

Statistical Inference of Biological Structure and Point Spread Functions in 3D Microscopy

Joseph Schlecht[†] Kobus Barnard[†] Barry Pryor[‡]

[†]Computer Science Department
Univerisity of Arizona
Tucson, AZ 85721
{schlecht, kobus}@cs.arizona.edu

[‡]Plant Sciences Department
University of Arizona
Tucson, AZ 85721
bmpryor@u.arizona.edu

Abstract

We present a novel method for detecting and quantifying 3D structure in stacks of microscopic images captured at incremental focal lengths. We express the image data as stochastically generated by an underlying model for biological specimen and the effects of the imaging system. The method simultaneously fits a model for proposed structure and the imaging system's parameters, which include a model of the point spread function.

We demonstrate our approach by detecting spores in image stacks of Alternaria*, a microscopic genus of fungus. The spores are modeled as opaque ellipsoids and fit to the data using statistical inference. Since the number of spores in the data is not known, model selection is incorporated into the fitting process. Thus, we develop a reversible jump Markov chain Monte Carlo sampler to explore the parameter space.*

Our results show that simultaneous statistical inference of specimen and imaging models is useful for quantifying biological structures in 3D microscopic images. In addition, we show that inferring a model of the imaging system improves the overall fit of the specimen model to the data.

1. Introduction

In this paper we detail a new method for automatically detecting and quantifying 3D structure of biological specimen that counteracts the blurring effects of a microscopic imaging system. Quantifying the structure of cells and organisms is important for many biological experiments, but this process can be expensive and very time consuming when done manually. A method to automatically detect, quantify, and classify the 3D structure of specimen in microscopic images would enable high-throughput data analysis, improved experimental efficiency, and possibly lead to increased frequency of scientific discoveries.

The challenges in creating such an algorithm for analyzing microscopic data lie not only in the detection of structure, but in understanding the image formation process of the microscope. Depending on the type of microscope used, images of a specimen under view may contain a significant amount of blur from out-of-focus regions. In a standard compound microscope with high magnification, this is a result of a shallow depth of field. Thus, the optical system of a microscope can make accurate localization of detected structure in images more difficult.

To detect and quantify the structure of biological specimen in microscopic images, we propose a model that stochastically generates the observed data. A set of 3D geometrical objects model the structure of the specimen under study, and a theoretical impulse response of the microscope models the optical system. Using Bayesian statistical inference and Markov chain Monte Carlo sampling, we fit both of these models simultaneously to microscopic image data with mutual benefit; information learned from inferred specimen structure is used to learn model parameters of the imaging system and vice versa.

The impulse response, or point spread function, of the microscope's optical system blurs the observed image data. Learning a model of the point spread function (PSF) enables an understanding of the image formation process in the microscope. This, in turn, permits us to hypothesize unblurred images of the specimen and obtain a more accurate fit to its structure.

Since a microscope captures multiple image sets over time, we can learn its PSF in conjunction with fitting structure in many data sets at once. When sufficiently fit, we could utilize the PSF to detect structure in future data sets with less computation. Moreover, using a model to learn the PSF from image data facilitates inferring structure that has been imaged under a range of optical systems.

Fitting a geometrical model to microscopic image data results in quantified information about a specimen's structure, such as volume, count, and eccentricity of shape. This

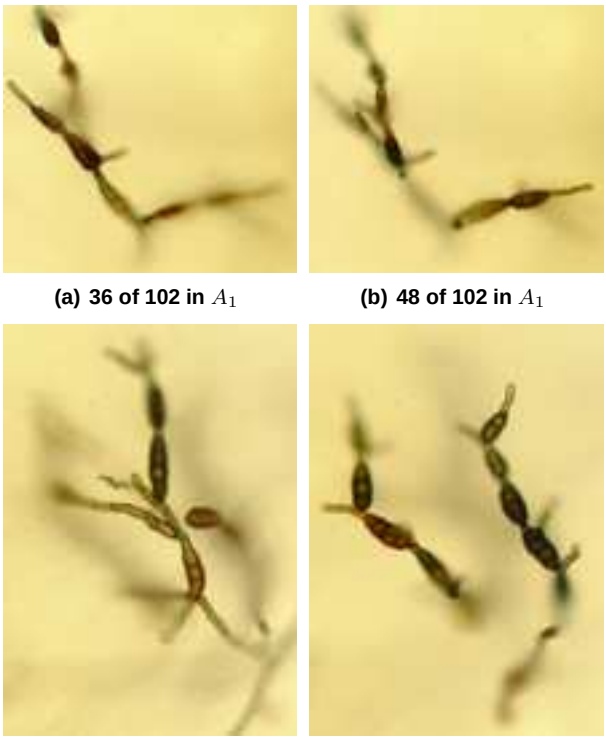

(a) 36 of 102 in A_1 (b) 48 of 102 in A_1

(c) 13 of 82 in A_2 (d) 53 of 82 in A_2

Figure 1. Images from *Alternaria* 3D data sets A_1 and A_2. In each image, the PSF of the brightfield transmitted-light microscope generated blur from nearby focal planes.

information enables scientific morphological analysis, automatic species or type classification, linkage to gene expression data, and 3D visualization. A model based visualization has significant advantages to standard surface detection methods, including noise and blur elimination, and reconstruction when data is missing.

