

LIKELIHOOD INFERENCE FOR PARTICLE LOCATION IN FLUORESCENCE MICROSCOPY*

BY JOHN HUGHES, JOHN FRICKS, AND WILLIAM HANCOCK

The Pennsylvania State University

We introduce a procedure to automatically count and locate the fluorescent particles in a microscopy image. Our procedure employs an approximate likelihood estimator derived from a Poisson random field model for photon emission. Estimates of standard errors are generated for each image along with the parameter estimates, and the number of particles in the image is determined using an information criterion and likelihood ratio tests. Realistic simulations show that our procedure is robust and that it leads to accurate estimates, both of parameters and of standard errors. This approach improves on previous *ad hoc* least squares procedures by giving a more explicit stochastic model for certain fluorescence images and by employing a consistent framework for analysis.

1. Introduction. The accurate and precise tracking of microscopic fluorescent particles attached to biological specimens (e.g., organelles, membrane proteins, molecular motors) can give insights into the nanoscale function and dynamics of those specimens. This tracking is accomplished by analyzing digital images produced by a CCD (charge-coupled device) camera attached to a microscope used to observe the specimens repeatedly. In this paper, we introduce an improved technique for analyzing such images over time. Our method, which applies maximum likelihood principles, improves the fit to the data, derives accurate standard errors from the data with minimal computation, and uses model-selection criteria to “count” the fluorophores in an image. The ability to automate the process and quickly derive standard errors should allow for the analysis of thousands of images obtained from a typical experiment and aid in methods to track individual fluorophores across sequential images.

In fluorescence microscopy, a specimen of interest is tagged with a fluorescent molecule or particle. The fluorescence microscope then irradiates the specimen with light at the excitation wavelength of the fluorophore, and when the excited electrons revert to the ground state they emit photons at

*This work was supported by the NSF/NIH joint initiative in mathematical biology (DMS 0714939).

Keywords and phrases: maximum likelihood methods, Poisson random field, fluorescence microscopy, particle tracking, organelle, molecular motor, nanotechnology

the emission wavelength. A filter separates the emitted light from the excitation light so that only the light from the fluorescent material can pass to the microscope’s eyepiece and camera system (Rost, 1992).

In general, the Rayleigh criterion implies that the maximum resolution for a light microscope should be roughly 250 nm (half of the wavelength of visible light); however, Selvin and his collaborators found that by fitting the center point of the point spread function one can locate a particle of interest. This technique is known as FIONA (Fluorescence Imaging with One-Nanometer Accuracy), which was introduced in Yildiz, Forkey, McKinney, Ha, Goldman, and Selvin (2003). The key element of FIONA is to focus attention on single fluorophores used as markers in biological specimens (Kural, Balci, and Selvin, 2005). By analyzing sequences of images, molecular motors (e.g., myosin VI and kinesin) and other specimens can be tracked through time giving researchers insight into their dynamics and biological function. For instance, Yildiz et al. used FIONA to find compelling support for the hypothesis that myosin V walks hand over hand and evidence to eliminate other hypotheses (Yildiz et al., 2003; Kural et al., 2005).

A number of analysis techniques have been proposed for FIONA images. In 2001, Cheezum, Walker, and Guilford compared four methods—cross-correlation, sum-absolute difference, centroid, and Gaussian fit—and ultimately recommended the Gaussian-fit method for single-fluorophore tracking. In the Gaussian-fit approach, the method of ordinary least squares (OLS) is used to fit a sum of symmetric bivariate Gaussian functions to the image. Least squares fitting is relatively efficient, and software to do it is widely available. Thompson, Larson, and Webb subsequently proposed a “Gaussian mask” algorithm that is easier to implement than the Gaussian-fit method, is computationally less intensive, and performs nearly as well in simulations (2002). The Gaussian-mask algorithm is essentially a centroid calculation that weights each pixel with the number of photons in the pixel and with a bivariate Gaussian function integrated over the pixel. In both cases, simulation studies using typical experimental values showed that sub-pixel or even nanometer resolution was possible.

The above mentioned Gaussian-fit and Gaussian-mask methods, while appealing, share two shortcomings. Since one or more beads may move out of frame for a particular image, the number of beads from one image to the next is not known *a priori* and must be determined for each image. Previous authors have attempted to solve this problem by means of a grid search, the first step of which is to scan the image for all pixels greater than some arbitrary threshold value. Each of these extreme pixels is taken to be a bead location, and some region surrounding each extreme pixel is

extracted from the image and processed by OLS or the Gaussian mask. Thompson et al. suggest a threshold that is eight standard deviations above the mean pixel value, but no explicit evidence is given in support of this choice (Thompson et al., 2002). The correct threshold level for a set of images could possibly be approximated using simulations, but this would be a complex and computationally intensive task that would be necessary for each set of images, since the level of background noise may vary significantly from one experiment to the next.

