

Bayesian Image Recovery for Dendritic Structures Under Low Signal-to-Noise Conditions

Geoff Fudenberg and Liam Paninski

Department of Statistics and Center for Theoretical Neuroscience

Columbia University

<http://www.stat.columbia.edu/~liam>

Abstract

Experimental research seeking to quantify neuronal structure constantly contends with restrictions on image resolution and variability. In particular, experimentalists often need to analyze images with very low signal-to-noise ratio (SNR). In many experiments dye toxicity scales with the light intensity; this leads experimentalists to reduce image SNR in order to preserve the viability of the specimen. In this work we present a Bayesian approach for estimating the neuronal shape given low-SNR observations. This Bayesian framework has two major advantages. First, the method effectively incorporates known facts about 1) the image formation process, including blur and the Poisson nature of image noise at low intensities, and 2) dendritic shape, including the fact that dendrites are simply-connected geometric structures with smooth boundaries. Second, we may employ standard Markov chain Monte Carlo (MCMC) techniques for quantifying the posterior uncertainty in our estimate of the dendritic shape. We describe an efficient computational implementation of these methods and demonstrate the algorithm’s performance on simulated noisy two-photon laser-scanning microscopy images.

1 Introduction

In neuroscience, what we see often limits what we know. Improving imaging capabilities holds great promise for innovative experimental work. For example, quantitative analysis of dendritic spine morphology has the potential to teach us a great deal about synaptic transmission [Engert and Bonhoeffer, 1999, Araya et al., 2006] and long-term

synaptic plasticity [Nimchinsky et al., 2004]. The opportunities for insight into the function of dendritic spines have already spurred computational work on automated analysis of dendritic spine morphology [Rodriguez et al., 2008, Koh et al., 2002, Farivar, 2005, Xu and Wong, 2006, Cheng et al., 2007].

Clearly, quantitative experiments can directly benefit from improved imaging procedures. In lieu of designing a new imaging apparatus, this paper focuses on methods for making better use of currently obtainable data, especially within a sparse low signal-to-noise (SNR) regime. As a key example, we focus on algorithms for recovering two-photon laser scanning microscopy (TPLSM) images. While TPLSM offers increased optical resolution without increased phototoxicity as compared with conventional confocal microscopy [Denk et al., 1994], phototoxicity persists within the focal slice and limits attainable resolution: to avoid damaging the sample, imaging at long time scales or high intensities is impossible. Since photon detection during imaging can be modeled as a Poisson process [van Kempen et al., 1997], in which the signal-to-noise varies directly with the total number of photons absorbed, these biophysical limitations on the maximal intensity force us to deal with low-SNR images.

In general, the image recovery problem may be posed as follows: we make noisy, blurred observations of some underlying “true” neural shape S_{true} — schematically,

$$I_{obs} = S_{true} * w + noise,$$

where I_{obs} denotes the observed image data, and $*$ denotes a convolution by w , the point-spread function (or PSF) — and our goal is to recover the true input image as faithfully as possible. Of course, exact recovery of this true input image is rarely feasible, and so we also need to know how uncertain we are about our estimate: in a sense, we would like to put “errorbars” around our estimate \hat{S} .

Many groups have contributed to the general problem of restoring noisy blurred images using maximum likelihood (ML) approaches [Conchello and McNally, 1996, Markham and Conchello, 2001]. For the case of Poisson noise, Richardson and Lucy [Richardson, 1972, Lucy, 1974] independently introduced iterative deconvolution algorithms.

Significant research has been devoted to developing Bayesian statistical machinery for use in astrophysical [Molina, 1994] and medical imaging problems, especially positron emission tomography (PET) [Vardi et al., 1985, Kontaxakis et al., 1999].