1.1. 3D Microscopic Image Data

The data used in this research are 3D images of *Alternaria*, a genus of fungus, captured by a standard brightfield transmitted-light microscope. The images are 3D in the sense that the mycologist who captured them continuously imaged the specimen while increasing the focal depth of the microscope, a process commonly referred to as 3D microscopy. Figure 1 shows images from two of these sets, A_1 and A_2. Notice the significant blur in the images, a result of the optics in the transmitted-light microscope.

The general structure of *Alternaria* is tree-like with varying branching patterns. It comprises tubular filaments, known as hyphae, and ellipsoid-shaped reproductive spores that are darkly pigmented. Branching typically occurs as a bifurcation of the hyphae, but it may occur in the spores as well. One of the goals of this research is to create a high-throughput system for automatically quantifying reproductive spores in 3D images of *Alternaria*.

1.2. Related work

It was recently shown that a 2D model of ellipses could be inferred from images of cartilage cells formed under a confocal microscope [2]. This analysis enabled quantification of the number of cells in the image. However, because the model is in two dimensions and the structure of the cells exists in three, further quantitative analysis was not possible. To ameliorate this, a 3D model for fitting the cells was proposed, but the method of inference resulted in poor performance [1].

The optical system of a confocal microscope attempts to minimize the aberrations, flare, and blurring potential of its PSF, thus reducing blur in its images [6, 14]. Previous studies in statistical inference of structure modeled the PSF of a confocal microscope with a Gaussian function, but the parameterization was obtained by preliminary, manual analysis of the image data [2, 1]. Furthermore, because of the minimal blur in the data, their PSF model was less critical to a good fit of the structure than it would have been under a standard transmitted-light microscope, such as the one imaging our data.

Depending on the immersion medium and microscope, a PSF can be measured and estimated by imaging a tiny bead of material, such as oil or latex. The resulting measurements can subsequently be used to deconvolve images formed by the microscope [13]. However, performing the measurements can be a very time consuming and tedious process and the results are microscope dependent.

Efforts have been made to learn the structure of a PSF without direct measurement for the sole purpose of image restoration [5, 9]. Results of this work have been somewhat successful. However, the images were formed under a confocal microscope. It has not been shown that these methods can effectively deconvolve images from a brightfield transmitted-light microscope.

2. Structure and Imaging Models

Our generative model for the 3D microscopic image data comprises a model for *Alternaria* spores, the PSF of the imaging system, and the background light intensity of the brightfield microscope. What follows is a description of each component in the model.

2.1. Spore Model

Figure 1 shows that *Alternaria* spores are elliptical in shape and darkly pigmented, so we model them as opaque ellipsoids. Thus, the i^{th} spore in the model has parameters

$$s_i = (x, y, z, a, b, c, \varphi, \vartheta, \psi, \lambda)\,,$$

where x, y, z give the center of the spore in a 3D imaging window $\mathcal{W}$; $a, b, c \in \mathbb{R}^{3+}$ give the semi-axis lengths; φ, ϑ, ψ are Euler rotation angles and vary over $[0, \pi]$; $\lambda \in [0, 1]$ represents the average opacity of the spore in the image data.

Denote the space containing all parameterizations of the i^{th} spore as S_i, and let the space for n spores be $\Psi^{(n)} = S_1 \times \cdots \times S_n$. Then an ordered set of n spores is given by

$$\psi^{(n)} = (s_1, \cdots, s_n)\,.$$

2.2. PSF Model

The image formation process in a microscope is a convolution of the true unobserved 3D image with the point spread function, or impulse response, of the imaging system. The PSF is the 3D response $h(x, y, z)$ of a point source of light in the system. Using constraints from previous empirical observations [13], we introduce a model for the PSF of a transmitted light microscope.

Let $\tilde{h}(\cdot)$ be a model of the actual PSF in the imaging system. The x, y-plane in the space containing the model is defined to be parallel to the focal plane and the z-axis aligned with the optical axis of the microscope. The function is defined as a sequence of weighted 2D Gaussians, each parallel to the x, y-plane and centered on the z-axis. Thus, it is symmetric about the x, y-plane and around the z-axis.