The second drawback is that estimates of precision are derived from simulation studies alone. If the probability model for the problem is misspecified, then error estimates based on simulations from that model will be inaccurate even if reasonable parameter values were used in the studies. And these values will vary from image to image due to changing experimental conditions, e.g., elevated background noise or slight changes in focus. A possible solution is to perform a Monte Carlo simulation study using parameter values derived from the current experiment. But, given that the fitting procedures themselves are time consuming, these approaches to standard error calculation may prove infeasible. It takes our algorithm several minutes—on a dual 2.8GHz Quad-Core Intel Xeon Mac Pro—to process an image with fifteen particles, which implies that bootstrapping standard errors for such an image would require hours of computation. Moreover, a full analysis of an experiment requires processing many hundreds of images.

In what follows, we present a new approach to counting and locating fluorophores. Our approach eliminates the need for a grid search and estimates standard errors from the data, without additional simulation, via standard likelihood tools. In Section 2, we present an explicit probability model for a FIONA image along with a maximum likelihood estimation procedure suitable for this model. In Section 3, we discuss the properties of the approximate likelihood estimator presented in Section 2. In Section 4, we discuss stepwise model selection, which allows our procedure to automatically determine the number of beads in an image in a consistent manner. In Section 5, we describe the results of realistic simulation studies that support the approximations presented in Section 2 and demonstrate the robustness of our procedure. Finally, in Section 6, we carry out a complete analysis of an experimentally collected FIONA image, introducing relevant diagnostic criteria for our fit.

2. Model for Data from a Single FIONA Image. We first develop a model for the photon emission from particles distributed over a microscope slide. Photon emission from a constant source generally follows a Poisson

distribution; this fact naturally leads to a Poisson random field model for the emission from a slide. We then express the effect of pixelation on the field, representing the emission as viewed from the digital camera, and employ a normal approximation to the Poisson distribution. We arrive at our final approximate model by accounting for additional error introduced by the camera and its associated equipment.

We begin with the standard model for the photon-emission pattern (as distorted by the point-spread function of a microscope objective) of a collection of fluorophores distributed at random over some region of \mathbb{R}^2 (Cheezum et al., 2001; Thompson et al., 2002).

Let N , a Poisson random field on a rectangular subset T of \mathbb{R}^2 , represent the emission pattern of the sample. The intensity function can be defined for any Borel set $R \subset T$ as

$$E\{N(R)\} = \iint_R \left\{ B + \sum_{j=0}^{J-1} A_j \cdot \exp \left(-\frac{(x-x_j)^2 + (y-y_j)^2}{S^2} \right) \right\} dx dy.$$

Thus, $N(R)$ is a Poisson random variable with mean equal to a sum of J Gaussian functions, one for each bead, with Gaussian function j symmetric about (x_j, y_j) (which is contained in T). In addition, there is a constant background intensity of magnitude B representing background fluorescence. Although the intensity function for a bead is more often modeled in the physics literature by an Airy function, a Gaussian function approximates the Airy function quite well, and so we take the Gaussian centered at (x_j, y_j) to represent the (distorted) emission of bead j (Saxton, 1997; Thompson et al., 2002).

The photons emitted by the sample are collected by a camera, the pixels of which can be represented by partitioning T into a uniform grid, where each pixel in the grid is a nm on a side. Then, for a given pixel Z_i with center (x_i, y_i) ,

$$E(Z_i) = \int_{y_i-a/2}^{y_i+a/2} \int_{x_i-a/2}^{x_i+a/2} \left\{ B + \sum_{j=0}^{J-1} A_j \cdot \exp \left(-\frac{(x-x_j)^2 + (y-y_j)^2}{S^2} \right) \right\} dx dy,$$

which is approximately equal to

$$\left\{ B + \sum_j A_j \cdot \exp \left(-\frac{(x_i-x_j)^2 + (y_i-y_j)^2}{S^2} \right) \right\} a^2.$$

Since a^2 is a constant, we allow the A_j and B to absorb it, arriving at

$$E(Z_i) \approx B + \sum_j A_j \cdot \exp \left(-\frac{(x_i-x_j)^2 + (y_i-y_j)^2}{S^2} \right) = f_i.$$

Moreover, B is generally large enough to justify using a normal approximation to the Poisson distribution. More precisely,

$$Z_i \sim N(f_i, f_i),$$

where we note that f_i depends on the parameters of interest. In this model, it is obviously important that B and the A_j 's be constrained so that f_i is non-negative.

The discretized Poisson random field described above is taken as the underlying model for the photon emission; however, additional error, which we will call instrumentation error, arises from various sources such as signal quantization and dark current, an electric current that flows through a CCD even when no light is entering the device (Bobroff, 1986; Thompson et al., 2002). If we model the instrumentation error as a $N(0, \theta)$ random variable independent of the intensity, then we have as a final approximate model for the data

$$Z_i \sim N(f_i, f_i + \theta).$$

Consequently, the approximate likelihood of a given image with n pixels is

$$L_n(\boldsymbol{\beta}) = \prod_{i=1}^n \frac{1}{\sqrt{2\pi(f_i + \theta)}} \exp\left(-\frac{(z_i - f_i)^2}{2(f_i + \theta)}\right),$$

where $\boldsymbol{\beta} = (x_0, y_0, A_0, \dots, x_{J-1}, y_{J-1}, A_{J-1}, S, B, \theta)^T$, the parameters of interest. This implies that the log-likelihood, without unnecessary constants, is

$$\ell_n(\boldsymbol{\beta}) = -\sum_i \ln(f_i + \theta) - \sum_i \frac{(z_i - f_i)^2}{f_i + \theta},$$

which we maximize with respect to $\boldsymbol{\beta}$ to obtain $\hat{\boldsymbol{\beta}}_{MLE}$.

