debris might be present in the area surrounding the grain, and (3) background and foreground image regions might have similar intensity and color.

We address these problems by using information from the whole multifocal image sequence. Our goal is to detect the surface separating pollen grain regions from the surrounding multifocal volume. We use a curve-propagation approach often applied in medical-image segmentation [13], [14]. The probabilistic curve-propagation framework presented in [13] is particularly appealing as it accounts for nondeterminism in the input image data. The method estimates a statistical description of both the background and foreground image regions while propagating a 2-D deformable curve that adapts itself to the shape of the pollen's multifocal subvolume. The deformable curve starts at the edges of the rectangular multifocal images, and iteratively deforms its shape while trying to maximize the foreground statistics inside the closed boundary surface. This segmentation method uses information from the entire multifocal sequence. Next, we provide a formal description of the boundary-segmentation method.

A. Variational image segmentation

We begin by describing the variational formulation proposed by Rousson and Deriche [13] for the segmentation of a multi-channel image. Let $S = (S_1, \ldots, S_k)$ be a sequence of multifocal images of a pollen grain obtained using measurements at k consecutive depths (i.e., image S_k is acquired at the lowest depth). We treat the multifocal sequence as a single multi-channel image. Thus, each data point is represented by a vector of pixel values from all focal slices. In the formulation proposed in [13], the segmentation of foreground and background regions in S can be accomplished by maximizing the posterior probability $p(\Omega)$, where $\Omega = \{\Omega_b, \Omega_f\}$ is the image domain divided into two partitions representing the background region Ω_b and the foreground region Ω_f. Additionally, let $u(\mathbf{x})$ be the measurement value at spatial location $\mathbf{x}$, where $u : \mathbb{R} \rightarrow \mathbb{R}^n$. Let p_b and p_f, be the probability density functions (PDF) for the value $u(\mathbf{x})$ to be in regions Ω_b and Ω_f, respectively. Following the method described in [13], segmentation is obtained by minimizing the following energy functional:

$$\begin{aligned} E(\Omega_b, \Omega_f) = &- \int_{\Omega_b} \log p_b\left(u\left(\mathbf{x}\right)\right) d\mathbf{x} \\ &- \int_{\Omega_f} \log p_f\left(u\left(\mathbf{x}\right)\right) d\mathbf{x} \\ &+ \text{length}(\partial\Omega), \end{aligned} \tag{1}$$

where $\partial\Omega$ is the boundary between Ω_b and Ω_f, and length$(\partial\Omega)$ is a normalization term (i.e., boundary energy) added to prevent oversegmentation. Prior to performing curve propagation, we preprocess all images using the non-linear diffusion filtering method in [15]. This pre-processing step helps reduce noise while preserving descriptive geometrical details present on the object's surface. In our algorithm, filtering is performed on the gray-level versions of the multifocal slices. Figure 2 shows an example of a non-linearly diffused sequence.

Noise and debris may still remain in the images, and can cause incorrect segmentation of the target pollen. Figure 3 shows a segmentation example for a multifocal sequence containing a non-uniform background region. Here, artifacts can cause significant gradient variations in background regions of the diffused images (e.g., dark and bright regions separated by the dashed curve in Figure 3.b), leading the segmentation method based on Equation 1 to segment dark and white regions well while failing to extract the pollen grain (Figure 3.c). Also, this method does not automatically label background and foreground regions. Thus, a post-segmentation labeling procedure is required.

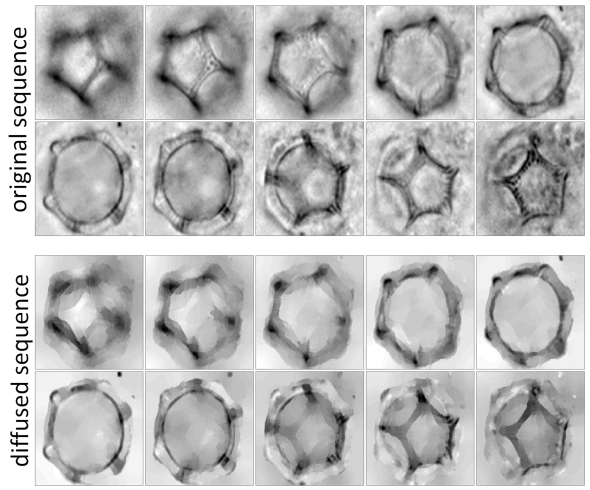

Fig. 2. Nonlinearly diffused multifocal sequence. The diffused sequence contains less noise while preserving edge information.