Formally, we define $\tilde{h}(\cdot)$ as a mixed function

$$\tilde{h}(x, y, z) = \frac{\alpha^{|z|}}{\sqrt{2\pi\,(\beta\,|z| + \gamma)}}\, e^{-\frac{x^2+y^2}{2(\beta\,|z|+\gamma)}} \qquad (1)$$

with $x, y \in \mathbb{R}^2$ and $z \in \mathbb{Z}$. The parameter γ gives the base variance for the Gaussians, and β scales the distance from the x, y-plane. Thus, each Gaussian in $\tilde{h}(\cdot)$ has a variance that is linear with respect to its distance from the x, y-plane. α is the base in a geometric distribution used to weight the Gaussians.

Denote the space over all PSF models as Φ, and let a parameterization of the model be

$$\phi = (\alpha, \beta, \gamma).$$

An approximate geometric description of (1) is two cones placed apex-to-apex at the origin of the x, y-plane. The values within the cones are weighted 2D Gaussians parallel to the x, y-plane (figure 2).

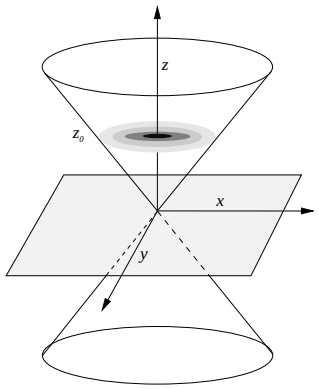

Figure 2. Diagram of the PSF model $\tilde{h}(x, y, z)$. The x, y-plane is the focal plane of the microscope, and the z-axis is aligned with the optical axis. The 2D Gaussians are stacked along the z-axis away from the focal plane with linearly increasing variance and geometrically decreasing weights. A Gaussian at distance z_0 is illustrated.

2.3. Imaging Background

Alternaria in the 3D image data occupy a relatively small region of the imaging window. Hence, many pixels in the data are saturated with the intensity of light used by the brightfield microscope. We denote the background intensity of the imaging system as υ and define it over the range $\Upsilon = [0, 1]$.

2.4. Image Model

Let $\Theta^{(n)} = \Psi^{(n)} \times \Phi \times \Upsilon$ be the parameter space over multi-spore, PSF, and background models, and let $\theta^{(n)} = (\psi^{(n)}, \phi, \upsilon)$ be an instance of that space. Then the solution space spanning all model configurations is

$$\Omega = \bigcup_{n>0} n \times \Theta^{(n)}.$$

For any $(n, \theta^{(n)}) \in \Omega$, we construct a model scene $I_\theta(i, j, k)$, which is a hypothesis of the unobserved 3D image data. Background pixels in the model scene have the highest saturation with value υ, and pixels belonging to the l^{th} spore with opacity λ_l have the value $\upsilon\,(1 - \lambda_l)$.

Given a model scene, pixels in the 3D image data $I(i, j, k)$ are modeled as independent and Gaussian distributed with means and variances defined by

$$\mu_{I_\theta}(i, j, k) = I_\theta \ast\ast\ast \hat{h} \qquad (2)$$

$$\sigma^2_{I_\theta}(i, j, k) = c_1\,|\mu_{I_\theta}(i, j, k) - \upsilon| + c_2\,\mu_{I_\theta}(i, j, k), \qquad (3)$$

where $***$ denotes 3D convolution and $\hat{h}$ is the quantized PSF model in (1).

The mean value for a pixel in $I(\cdot)$ is a weighted average of the model scene pixel intensities by the PSF model. The constants c_1 and c_2 scale the variance in a linear combination of spore and pixel intensities. A spore's pigment is not uniform across occupying pixels, and spores with greater opacity tend to have higher variability. The second term in (3) approximates pixel intensity variations due to Poisson noise in the imaging system. The scaling constants are set to small values and obtained by preliminary analysis. We found that the system is not very sensitive to them.

3. Model Inference

Given a stack of *Alternaria* image data $I(i,j,k)$ in the 3D window $\mathcal{W}$, the task is to find the model $(n, \theta^{(n)}) \in \Omega$ that best fits the data. To accomplish this, Bayesian statistical inference is used; we define a probability distribution over the solution space given the image data and find a maximum. Specifically, we define a *posterior*

$$p\left(n, \theta^{(n)} \mid I\right) = k_p \, L\left(I \mid n, \theta^{(n)}\right) \, \pi\Big(n, \theta^{(n)}\Big), \quad (4)$$

where k_p is a normalization constant, $L(\cdot \mid \cdot)$ is the likelihood of the image data, and $\pi(\cdot)$ is the model prior.