3. Estimating Standard Errors. From the theory of maximum likelihood estimators we know that, provided certain regularity conditions are met, a properly scaled MLE converges asymptotically to a normally distributed random variable. In our case,

$$[\mathcal{I}_n(\boldsymbol{\beta})]^{1/2}(\hat{\boldsymbol{\beta}}_n - \boldsymbol{\beta}) \xrightarrow{D} N_p(\mathbf{0}, \mathbf{I}_p),$$

where $\mathcal{I}_n(\boldsymbol{\beta})$ is the Fisher information about $\boldsymbol{\beta}$ contained in a sample of size n , p is the dimension of $\boldsymbol{\beta}$, and \mathbf{I} is the $p \times p$ identity matrix. This implies that

$$\hat{\boldsymbol{\beta}}_n \sim N_p(\boldsymbol{\beta}, [\mathcal{I}_n(\boldsymbol{\beta})]^{-1})$$

for large n , and so the diagonal elements of $[\mathcal{I}_n(\boldsymbol{\beta})]^{-1}$ are approximate sampling variances for the estimators $\hat{\boldsymbol{\beta}}_n$. However, we do not know $[\mathcal{I}_n(\boldsymbol{\beta})]^{-1}$ because the true $\boldsymbol{\beta}$ is unknown and analytical calculation of the information is prohibitively complicated.

Consequently, we use the standard substitution $[\mathcal{I}_n(\boldsymbol{\beta})]^{-1}$ using the observed information, $\mathcal{J}_n(\hat{\boldsymbol{\beta}}_n) = \left[-\frac{\partial^2}{\partial \beta_i \partial \beta_j} \ell_n(\hat{\boldsymbol{\beta}}_n) \right]$, i.e., $[\mathcal{I}_n(\boldsymbol{\beta})]^{-1} \approx [\mathcal{J}_n(\hat{\boldsymbol{\beta}}_n)]^{-1}$. Estimating $[\mathcal{I}_n(\boldsymbol{\beta})]^{-1}$ by inverting $\mathcal{J}_n(\hat{\boldsymbol{\beta}}_n)$ has the advantage of avoiding closed form derivatives which are unwieldy in this case.

The preceding implies that $\hat{x}_j \sim N[x_j, \widehat{Var}(\hat{x}_j)]$ and $\hat{y}_j \sim N[y_j, \widehat{Var}(\hat{y}_j)]$, where $\widehat{Var}(\hat{x}_j)$ and $\widehat{Var}(\hat{y}_j)$ represent the diagonal elements of $[\mathcal{J}_n(\hat{\boldsymbol{\beta}}_n)]^{-1}$ corresponding to \hat{x}_j and \hat{y}_j , respectively. Moreover, the joint distribution of \hat{x}_j and \hat{y}_j is approximately bivariate normal. More precisely,

$$(\hat{x}_j, \hat{y}_j)^T \sim N_2[(x_j, y_j)^T, \boldsymbol{\Sigma}],$$

where

$$\boldsymbol{\Sigma} = \begin{bmatrix} \widehat{Var}(\hat{x}_j) & 0 \\ 0 & \widehat{Var}(\hat{y}_j) \end{bmatrix}.$$

Since the contours, i.e., the equidensity curves, of the bivariate normal distribution are ellipses, an approximate 95% confidence region for the location of bead j also takes the form of an ellipse:

$$\frac{(x - \hat{x}_j)^2}{\widehat{Var}(\hat{x}_j)} + \frac{(y - \hat{y}_j)^2}{\widehat{Var}(\hat{y}_j)} = 3.8416 = z_{0.975}^2 = [\Phi^{-1}(0.975)]^2,$$

where Φ^{-1} denotes the standard normal quantile function. For an image with multiple beads, the typical case, we may want a collection of random ellipses the union of which will enclose all J beads with probability 0.95. We can accomplish this by using the Bonferroni correction, which assigns to each bead an error rate of $\frac{0.05}{J}$, thereby making the image wide error rate 0.05. The resulting collection of simultaneous confidence ellipses is given by

$$\frac{(x - \hat{x}_j)^2}{\widehat{Var}(\hat{x}_j)} + \frac{(y - \hat{y}_j)^2}{\widehat{Var}(\hat{y}_j)} = z_{1-0.025/J}^2 = \left[\Phi^{-1} \left(1 - \frac{0.025}{J} \right) \right]^2.$$