To resolve the background- and foreground-labeling ambiguity, we incorporate a non-informative background prior into the segmentation process by adding a regularization term to the background energy in (1). This term measures the information content for pixels inside region Ω_b (in the sense of Shannon's theorem [16]). We can then re-write Equation 1 as:

$$\begin{aligned} E(\Omega_b, \Omega_f) = -\mathcal{H}\left(\Omega_b\right) &- \int_{\Omega_b} \log p_b\left(u\left(\mathbf{x}\right)\right) d\mathbf{x} \\ &- \int_{\Omega_f} \log p_f\left(u\left(\mathbf{x}\right)\right) d\mathbf{x} \\ &+ \text{length}\left(\partial\Omega\right), \end{aligned} \tag{2}$$

where $\mathcal{H}\left(\Omega_b\right)$ is the information entropy calculated in the background region Ω_b. The background entropy is given by:

$$\mathcal{H}\left(\Omega_b\right) = -\int_{\Omega_b} p_b\left(u\left(\mathbf{x}\right)\right) \log\, p_b\left(u\left(\mathbf{x}\right)\right) d\mathbf{x}. \tag{3}$$

The entropy term in Equation 2 is designed to enforce the background region's non-informity. We substitute (3) into (2),

and combine the integrals to obtain:

$$E(\Omega_b, \Omega_f) = -\int_{\Omega_b} (1 - p_b(u(\mathbf{x}))) \log p_b(u(\mathbf{x}))d\mathbf{x} - \int_{\Omega_f} \log p_f(u(\mathbf{x}))d\mathbf{x} + \text{length}(\partial\Omega). \quad (4)$$

The integrals in (4) can be extended to the entire multifocal volume by using the level-set function $\phi : \Omega \rightarrow \mathbb{R}$ defined in terms of the distance transform $\mathcal{D}$ as:

$$\begin{cases} \phi(\mathbf{x}) = \mathcal{D}(\mathbf{x}, \partial\Omega), & \text{if } \mathbf{x} \in \Omega_b \\ \phi(\mathbf{x}) = -\mathcal{D}(\mathbf{x}, \partial\Omega), & \text{if } \mathbf{x} \in \Omega_f. \end{cases} \quad (5)$$

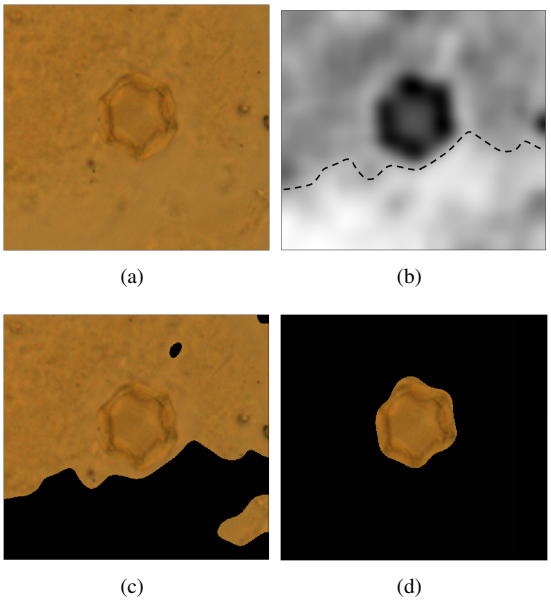

Fig. 3. Segmentation of *Alternanthera* pollen. A single image of the segmented multifocal volume is shown. Segmentation is performed on the whole volume. (a) Original image, (b) Filtered image, (c) Segmentation without using a background prior (i.e., Equation 1), and (d) Segmentation using our background prior (i.e., Equation 4).