The independence assumption among pixels in the model of the image data results in a product of Gaussians for the likelihood function. Using the image model means (2) and variances (3), the likelihood is defined as

$$L\left(I \mid n, \theta^{(n)}\right) = \prod_{i,j,k} \frac{\sigma_{I_\theta}^{-1}}{\sqrt{2\pi}} \, e^{-\frac{1}{2}\left[\frac{I(i,j,k) - \mu_{I_\theta}(i,j,k)}{\sigma_{I_\theta}}\right]^2}. \quad (5)$$

3.1 Model priors

The prior over the model space Ω assumes independence between the spore and imaging models and is defined as

$$\pi\Big(n, \theta^{(n)}\Big) = \pi_\Psi\Big(n, \psi^{(n)}\Big) \; \pi_\Phi(\phi) \; \pi_\Upsilon(\upsilon)\,. \quad (6)$$

The priors for the imaging model parameters ϕ and υ are i.i.d. Gaussian. The spores in *Alternaria* generally have the same shape, opacity and count, but their position and orientation is quite varied. We integrate this information into the spore model prior as follows.

The position of a spore ranges uniformly over the window $\mathcal{W}$ with volume $V_\mathcal{W}$. The rotation angles are modeled as independent and uniformly distributed over $[0, \pi]$. Since the spore sizes tend to be roughly the same, with a major axis and two minor axes of similar length, we define independent Gaussians over them with means μ_a for the major axis and μ_{bc} for the two minor axes. We model the spore opacity as a truncated Gaussian over $(0, 1]$. Thus, the density function for a spore s_i is

$$f(s_i) = f_{x,y,z}(s_i) \, f_{a,b,c}(s_i) \, f_{\varphi,\vartheta,\psi}(s_i) \, f_\lambda(s_i), \quad (7)$$

where

$$f_{x,y,z}(s_i) = \frac{1}{V_\mathcal{W}},$$

$$f_{a,b,c}(s_i) = \frac{\sigma_a^{-1}\,\sigma_{bc}^{-2}}{(2\pi)^{\frac{3}{2}}} \, e^{-\left[\frac{(a_i-\mu_a)^2}{2\sigma_a^2} + \frac{(b_i-\mu_{bc})^2 + (c_i-\mu_{bc})^2}{2\sigma_{bc}^2}\right]},$$

$$f_{\varphi,\vartheta,\psi}(s_i) = \frac{1}{\pi^3},$$

$$f_\lambda(s_i) = \frac{\sigma_\lambda^{-1}}{\sqrt{2\pi}} \, e^{-\frac{1}{2}\left[\frac{(\lambda_i - \mu_\lambda)}{\sigma_\lambda}\right]^2}.$$

The existence of a spore in $\mathcal{W}$ follows a Poisson process, so n is Poisson distributed with intensity ν. For this work, the value of ν was set to 10. Finally, we restrict the interaction between spores so they do not intersect. The spore model prior is then

$$\pi_\Psi\Big(n, \psi^{(n)}\Big) = k_\pi^n \, \frac{\nu^n \, e^{-\nu}}{n!} \, \prod_{i=1}^{n} f(s_i) \, \chi(s_i \nvdash s_{j \neq i}), \quad (8)$$

where k_π^n is a normalization constant for the truncated Gaussians, $\vdash$ denotes geometric intersection, and $\chi(\cdot)$ is the characteristic function giving 1 for true and 0 otherwise.

4. Inference via Sampling

Inferring the most likely model given *Alternaria* image data is a challenging task; the posterior (4) is a complex distribution virtually impossible to evaluate analytically or numerically. Thus, we employ Markov chain Monte Carlo (MCMC) sampling to explore the model solution space in search of a maximum under the posterior [3, 11].

The sampler iteratively generates random, unbiased model samples from the solution space Ω. It consists of a set of moves, or Markov chain, that create new model samples by proposing changes to parameters in a previous sample. The sampler moves fall into two categories: changes to a spore, the PSF, or the background; and changes to the number of spores in the model. The latter are referred to as *diffusion* moves and the former *jump* moves [15].

At each iteration of the sampler, the m^{th} move is selected for execution with probability $r(m)$ and a new model

$(n, \tilde{\theta}^{(n)})$ is proposed. In this paper, a uniform distribution was used for $r(\cdot)$. Depending on how likely the new model is under the posterior and to have been proposed, it is accepted or rejected. This is the Metropolis-Hastings (MH) algorithm for MCMC [10, 8] and it is used for both diffusion and jump moves.

4.1. Diffusion moves

The diffusion moves for modifying a spore and proposing a new model are shift, resize, rotate, and opacity modification, as well as moves to update PSF and background parameters. The proposal distributions for diffusion moves are obtained by modifying the prior in (6). For parameters updated in a move, we replace their subdensity in the prior with a Gaussian that has means equal to corresponding parameters in the previously accepted model.