We evaluated our standard-error estimation and the convergence of our estimator by way of a simulation study. Ten thousand 100×100 single-bead images were simulated, each image having its lone bead located at (7823, 3353) where the coordinates are given in nanometers from the lower left corner. (For an idea about the nature of the data, see Figure 1 which has

TABLE 1
Estimation of Standard Errors

Parameter	Value	Estimate	\widehat{SE}	True SE
x_0	7823	7823.4	0.422	0.420
y_0	3353	3352.9	0.421	0.424
A_0	15000	14900.2	61.2	77.4
S	200	200.2	0.341	0.393
B	200	199.9	0.171	0.173
θ	100	93.3	4.16	4.26

four beads.) Table 1 shows $\hat{\beta}_{MLE}$ for a single image along with standard-error estimates for that image and the true standard errors gleaned from all 10,000 images. Our estimated standard errors are in close agreement with the true standard errors.

Figure 2 shows estimated densities for the sampling distributions of \hat{x}_0 , \hat{y}_0 , \hat{A}_0 , \hat{S} , \hat{B} , and $\hat{\theta}$, respectively. Appropriate normal densities (dashed) are shown superimposed. Each normal density is centered at the true value for its parameter. It is clear that the sampling distributions converge to normality, but the estimators of S and A_0 are slightly biased in opposing directions; intuitively, estimation of S works in opposition to that of A_0 . This is because the fitting procedure is attempting to simultaneously conform \hat{S} to the base of the Gaussian peak and \hat{A}_0 to the peak's height, and an adjustment of either estimate nudges the other in the opposite direction.

4. Model Selection: How Many Beads are in the Image. Previous authors have suggested that the number of beads in an image be determined by applying a grid search algorithm prior to fitting (Cheezum et al., 2001; Thompson et al., 2002). As we mentioned in Section 1, any pixel with an intensity above some threshold is identified with a bead, and then some region in the vicinity of the pixel is extracted from the image and fitted. This thresholding approach may be adequate for producing initial estimates of bead locations, but thresholding prevents full automation because the threshold must be chosen by the investigator. And even a seemingly well-chosen threshold may be too large to distinguish dim beads from background noise.

Our procedure eliminates these problems by fitting first and selecting the number of beads based on those fits. In our scheme, an approximate information criterion derived from an OLS fit is used to approximate the size of the model, and then, because the candidate models are nested, likelihood ratio tests are used to select the final model. As we will show in Section 5, this approach is able to identify all of the beads, even very dim ones, in a fully automatic way.





FIG 1. An image with four beads, only three of which are visible to the naked eye. Our procedure locates all four beads.