Fig. 4. Pollen grains and corresponding detected contours using our method. Contours are displayed for a single slice in the multifocal volume. Actual segmentation is performed on the entire sequence.

Using the regularized form $H(z)$ of the heaviside function as proposed in [14], we obtain the following evolution equation for level-set surface ϕ:

$$\phi_t(\mathbf{x}) = \delta(\phi(\mathbf{x})) \times [\nu.\text{div}\left(\frac{\nabla\phi}{|\nabla\phi|}\right) + (1 - p_b(u(\mathbf{x}))) \log p_b(u(\mathbf{x})) - \log p_f(u(\mathbf{x}))], \quad (6)$$

where $\delta(\phi) = H'(\phi)$. As in [13], for simplicity, we assume that the distributions p_b and p_f are Gaussian. In general, gaussianity assumption may not entirely hold for distributions p_b and p_f in a particular focal slice. But, because we segment all focal slices simultaneously, our method works well, while keeping the model simple. The densities' parameters are updated at each iteration of Equation 6 as follows:

$$\begin{cases} \mu_i(\phi) = \frac{\int_\Omega u(\mathbf{x})\mathcal{Y}_i d\mathbf{x}}{\int_\Omega \mathcal{Y}_i d\mathbf{x}} \\ \Sigma_i(\phi) = \frac{\int_\Omega (\mu_i - u(\mathbf{x}))(\mu_i - u(\mathbf{x}))^T \mathcal{Y}_i d\mathbf{x}}{\int_\Omega \mathcal{Y}_i d\mathbf{x}}, \end{cases} \quad (7)$$

with $\mathcal{Y}_b = H(z)$ and $\mathcal{Y}_f = 1 - H(z)$. Here, $\mu_i(\phi)$ and $\Sigma_i(\phi)$ are the mean vectors and covariance matrices of p_b and p_f. When updating the parameters in (7), we assume statistical independence of spatial information provided by different focal slices. This assumption simplifies the computations of the inference process. Figure 4 shows examples of detected contours using our method. The output of this segmentation process is a volume containing the pollen grain regions. With the segmented multifocal volume of the pollen at hand, we can create a representation of pollen's surface texture in all focal images.

IV. Modeling the Texture Appearance

The pollen's surface contains descriptive features including the symmetric placement of blob-shaped regions and elongated shapes, that can appear at various scales and orientations. We represent the visual texture in these images in terms of the responses of a bank of multi-scale and multi-orientation filters. Each image is convolved with the filters in the bank.

Let $F = (f_1, \ldots, f_{N_f})$ be a filter bank containing N_f filters of different orientations and shapes. The filter bank contains 7 filters (i.e., $N_f = 7$) including 4 filters with different orientations at a single scale and 3 isotropic (i.e., rotation-invariant) filters of different scales.

Image S_i is then represented by a feature vector $\mathbf{r} = (r_1, r_2, \ldots, r_{N_f})^\mathsf{T}$ describing local texture appearance in terms of filter responses. The components in $\mathbf{r}$ are obtained by convolving filter bank F with image S as follows:

$$r_j = f_j * S, \qquad j = 1, \ldots, N_f. \quad (8)$$

where the symbol $*$ denotes the convolution operation. Figure 5(c–f) shows the results of convolving the original sequences in Figure 1 with some of the multi-scale multi-orientation filters in filter bank F.

Pollen appear at various sizes and rotations when seen under the microscope. To reduce the effects of scaling and orientation, we convert the segmented pollen regions into a scale-normalized representation using polar mapping before