For example, the proposal distribution for randomly selecting the j^{th} spore in a model $(n, \theta^{(n)})$ and shifting its position is given by

$$q_{\mathsf{shift}}\left(\tilde{\theta}^{(n)} \mid \theta^{(n)}\right) = \frac{1}{n} \frac{\pi\left(n, \tilde{\theta}^{(n)}\right)}{f_{x,y,z}(\tilde{s}_j)} \frac{\sigma_{x,y,z}^{-3}}{(2\pi)^{\frac{3}{2}}} e^{-\frac{(\tilde{x}_j - x_j)^2 + (\tilde{y}_j - y_j)^2 + (\tilde{z}_j - z_j)^2}{2\,\sigma_{x,y,z}^2}}, \quad (9)$$

where $\sigma_{x,y,z}^2$ is a small variance. The proposal distributions for other diffusion moves are similarly constructed.

Under the MH algorithm for the m^{th} diffusion move, the acceptance probability for a proposed model is

$$\alpha\left(n, \tilde{\theta}^{(n)}\right) = \min\left\{1, \frac{p(n, \tilde{\theta}^{(n)} \mid I)\, q_{\mathsf{m}}(\theta^{(n)} \mid \tilde{\theta}^{(n)})}{p(n, \theta^{(n)} \mid I)\, q_{\mathsf{m}}(\tilde{\theta}^{(n)} \mid \theta^{(n)})}\right\}. \quad (10)$$

The definition is derived to maintain a detailed balance condition in the Markov chain, which is a sufficient condition for convergence to the posterior [11].

By expansion, most of the terms in (10), including the normalization constants, cancel. Thus, the acceptance probability for a shift move of the j^{th} spore becomes

$$\alpha\left(n, \tilde{\theta}^{(n)}\right) = \min\left\{1, \frac{L(I \mid n, \tilde{\theta}^{(n)})\, f_{x,y,z}(\tilde{s}_j)}{L(I \mid n, \theta^{(n)})\, f_{x,y,z}(s_j)} \prod_{i \neq j} \chi(s_i \not\vdash \tilde{s}_j)\right\}. \quad (11)$$

As with the proposals, the acceptance probabilities for other diffusion moves are similar; hence, their definitions are omitted.

4.2. Jump moves

The jump moves in the sampler are birth and death of a spore. In both moves, the dimensionality of the model is modified as a spore is added to or removed from the model. For a birth move, the proposal distribution for a new spore $\tilde{s}$ is defined as the normalized spore density (7) in the model prior

$$q_{\mathsf{birth}}(\tilde{s}) = k_\pi\, f(\tilde{s}). \quad (12)$$

During a death move, a spore is randomly selected for deletion, so a proposal distribution is not needed.

In order to use the MH algorithm for jump moves, we redefine the acceptance probability. Following the guidelines for reversible-jump MCMC [7], the acceptance probability for a birth move becomes

$$\alpha\left(n+1, \tilde{\theta}^{(n+1)}\right) = \min\left\{1, \frac{p(n+1, \tilde{\theta}^{(n+1)} \mid I)}{p(n, \theta^{(n)} \mid I)} \frac{r(\mathsf{death})}{r(\mathsf{birth})\, q_{\mathsf{birth}}(\tilde{s})} \left|\frac{\partial(\tilde{\theta}^{(n+1)})}{\partial(\theta^{(n)}, \tilde{s})}\right|\right\}. \quad (13)$$

Since the change in dimensionality is a one-to-one mapping from $(\tilde{s}, \theta^{(n)}) \rightarrow \tilde{\theta}^{(n+1)}$ and a uniform distribution is used for $r(\cdot)$, the Jacobian is 1 and the move probabilities cancel. Thus, the acceptance probability for birth reduces to

$$\alpha\left(n+1, \tilde{\theta}^{(n+1)}\right) = \min\left\{1, \frac{L(I \mid n+1, \tilde{\theta}^{(n+1)})}{L(I \mid n, \theta^{(n)})} \frac{\nu}{n+1} \prod_{i=1}^{n} \chi(s_i \not\vdash \tilde{s})\right\}. \quad (14)$$

Because birth and death moves are dual, the acceptance probability for a death move is the inverse of the second argument to the minimum function in (14).

As with the diffusion moves, the jump move acceptance probabilities maintain the detailed balance condition [7]. Thus, the posterior will be the stationary distribution of the trans-dimensional Markov chain followed by the sampler.