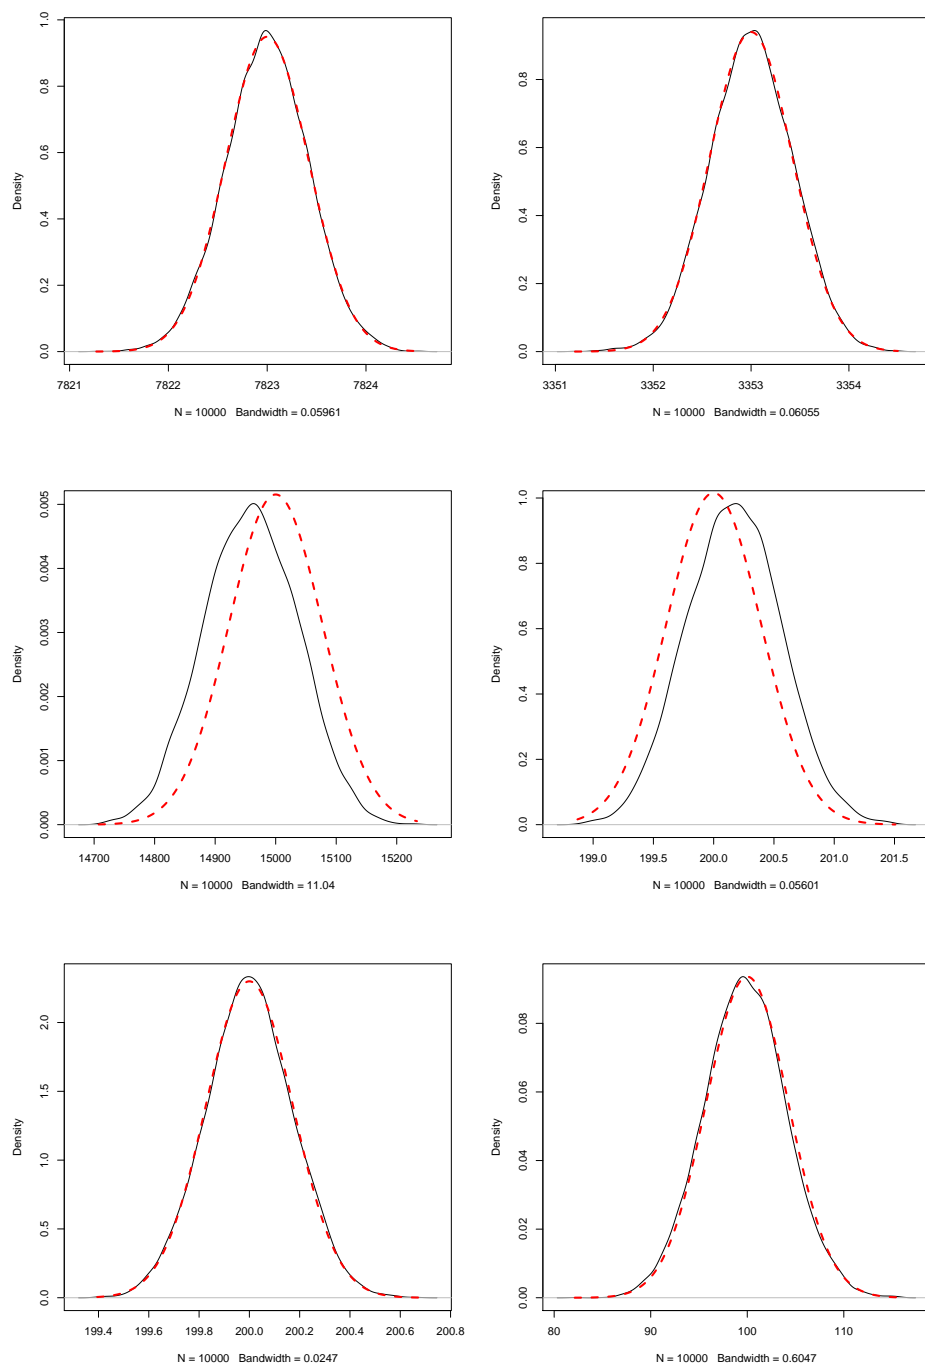


FIG 2. These plots show density estimates for the sampling distributions of \hat{x}_0 , \hat{y}_0 , \hat{A}_0 , \hat{S} , \hat{B} , and $\hat{\theta}$, respectively, with normal densities (dashed) superimposed.

Our algorithm has a preliminary stage for estimating the number of beads and producing initial estimates of all parameters except θ , and a final stage for estimating θ , giving maximum likelihood estimates of the other parameters, and accurately determining the number of beads. The preliminary stage assumes zero beads initially and fits $f(x, y) = B$ to the image using OLS. Using the least squares fit at each stage, the information criterion

$$IC^{(k)} = n \ln \left(\frac{RSS}{n} \right) + p \sqrt{n}$$

is computed, where k is the (assumed) number of beads, n is the sample size, RSS is the residual sum of squares, and p is the number of free parameters. Note that IC is an increasing function of RSS and p , which implies that IC rewards a better fit (smaller RSS) and penalizes more free parameters. On the next iteration, one bead is assumed, and so

$$f(x, y) = B + A_0 \cdot \exp \left(-\frac{(x - x_0)^2 + (y - y_0)^2}{S^2} \right)$$

is fit to the image, producing $IC^{(1)}$. Iteration continues until $IC^{(k)} > IC^{(k-1)}$, which indicates that the image contains $k - 1$ beads.

Note that IC is a non-standard information criterion. We found that even the Bayesian Information Criteria (BIC) does not penalize additional parameters sufficiently. This can allow the initial stage of the algorithm to significantly overestimate the correct number of beads, causing much unnecessary computation during the final stage of the algorithm. Our simulations showed that replacing BIC's $\ln(n)$ with \sqrt{n} minimizes overfitting.

As the algorithm makes an initial sweep over possible models, the OLS parameter estimates are saved. Those estimates are used to initialize the maximum likelihood estimation carried out in the final stage. Providing the MLE routine with good initial estimates of all parameters except θ allows the MLE to converge faster than it otherwise would.

The parameter fits and the selection criteria initially computed are then used to find the final parameter estimates and make the final model selection. The key differences are that OLS is replaced by MLE and IC is replaced by the likelihood ratio statistic

$$G_{(beads)}^2 = -2 \left\{ \ell_n(\hat{\beta}_{MLE}^{(beads)}) - \ell_n(\hat{\beta}_{MLE}^{(beads+1)}) \right\},$$

which should be approximately χ^2 distributed with three degrees of freedom (because each bead is associated with three parameters: A_j , x_j , and y_j). The full algorithm is given in pseudocode below.