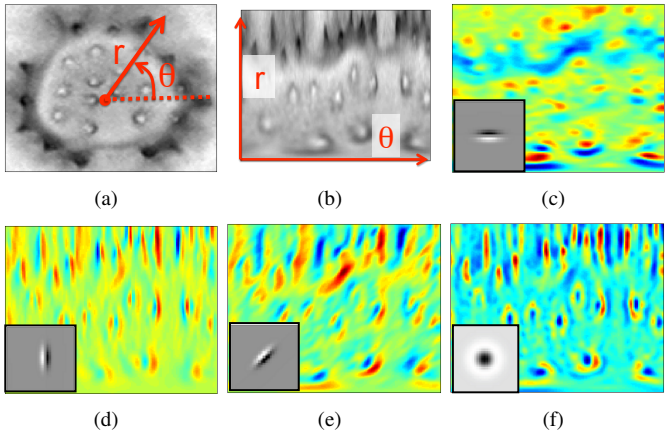

Fig. 5. Normalized representation of pollen texture. (a) Original pollen grain. (b) Polar representation (r, θ), where r is the distance from a boundary pixel to the estimated centroid of the boundary, and θ is the angle of r. (c–f) Responses from four filters in F. The insets display the corresponding filters (actual filter size is 64×64). Filter responses for each focal image in sequence S were obtained using Equation 8.

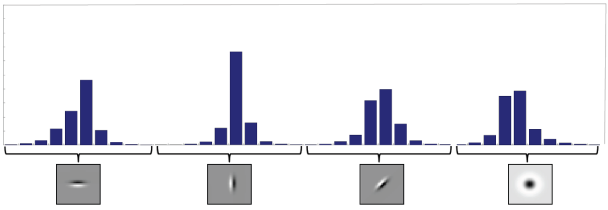

Fig. 6. Example of appearance model of a single focal slice. Model consists of concatenated normalized frequency histograms of filter responses (Figure 5.c–f).

applying filter convolution. This transformation flattens the near-circular pollen region into a rectangular shape of predefined size. Here, each segmented pollen is represented by a distance-angle map (r, θ), where r is the distance from boundary pixels to the grain's estimated centroid, and θ is the angle of vector r. The resulting polar map is scale invariant, and rotation of the grain is mapped into a simple translation. Indeed, in the specific case of our application, rotation ambiguity is significantly reduced by pollen texture inherent symmetry. Figure 5(b) shows an example of the polar mapping for the pollen shown in Figure 5(a).

A simple appearance descriptor for each focal images is then obtained by concatenating the normalized frequency histograms of filter response values for all filters. Figure 6 illustrates this descriptor for a single focal slice. Finally, each multifocal volume is represented by a data vector sequence $V = (v_1, \ldots, v_k)$, where v_i is a descriptor for the filter-based decomposition of the i-th polar-mapped focal slice.

V. Classification by Sequence Matching

The Dynamic-Time Warping (DTW) algorithm [17] is widely used for aligning sequences of data points. DTW provides a similarity measure between two time series while achieving sequence matching using temporal warping [18]. However, a key challenge in matching multifocal pollen sequences is that different focal lengths may be required to produce similar focal slices of two pollen grains of the same kind. In this case, standard DTW-based alignment will be significantly distorted as the algorithm will attempt to assign point correspondences to all data points. Thus, rather than using the traditional DTW alignment, we use the more robust longest Common Subsequence (LCSS) matching scheme introduced in [12]. We begin by defining a distance measure between two sequences. Let $V = (v_1, \ldots, v_m)$ and $W = (w_1, \ldots, w_n)$ represent two multifocal image sequences for two different pollen grains. The measure used in the LCSS scheme can be defined for an integer δ and a real number $0 < \epsilon < 1$ as follows:

$$LCSS^{m,n}_{\delta,\epsilon} = \begin{cases} 0 \quad \text{if} \quad m = 0 \ \text{ or } \ n = 0 \\ 1 + LCSS^{m-1,n-1}_{\delta,\epsilon} \\ \quad \text{if } ||v_m, w_n|| < \epsilon \ \text{ and } \ |n - m| \leq \delta \\ \max\left(LCSS^{m-1,n}_{\delta,\epsilon}, LCSS^{m,n-1}_{\delta,\epsilon}\right), \\ \quad \text{otherwise.} \end{cases} \tag{9}$$

Here, $LCSS^{i,j}_{\delta,\epsilon} = LCSS_{\delta,\epsilon}((v_1, \ldots, v_i), (w_1, \ldots, w_j))$, and $||.,.||$ is the distance between points v_m and w_n. The parameter ϵ is the threshold on the maximum distance between a pair of matched points, and δ controls the maximum number of consecutive points without a match. In the original LCSS method in [12], the parameter ϵ controls the maximum absolute difference between each of the data points. The similarity between the two sequences is given by:

$$D_{\delta,\epsilon}(V, W) = 1 - \frac{LCSS^{m,n}_{\delta,\epsilon}}{\min(m, n)}. \tag{10}$$