4.3. Data-driven Birth Move

Since the prior is uniform over spore position and orientation, birth proposals based on it have a high rejection rate. This causes an increase in the number of iterations required for the sampler to converge. To solve this problem, we improve the birth proposals by doing preliminary data analysis to construct a more informative proposal distribution, so called data-driven MCMC [15].

The replacement proposal distribution for birth moves is generated from the *Alternaria* image data. We apply surface detection and a Hough transform for ellipsoids to obtain rough estimates of spores in the data. The estimates are subsequently collected into a spore likelihood table, which is normalized and used as the new proposal distribution.

The surface detector is an extension of the standard Canny edge detection algorithm [4] to three-dimensions. The gradient of a 3D Gaussian is convolved with the image data to produce 3D pixel gradient vectors. We apply non-maximal suppression and hysteresis to the resulting vectors to detect surface points.

A Hough transform is used to find ellipsoids from the detected surface. We reduce the number of parameters defining a spore from 9 to 7 by assuming equal minor axes and use a very coarse parameter quantization for the spores. Although this results in coarse estimates of spores in the data, it reduces the size and complexity of the Hough transform. Furthermore, coarse estimates are tolerable because diffusion moves in the sampler will perfect the fit of proposed spores.

We construct the Hough transform by iterating over the quantized spores at each surface point and incrementing a counter in the Hough array $\mathcal{H}$. The counts in the array are normalized to produce a likelihood over spores. Thus, we redefine the birth proposal distribution as

$$q_{\text{birth}}(s) = \mathcal{H}[i_s], \tag{15}$$

where i_s is the index into the normalized Hough array that has minimal Euclidean parameter distance to the spore s.

5. Results

We evaluated the effectiveness of the model sampler on *Alternaria* image sets A_1 and A_2. In addition, we tested the sampler on synthetic spore data to obtain a comparative measure for its performance on ideal data.

5.1. Synthetic Data Evaluation

The synthetic data were randomly generated from our model of the imaging system and spores. We created ten data sets and optically sectioned them into 80 images of size 300×300 pixels. Each set contained 10 randomly generated spores.

The MCMC sampler was run for 3000 iterations at a resolution of 20%, followed by 1000 iterations at 50%. It was run on lower resolutions of the images in order to decrease run-time. This process was repeated four times on all the data sets using a different random seed each time. The average number of spores correctly fit to a data set was 8. The results for each set are shown in figure 3.

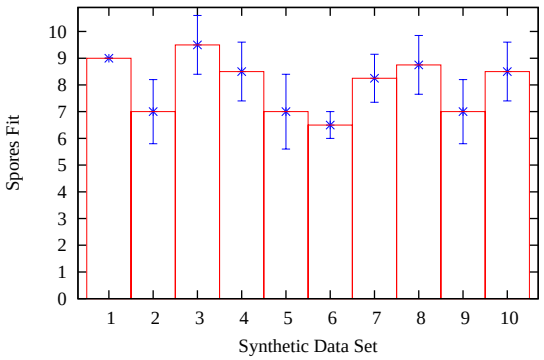

Figure 3. Mean number of spores correctly fit to the synthetic data sets, standard deviation bars are shown. Each synthetic set contained 10 spores and was run with four different random seeds.

Difficulties in detection arose when two or more spores in the data were nearly parallel in their major axes and very close together, in which case one model spore was sometimes fit to both. Occasionally, multiple model spores were fit to one in the data. In a few cases, no proposal was made for a spore in the data, most likely because it had low likelihood in the birth proposal distribution and would require more iterations to be proposed.

5.2. *Alternaria* Evaluation

Two sets of *Alternaria* image stacks were evaluated: A_1 composed of 102 images of size 800×800 pixels and A_2 with 80 images of size 700×700. Images of these sets can be seen in figure 1. The number of spores in the data sets were manually counted and found to be 17 and 21 for A_1 and A_2 respectively.

We ran the sampler on both *Alternaria* data sets for 500 iterations at a resolution of 20%. As with the synthetic evaluation, four instances of the sampler were run on the data sets, each time with a different random seed. Figure 4 shows a 3D rendering of a fit model for each data set next to the detected *Alternaria* surface used for generating data-driven birth proposals.

In A_1 the average number of spores detected was 6 and approximately 9 for A_2. While we did not achieve 80% accuracy, as in the synthetic case, the results are still noteworthy considering the amount of non-spore structure and substantial blur in the data.