Algorithm 4.1: LOCATEBEADS($pixels$)

$beads \leftarrow 0$

$\{\hat{\beta}_{OLS}^{(beads)}, IC^{(beads)}\} \leftarrow OLS(pixels, beads)$

repeat

$\left\{ \begin{array}{l} beads \leftarrow beads + 1 \\ \{\hat{\beta}_{OLS}^{(beads)}, IC^{(beads)}\} \leftarrow OLS(pixels, beads) \end{array} \right.$

until $IC^{(beads)} > IC^{(beads-1)}$

$beads \leftarrow beads - 1$

$\{\hat{\beta}_{MLE}^{(beads)}, \ell_n(\hat{\beta}_{MLE}^{(beads)})\} \leftarrow MLE(pixels, beads, \hat{\beta}_{OLS}^{(beads)})$

repeat

$\left\{ \begin{array}{l} beads \leftarrow beads - 1 \\ \{\hat{\beta}_{MLE}^{(beads)}, \ell_n(\hat{\beta}_{MLE}^{(beads)})\} \leftarrow MLE(pixels, beads, \hat{\beta}_{OLS}^{(beads)}) \\ G_{(beads)}^2 \leftarrow -2 \left\{ \ell_n(\hat{\beta}_{MLE}^{(beads)}) - \ell_n(\hat{\beta}_{MLE}^{(beads+1)}) \right\} \end{array} \right.$

until $G_{(beads)}^2 > 7.81 = \chi_{0.95,3}^2$

return $(\hat{\beta}_{MLE}^{(beads+1)})$

5. Simulated Examples. In this section we present a series of simulated examples. We first apply our procedure to a typical image simulated from the Poisson plus Gaussian model presented above. Then we examine the robustness of our procedure by applying it in three atypical scenarios: dim beads, beads in close proximity, and beads that are not entirely contained by the image. Finally, we investigate the sensitivity of our procedure to misspecification of the instrumentation error.

First, we fit an image with fifteen roughly even-spaced fluorophores to verify that our method can handle the substantial numbers that are sometimes found in experimental data. The parameter estimates and their approximate standard errors appear in Table 2. While the time to fit such a larger example is considerable, the method works well and finds the correct number

TABLE 2
Localization for a Typical FIONA Image

Parameter	Value	Estimate	\widehat{SE}	Parameter	Value	Estimate	\widehat{SE}	Parameter	Value	Estimate	\widehat{SE}
x_0	23566	23565.1	0.422	x_5	17753	17753.3	0.421	x_{10}	25394	25393.9	0.422
y_0	4852	4851.1	0.423	y_5	28518	28517	0.423	y_{10}	28468	28468.4	0.421
A_0	15000	15029.4	43.8	A_5	15000	14952.6	43.5	A_{10}	15000	15025.1	43.7
x_1	2522	2522.6	0.423	x_6	28956	28955.8	0.421	x_{11}	29112	29111.9	0.423
y_1	18672	18671.9	0.421	y_6	6771	6770.8	0.421	y_{11}	28770	28770.6	0.422
A_1	15000	15020.7	43.8	A_6	15000	15045.2	43.7	A_{11}	15000	15035.8	43.8
x_2	10475	10475.6	0.424	x_7	27512	27511.7	0.419	x_{12}	18028	18028	0.422
y_2	4858	4858.05	0.423	y_7	3454	3454.6	0.419	y_{12}	21796	21795.9	0.422
A_2	15000	14933.6	43.6	A_7	15000	15104.1	43.8	A_{12}	15000	14928	43.5
x_3	16643	16643.3	0.422	x_8	4165	4164.7	0.422	x_{13}	28318	28317.8	0.434
y_3	19505	19504.4	0.423	y_8	13466	13466.5	0.422	y_{13}	12251	12250.8	0.428
A_3	15000	15073.3	43.9	A_8	15000	15046.2	43.8	A_{13}	15000	15059.3	44.2
x_4	6842	6841.6	0.423	x_9	28960	28959.6	0.421	x_{14}	27757	27756.2	0.432
y_4	16060	16059.9	0.421	y_9	11712	11712.5	0.421	y_{14}	11937	11936.5	0.426
A_4	15000	15013.4	43.7	A_9	15000	14970.3	43.6	A_{14}	15000	14978.9	43.9
S	200	199.9	0.0879	B	200	200.1	0.0678	θ	100	98.3	1.66

TABLE 3
Localization of a Dim Bead

Parameter	Value	Estimate	\widehat{SE}
x_0	8637	8636.8	0.423
y_0	748	748	0.423
A_0	15000	14956.3	49.5
x_1	3372	3372.1	0.422
y_1	1452	1452	0.422
A_1	15000	14975.1	49.4
x_2	5284	5284.6	0.423
y_2	10770	10769.9	0.423
A_2	15000	14941.9	49.4
x_3	6475	6476	5.08
y_3	1812	1806.4	5.14
A_3	400	392	11
S	200	200.1	0.198
B	200	200	0.175
θ	100	101.9	4.30

without difficulties.

Figure 1 shows an image that contains a bead that is very dim ($A = 400$) relative to the image's other three beads ($A = 15000$). The fourth bead is so dim, in fact, that it is invisible to the naked eye, yet our procedure was able to estimate the bead's location to within approximately five nm, which, relative to the other beads, represents a tenfold decrease in resolution for a fortyfold decrease in brightness. Table 3 gives the results for this image. Additionally, simulations showed our algorithm capable of consistently locating (against a background of 200) beads as dim as $A = 75$, which implies a contrast ratio, i.e., the ratio of the brightest pixel value and the background value, equal to 1.4 (versus 75 for a typical bead).