We build upon this extensive previous imaging literature by incorporating important prior information about dendrites. In particular, we know that dendrites are simply connected geometrical structures, with fairly smooth boundaries. We may further restrict our attention to binary images: a given pixel may be either inside or outside the dendrite¹. Furthermore, we focus our attention on cases where it is possible to generate an initial binary image with the correct topology. Thus, conceptually, to recover the true neural shape we need only search over the space of topologically equivalent binary images with smooth boundaries (of course, we do not attempt a brute-force search over this space, which would be computationally impractical). By combining the statistical model for image degradation with this *a priori* information, we can apply powerful likelihood-based tools from Bayesian statistics to the problem of optimally recovering dendritic shape (including spine size, etc.). In particular, we develop efficient Markov chain Monte Carlo (MCMC) methods for computing the optimal estimate of the underlying neuronal shape, while at the same time quantifying our uncertainty about this estimate.

2 The blurred Poisson image formation model

We begin by describing our model of how the observed image depends on the true underlying neuronal shape. Define $I_{obs}(s, t, u)$ to be the photon count observed in the (s, t, u) -th pixel, and $S_{true}(x, y, z)$ to be the true underlying (unobserved) neuronal shape: $S_{true}(x, y, z)$ is either one or zero, depending on whether the location (x, y, z) is inside or outside of the neuron, respectively. (Note that the pixelization of S_{true} is user-defined — we can attempt to reconstruct the neuronal shape at any resolution we desire — while the pixelization of I_{obs} is set by the imaging apparatus. Therefore, the observed pixelization may generally be coarser than that of S_{true} .)

¹Note that this is equivalent to the assumption that neuronal fluorescence is uniform inside the neuron, i.e., the dye concentration equilibrates completely. This assumption has been debated in the literature [Simoni et al., 2004, Svoboda, 2004] but provides a reasonable starting point for the analysis presented here.

The observed photon counts per pixel $I_{obs}(s, t, u)$ are assumed to be discretized observations of a Poisson process; the rate of each individual Poisson count at pixel (s, t, u) is given by the convolution

$$\lambda(s, t, u) = \int_{s-ds/2}^{s+ds/2} \int_{t-dt/2}^{t+dt/2} \int_{u-du/2}^{u+du/2} [\lambda_{out} + (\lambda_{in} - \lambda_{out})S_{true}(x, y, z)]w(s-x, t-y, u-z)dx dy dz, \quad (1)$$

where $w(x, y, z)$ is the point-spread (blur) function, or PSF, and λ_{out} represents the baseline fluorescence outside of the neuron and λ_{in} represents the internal fluorescence; typically, λ_{in} is greater than λ_{out} . A reasonable definition of the signal-to-noise ratio here is

$$SNR = \frac{\lambda_{in} - \lambda_{out}}{\sqrt{\lambda_{out}}};$$

this is the difference between the mean fluorescence inside and outside the neuron, normalized by the standard deviation of the baseline (extracellular) photon count in a pixel of size one.

The image space \mathcal{S} is the space of all binary simply connected shapes; that is, all possible shapes S such that $S(x, y, z)$ is either one or zero (inside or outside the neuron, respectively) and such that the interior (the collection of pixels (x, y, z) satisfying $S(x, y, z) = 1$) is simply connected: we may connect any two “inside” pixels by a (possibly curved) continuous path which lies entirely inside the neuron. Thus we implicitly assume that the pixelization (x, y, z) is sufficiently fine that no pixels are partially inside and partially outside the neuron. We further assume (as discussed in the introduction above) that the fluorescent dye is fully equilibrated in the sample: there are no spatial fluctuations in λ_{in} (caused, e.g., by fluctuations in internal dye concentration) or λ_{out} .

Thus the model is fully specified by the parameters $(S, \lambda_{in}, \lambda_{out}, w)$. This model is fairly standard in the imaging literature, and may be extended in a straightforward manner in a number of natural ways: non-Poisson noise, non-homogeneous blurring, non-homogeneous baseline fluorescence, etc. However, for simplicity we will only treat the homogeneous Poisson case here. We will further assume the parameters $(\lambda_{in}, \lambda_{out}, w(., .))$ to be known properties of the fluorescent dye and imaging apparatus, respectively (though of course we may in general attempt to simultaneously infer these parameters, too; e.g., Fig. 5). Finally, we will restrict

our attention to the two-dimensional case for illustrative purposes.