The average inferred background intensity for A_1 and A_2 was 0.78 and 0.75 with a negligible standard deviation. Table 1 gives the inferred PSF model parameters for the data sets. The standard deviation is relatively high for PSF parameters β and γ. This is most likely due to the sampler

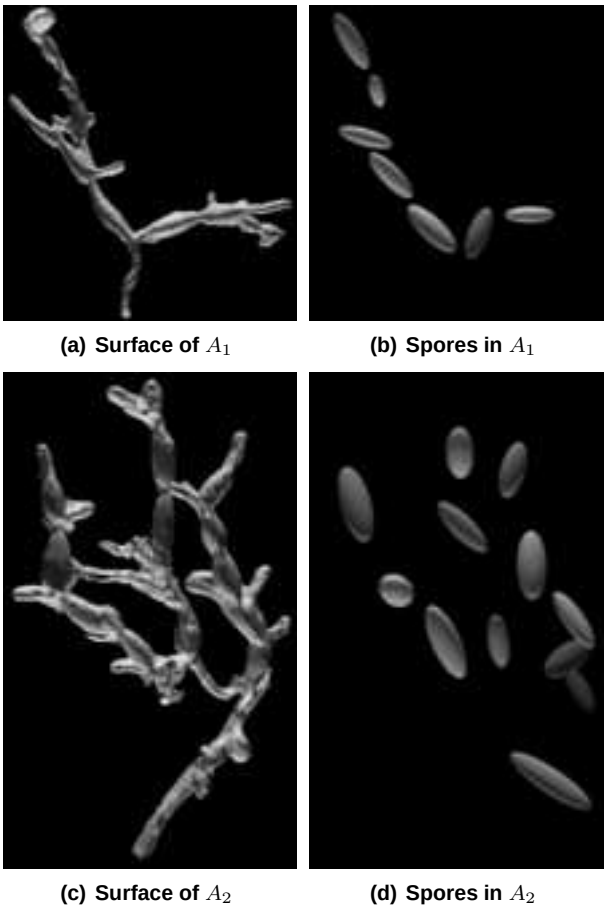

(a) Surface of A_1 **(b) Spores in A_1**

(c) Surface of A_2 **(d) Spores in A_2**

Figure 4. Reconstructed surface of *Alternaria* in the image stacks and 3D renderings of corresponding inferred spore models. The surface detection algorithm for data-driven birth proposals generated the views in (a) and (c). Perceived structure in these images is known only to the viewer. Sub-figures (b) and (d) represent detected spore structure.

adding variance to the PSF in order to accommodate the large quantity of non-spore structure in the images.

We tested the effect of using our model of the PSF for fitting spores versus using a 3D Gaussian and a delta function. In the case of the Gaussian, a σ of 0.6 was chosen for all dimensions by empirical analysis of the *Alternaria* image data. This is consistent with work done previously to approximate the PSF [2, 1, 13].

Table 2 lists the average number of spores detected using each of the PSFs during model inference. Figure 5 shows images in A_1 compared to inferred model scene images that are convolved with each of the PSF types. These results combined show that fitting our model of the PSF to the image data most closely resembles the imaging effects of the microscope and enables a more accurate estimate of structure in the images.

	α		β		γ	
	mean	stdev	mean	stdev	mean	stdev
A_1	0.93	0.001	0.82	0.44	1.31	0.30
A_2	0.93	0.030	1.06	0.14	1.35	0.33

Table 1. Mean PSF model parameters inferred from the *Alternaria* data from four random starting states. As expected, the fit parameters are similar for the data sets, which were imaged under the same microscope.

	Model PSF		Gaussian PSF		Delta PSF	
	mean	stdev	mean	stdev	mean	stdev
A_1	6.0	0.9	4.75	0.4	2.5	0.5
A_2	8.75	0.8	7.25	2.9	6.25	2.2

Table 2. Mean number of spores correctly fit to *Alternaria* data using a delta function, Gaussian, and our model as the PSF. Four random starting states were used for each data set. It is clear that fitting our model of the PSF improves spore detection.

6. Future Work

We plan to extend this work to model the overall structure of *Alternaria* with a grammar of its growth. Our idea is to fit substructures in the data, such as spores and hyphae, under the constraint that their combined structure is an instance of the grammar. One type of grammar that may be useful for this task is a stochastic L-system [12], which is commonly used for plant growth models to generate realistic instances of plants in computer graphics.

Currently, the sampler fits a PSF model and the structure in only one data set at a time. Although the learned PSF can be used when fitting other data sets sequentially, the sampler will be extended to fit a PSF and multiple data sets simultaneously. This will reduce the possibility of over-fitting the PSF.

References

[1] F. Al-Awadhi. *Statistical image analysis and confocal microscopy*. PhD thesis, Department of Mathematical Sciences, University of Bath, Bath, UK, 2001.

[2] F. Al-Awadhi, C. Jennison, and M. Hurn. Statistical image analysis for a confocal microscopy two-dimensional section of cartilage growth. *Journal of the Royal Statistical Society: Applied Statistics*, 53(1):31–49, 2004.