We define the distance between two focal slices as required by Equation 9 as:

$$||v_i, w_j|| = \chi^2(v_i, w_j), \tag{11}$$

where χ^2 is the chi-square histogram similarity [19]. Finally, pollen classification is achieved as follows. First, we create multifocal appearance models for a set of known pollen grains. These models, called templates, are stored to be compared with unknown sequences for recognition. This is the *learning stage* of our method. Formally, given templates $\{\widehat{V}_1, \ldots, \widehat{V}_M\}$, the identification of an unknown test sequence V is done using the adapted LCSS method in Equation 9. We use a nearest-neighbor classifier based on the similarity measure produced by the LCSS algorithm.

VI. Experiments

To assess the classification performance of our algorithm, we obtained multifocal sequences for the *Alnus*, *Alternanthera*, *Darwiniothamnus*, *Scalesia*, and *Waltheria* pollen. The sequences were extracted at 25 consecutive focal lengths. For

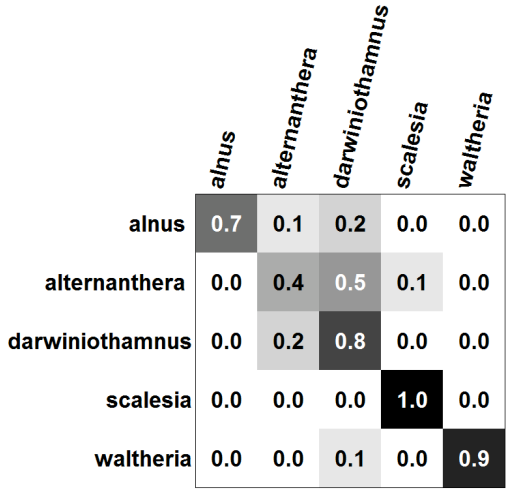

Fig. 7. Confusion matrix (76.0% recognition rate). The highest confusion occurred between the *Alternanthera* and *Darwiniothamnus* grains. Method was successful in identifying pollen types that are usually difficult to classify by visual inspection.

each pollen type, we extracted 20 multifocal sequences. We applied a "leave-one-out" classification scheme for evaluation. Here, we selected one image sequence from the dataset to be the test sequence. Templates for each class were then obtained from the remaining sequences. The templates were used to classify the test sequence, and the classification error was computed based on the number of correct recognition.

The confusion matrix generated by our classification results is presented in Figure 7. This matrix describes the average percentage of correct classification (i.e., matrix's main diagonal) for test images with respect to all other templates in the dataset. The off-diagonal values in the matrix show the average incorrect classification percentage. Overall, our algorithm accomplished 76.0% recognition performance. The highest confusion occurred between the *Alternanthera* and the *Darwiniothamnus* pollen. This confusion can be the result of similarity in outer boundary of the pollen. These results suggest that more descriptive information about the internal geometry and symmetry of surface features may be needed to improve classification of similar pollen types.

Acknowledgments. The authors acknowledge support from National Science Foundation (NSF) grant No. 1152306.

VII. Conclusion

We presented a pollen-classification algorithm based on matching sequences of multifocal images. Our appearance model captures scale and orientation information of local geometric features on the pollen's surface. These geometric information is represented by the responses of a bank of multi-scale oriented and isotropic filters. Rotation and scale invariance are achieved by a polar mapping of segmented images. The novelty of our approach lies in addressing the optical microscopy recognition problem from a sequence-alignment perspective, with the multi-focal texture modeling. The method uses optical microscopy and does not rely on specialized imaging procedures such as florescence and deconvolution. It is therefore applied to a wide variety of recognition tasks in day-to-day laboratory activities. Additionally, image sequences of any size can be used, with no requirements on focal-step uniformity during image acquisition. Finally, the appearance models can be easily extended to represent information such as color, geometry, and feature symmetry.