Figure 3 shows an image with two beads whose centers are separated by only 400 nm. Our algorithm was able to distinguish the two beads with only

TABLE 4
Localization of Beads in Close Proximity

Parameter	Value	Estimate	SE
x_0	7149	7148.7	0.421
y_0	9054	9054.6	0.422
A_0	15000	15075.9	47.2
x_1	4343	4342.9	0.447
y_1	4768	4766.8	0.555
A_1	15000	15062.7	53.5
x_2	1782	1782.06	0.421
y_2	10322	10322.3	0.421
A_2	15000	15042.4	47
x_3	10100	10099.9	0.423
y_3	5043	5043	0.422
A_3	15000	14962.5	46.9
x_4	4343	4342.9	0.447
y_4	5168	5168.3	0.555
A_4	15000	14964.9	53.2
S	200	199.8	0.158
B	200	200.3	0.175
θ	100	99.5	4.28

TABLE 5
Localization of a Partial Bead

Parameter	Value	Estimate	SE
x_0	5439	5438.9	0.521
y_0	11650	11650.1	0.881
A_0	15000	14986.3	72.5
x_1	2673	2672.7	0.422
y_1	8711	8711.4	0.422
A_1	15000	15033.1	57
S	200	200	0.290
B	200	200	0.174
θ	100	100.8	4.27

a slight loss of precision in the direction of the line between the beads, as is shown in Table 4.

The second image in Figure 3 shows a bead whose center is only 50 nm from the image's edge. Our algorithm was able to localize the bead with a loss of precision in the y direction that is quite acceptable and perhaps even surprisingly small given that nearly half of the bead is missing. Table 5 reports the results.

Maximum likelihood estimation is often sensitive to model misspecification, and so we investigate the performance of our procedure when the instrumentation error is not $N(0, \theta)$. The instrumentation error for each image has mean zero and variance θ , but otherwise the errors are distributed rather differently. The model was simulated, but with heavy-tailed (t_3 distributed) and asymmetric (exponentially distributed) instrumentation error.





FIG 3. *The upper image shows two beads in close proximity. The lower is an image with a partial bead.*

TABLE 6
Estimation When Instrumentation Error is Heavy-Tailed.

Parameter	Value	Estimate	\widehat{SE}
x_0	7823	7822.6	0.422
y_0	3353	3353	0.423
A_0	15000	15005.8	61.9
S	200	199.9	0.341
B	200	199.9	0.174
θ	100	101.4	4.28

TABLE 7
Estimation When Instrumentation Error is Asymmetric

Parameter	Value	Estimate	\widehat{SE}
x_0	7823	7823	0.424
y_0	3353	3353.1	0.424
A_0	15000	14921.3	61.5
S	200	200.7	0.343
B	200	200.1	0.176
θ	100	108.6	4.38

More specifically, the model was simulated according to

$$Z_i \sim Poi(f_i) + \sqrt{\theta/3} t_3$$

and according to

$$Z_i \sim Poi(f_i) + Exp(\sqrt{\theta}) - \sqrt{\theta},$$

where $Poi(\lambda)$ denotes the Poisson distribution with rate λ , t_ν denotes the t distribution with ν degrees of freedom, and $Exp(\lambda)$ denotes the exponential distribution with mean λ . Tables 6 and 7 show that localization was not affected by these misspecifications.

6. Analysis of an Experimentally Observed FIONA Image. In this section we apply our procedure to an experimentally observed FIONA image shown in Figure 4. Table 8 shows our parameter estimates for this image.

To verify the fit of the image, we ran numerous diagnostics to verify that the observed data originates from our random field model. Our approximate model implies that for the i th pixel

$$Z_i \overset{\sim}{\sim} N(f_i, f_i + \theta),$$

spatially independent of the other pixels. If our model is correct, then we should have, for the i th error,

$$\epsilon_i = \frac{Z_i - f_i}{\sqrt{f_i + \theta}} \overset{\sim}{\sim} N(0, 1),$$





FIG 4. An experimentally observed FIONA image.

TABLE 8
Parameter Estimates for an Experimentally Observed FIONA Image.

Parameter	Estimate	\widehat{SE}
x_0	12168.1	4.83
y_0	7570.4	4.79
A_0	269.8	10.5
x_1	12296.1	4.71
y_1	16509	4.63
A_1	275.7	10.3
S	175.7	3.31
B	33.4	0.0418
θ	80.8	0.632

spatially independent of the other errors. This implies that the variogram

$$\gamma(\mathbf{h}) = \frac{1}{2} \text{Var} \{ \epsilon(\mathbf{s} + \mathbf{h}) - \epsilon(\mathbf{s}) \}$$

should equal one for all locations \mathbf{s} and lag (displacement) vectors \mathbf{h} .

We plot empirical (residual) variograms to determine spatial independence and use a normal probability plot to check for normality (Cressie and Hawkins, 1980). The plots for our example image are shown in Figure 5. Except for some anomalous features in the lower tail of the probability plots, the diagnostics give a good indication that our proposed model is sound. The standardized residuals were also checked and no blatant violations of what would be expected for independent, identically distributed data were found.