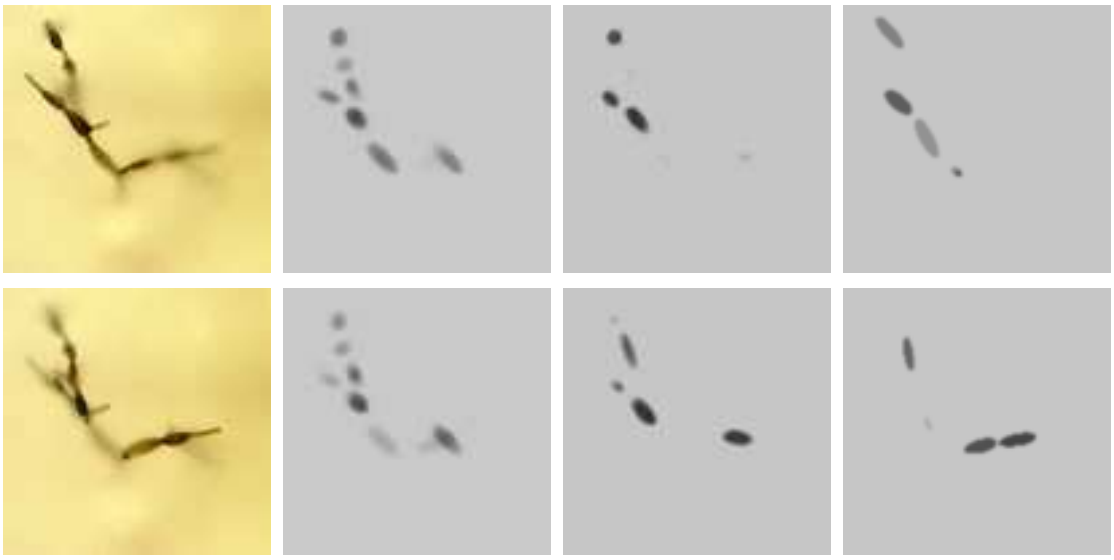

(a) Data (b) Model PSF (c) Gaussian PSF (d) Delta PSF

Figure 5. Illustration of the effects for three different PSFs used to detect spores. From top to bottom, column (a) contains images 36 and 48 out of 102 from *Alternaria* data set A_1. The other columns are corresponding images from inferred model scenes convolved with the learned PSF model (b), a 3D Gaussian (c), and a delta function (d). Notice that the images in (b) most closely resemble the *Alternaria* data in (a), indicating that our PSF model is substantively closer to the true PSF than a Gaussian and a delta function.

[3] C. Andrieu, N. de Freitas, A. Doucet, and M. I. Jordan. An introduction to MCMC for machine learning. *Machine Learning*, 50(1):5–43, 2003.

[4] J. Canny. A computational approach to edge detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 8:679–698, 1986.

[5] J. Conchello. Superresolution and convergence properties of the expectation-maximization algorithm for maximum-likelihood deconvolution of incoherent images. *Journal of Optical Society of America A*, 15(10):2609–2619, 1998.

[6] P. M. Conn, editor. *Confocal microscopy*, volume 307 of *Methods in enzymology*. Academic Press, San Diego, CA, 1999.

[7] P. J. Green. Reversible jump Markov chain Monte Carlo computation and Bayesian model determination. *Biometrika*, 82(4):711–732, 1995.

[8] W. K. Hastings. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika*, 57:97–109, 1970.

[9] T. J. Holmes. Blind deconvolution of quantum-limited incoherent imagery: maximum-likelihood approach. *Journal of Optical Society of America A*, 9:1052–1061, 1992.

[10] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller. Equations of state calculations by fast computing machines. *Journal of Chemical Physics*, 21:1087–1092, 1953.

[11] R. M. Neal. Probabilistic inference using Markov chain Monte Carlo methods. Technical Report CRG-TR-93-1, University of Toronto, 1993.

[12] P. Prusinkiewicz, A. Lindenmayer, and J. Hanan, editors. *The algorithmic beauty of plants*. Springer-Verlag, 1990.

[13] P. J. Shaw and D. J. Rawlins. The point-spread function of a confocal microscope: its measurement and use in deconvolution of 3-d data. *Journal of Microscopy*, 163(2):151–165, 1991.

[14] R. H. Webb. Theoretical basis of confocal microscopy. In P. M. Conn, editor, *Confocal microscopy*, volume 307 of *Methods in enzymology*, pages 3–20. Academic Press, San Diego, CA, 1999.

[15] S.-C. Zhu, R. Zhang, and Z. Tu. Integrating top-down/bottom-up for object recognition by data driven Markov chain Monte Carlo. In *IEEE Computer Vision and Pattern Recognition*, 2000.