References

[1] M. B. Bush and M. B. Wengs, "Introducing a new (freeware) tool for Palynology," *Journal of Biogeography. (In Press)*, 2006.

[2] P. Carrion, E. Cernadas, J. F. Galvez, M. Damian, and P. de Sa-Otero, "Classification of honeybee pollen using a multiscale texture filtering scheme," *Machine Vision and Applications*, vol. 15, no. 4, pp. 186–193, 2004.

[3] P. Li, W. J. Treloar, J. R. Flenley, and L. Empson, "Towards automation of palynology 2: the use of texture measures and neural network analysis for automated identification of optical images of pollen grains," *J. of Quaternary Science*, vol. 19, no. 8, pp. 755–762, 2004.

[4] Y. Zhang, D. W. Fountain, R. M. Hodgson, J. R. Flenley, and S. Gunetileke, "Towards automation of palynology 3: pollen pattern recognition using gabor transforms and digital moments," *J. of Quaternary Science*, vol. 19, no. 8, pp. 763–768, 2004.

[5] O. Ronneberger, E. Schultz, and H. Burkhardt, "Automated pollen recognition using 3d volume images from fluorescence microscopy," *Aerobiologia*, vol. 18, no. 2, pp. 107–115, 2002.

[6] A. Boucher, P. Hidalgo, M. Thonnat, J. Belmonte, C. Galan, P. Bonton, and R. Tomczak, "Development of a semi-automatic system for pollen recognition," *Aerobiologia*, vol. 18, no. 3-4, pp. 195–201, 2002.

[7] W. J. Treloar, G. E. Taylor, and J. R. Flenley, "Towards automation of palynology 1: analysis of pollen shape and ornamentation using simple geometric measures, derived from scanning electron microscope images," *J. of Quaternary Science*, vol. 19, no. 8, pp. 745–754, 2004.

[8] C. M. Travieso, J. C. Briceno, J. R. Ticay-Rivas, and J. B. Alonso, "Pollen classification based on contour features," in *Intelligent Engineering Systems (INES), 2011 15th IEEE International Conference on*. IEEE, 2011, pp. 17–21.

[9] I. Fogel and D. Sagi, "Gabor filters as texture discriminator," *Biological Cybernetics*, vol. 61, no. 2, pp. 103–113, June 1989.

[10] T. Leung and J. Malik, "Recognising surfaces using three-dimensional textons," in *ICCV99*, 1999, pp. 1010–1017.

[11] ——, "Representing and recognizing the visual appearance of materials using three-dimensional textons," *International Journal of Computer Vision*, vol. 43, pp. 29–44, 2001.

[12] M. Vlachos, M. Hadjieleftheriou, D. Gunopulos, and E. Keogh, "Indexing multi-dimensional time-series with support for multiple distance measures," in *KDD*. New York, NY, USA: ACM, 2003, pp. 216–225.

[13] M. Rousson and R. Deriche, "A variational framework for active and adaptative segmentation of vector valued images," in *Workshop on Motion and Video Computing*. IEEE, 2002, p. 56.

[14] T. F. Chan and L. A. Vese, "Active contours without edges," *Image Processing, IEEE Transactions on*, vol. 10, no. 2, pp. 266–277, 2001.

[15] G. Gerig, O. Kubler, R. Kikinis, and F. A. Jolesz, "Nonlinear anisotropic filtering of mri data," *Medical Imaging, IEEE Transactions on*, vol. 11, no. 2, pp. 221–232, 1992.

[16] D. J. C. MacKay, *Information Theory, Inference, and Learning Algorithms*. Cambridge University Press, 2003.

[17] D. J. Berndt and J. Clifford, "Finding patterns in time series: a dynamic programming approach," *Advances in Knowledge Discovery and Data Mining*, pp. 229–248, 1996.

[18] J. Blackburn and E. Ribeiro, "Human motion recognition using isomap and dynamic time warping," in *Workshop on Human Motion*, 2007, pp. 285–298.

[19] M. Varma and A. Zisserman, "A statistical approach to texture classification from single images," *International Journal of Computer Vision*, vol. 62, pp. 61–81, 2004.