7. Conclusions. The method outlined in this paper allows for the automated analysis of FIONA images, including the ability to select the number of fluorophores in an image. By using a likelihood framework, the method also allows for standard errors to be calculated simultaneously with the estimates. The method was then verified through simulation and the analysis of collected data. We hope that this case study will serve as an example of applying traditional statistical theory to enhance the analysis of nanoscale experimental methods where algorithmic approaches have been favored.

Since this method is largely automated through model selection techniques, it can handle the analysis of “movies” by processing each frame. Since the method also returns standard errors for the locations of the fluorophores, this opens the possibility of creating tracking methods to follow dynamic specimens using not only the position data but the information on observational errors which are given.

Estimation and model selection can be done using our free C++ application, *beads*, which makes extensive use of the GNU Scientific Library, and fit

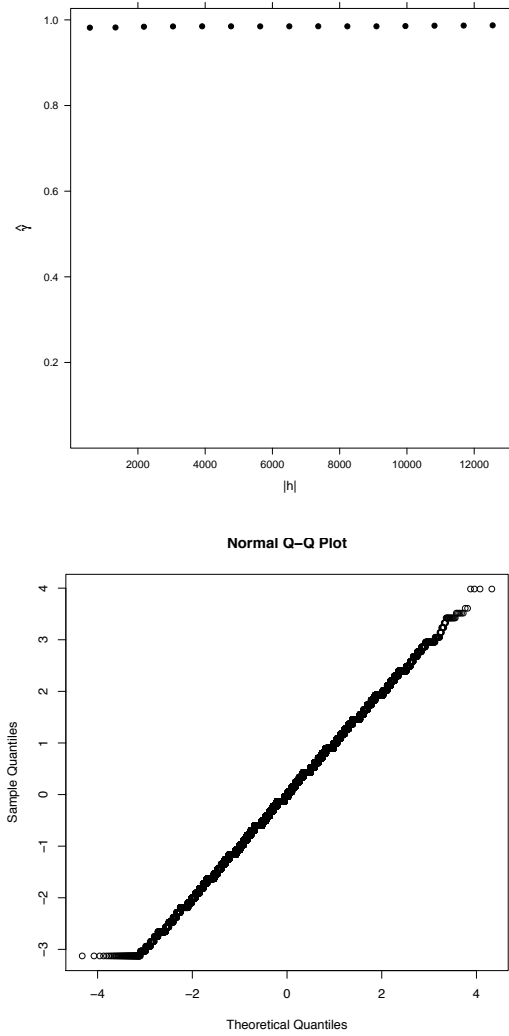


FIG 5. A variogram plot and a normal probability plot of the standardized residuals for the real FIONA image.

diagnostics can be carried out using our free R software package, FIONAdiag (Galassi, Davies, Theiler, Gough, Jungman, Booth, and Rossi, 2008).

References.

- Norman Bobroff. Position measurement with a resolution and noise-limited instrument. *Review of Scientific Instruments*, 57(6):1152–1157, 1986.
- Michael K. Cheezum, William F. Walker, and William H. Guilford. Quantitative comparison of algorithms for tracking single fluorescent particles. *Biophysical Journal*, 81(4):2378–2388, 2001.
- Noel Cressie and Douglas M. Hawkins. Robust estimation of the variogram: I. *Mathematical Geology*, 12(2):115–125, 1980.
- Mark Galassi, Jim Davies, James Theiler, Brian Gough, Gerard Jungman, Michael Booth, and Fabrice Rossi. *GNU Scientific Library Reference Manual*, 1.11 edition, February 2008.
- Comert Kural, Hamza Balci, and Paul R. Selvin. Molecular motors one at a time: Fiona to the rescue. *Journal of Physics: Condensed Matter*, 17(47):S3979–S3995, 2005.
- F.W.D. Rost. *Fluorescence Microscopy*. Cambridge University Press, 1992.
- M. J. Saxton. Single-particle tracking: the distribution of diffusion coefficients. *Biophysical Journal*, 72(4):1744–1753, 1997.
- Russell E. Thompson, Daniel R. Larson, and Watt W. Webb. Precise nanometer localization analysis for individual fluorescent probes. *Biophysical Journal*, 82(5):2775–2783, 2002.
- Ahmet Yildiz, Joseph N. Forkey, Sean A. McKinney, Taekjip Ha, Yale E. Goldman, and Paul R. Selvin. Myosin v walks hand-over-hand: Single fluorophore imaging with 1.5-nm localization. *Science*, 300(5628), 2003.

DEPARTMENT OF STATISTICS
PENN STATE UNIVERSITY
UNIVERSITY PARK, PA 16802
E-MAIL: jph264@psu.edu
E-MAIL: fricks@stat.psu.edu

DEPARTMENT OF BIOENGINEERING
THE PENNSYLVANIA STATE UNIVERSITY
UNIVERSITY PARK, PA 16802
E-MAIL: wohbio@engr.psu.edu