3 Computing the likelihood

Now that our model is defined properly, our task is to quantify the posterior distribution $p(S|I_{obs}, \lambda_{in}, \lambda_{out}, w)$ over neuronal shapes S given the observed data I_{obs} and the imaging parameters $(\lambda_{in}, \lambda_{out}, w)$. By Bayes' rule we know that we can express this posterior as a normalized product of two terms,

$$p(S|I_{obs}, \lambda_{in}, \lambda_{out}, w) \propto p(I_{obs}|S, \lambda_{in}, \lambda_{out}, w)p(S),$$

where the first term is the likelihood of observing I_{obs} given S and the second term is the prior probability of S .

We will begin by examining the likelihood. Our model states that the photon counts in each individual pixel constitute independent Poisson random variables. Thus the likelihood of observing a count of n_i in the i -th pixel is given by the Poisson distribution with rate λ_i ,

$$Poiss(n_i|\lambda_i) = \frac{e^{-\lambda_i} \lambda_i^{n_i}}{n_i!} \quad n_i = 0, 1, 2, \dots,$$

where λ_i is calculated via equation (1). Now to obtain the likelihood of the observed array $\{n_i\}$ of counts (where i indexes every pixel in $I_{obs}(s, t, u)$) we simply form the product

$$p(I_{obs}|S, \lambda_{in}, \lambda_{out}, w) = \prod_i \frac{e^{-\lambda_i} \lambda_i^{n_i}}{n_i!}.$$

It is convenient to work with the log-likelihood instead:

$$\log p(I_{obs}|S, \lambda_{in}, \lambda_{out}, w) = \sum_i (n_i \log \lambda_i - \lambda_i) + \text{const.}, \quad (2)$$

where the constant does not depend on S , and may therefore be ignored.

4 Maximum likelihood estimation

Before we describe methods for computing the posterior $p(S|I_{obs})$, it is worth examining the somewhat simpler problem of computing the maximum likelihood estimate for S ,

$$\hat{S}_{MLE} = \arg \max_{S \in \mathcal{S}} p(I_{obs}|S)$$

(we have suppressed the dependence of the likelihood function $p(I_{obs}|S)$ on the imaging parameters $(\lambda_{in}, \lambda_{out}, w)$ for simplicity).

To compute \hat{S}_{MLE} , we must search over the image space \mathcal{S} . Of course, this space is far too large to search directly; thus some kind of local search algorithm is necessary. We will describe the simplest version of this local search, in which we iteratively change the state of one pixel at a time in our current estimate of the true underlying neuronal shape S . It is important to note that \mathcal{S} is connected: any two simply connected shapes are topologically equivalent, in the sense that we can continuously deform one shape into another (no “tearing” or “gluing” is allowed). In this setting, continuous deformations are constructed by composing a sequence of single pixel flips which do not “tear” S (break it up into two disconnected pieces) or “glue” S on itself (form a loop of the dendritic shape onto itself). However, while the underlying space \mathcal{S} is connected in this sense, the likelihood function typically has many local optima, and therefore local search techniques must be supplemented with a randomized or multistart strategy (e.g., simulated annealing) in order to find the global optimum. We provide details of the local search strategy in the following subsections. For information on our particular initialization procedure, see the appendix.

4.1 The boundary set

On each iteration we choose a pixel randomly and decide whether or not to “flip” it (i.e., add the pixel to the interior set if it is currently outside, or remove the pixel if it is inside). Before deciding whether or not to flip a given pixel, we must first test whether flipping the pixel will preserve the simply connected structure of S . Clearly, we are not allowed to flip pixels which are not located on the boundary of S , since flipping a pixel which is not touching S will

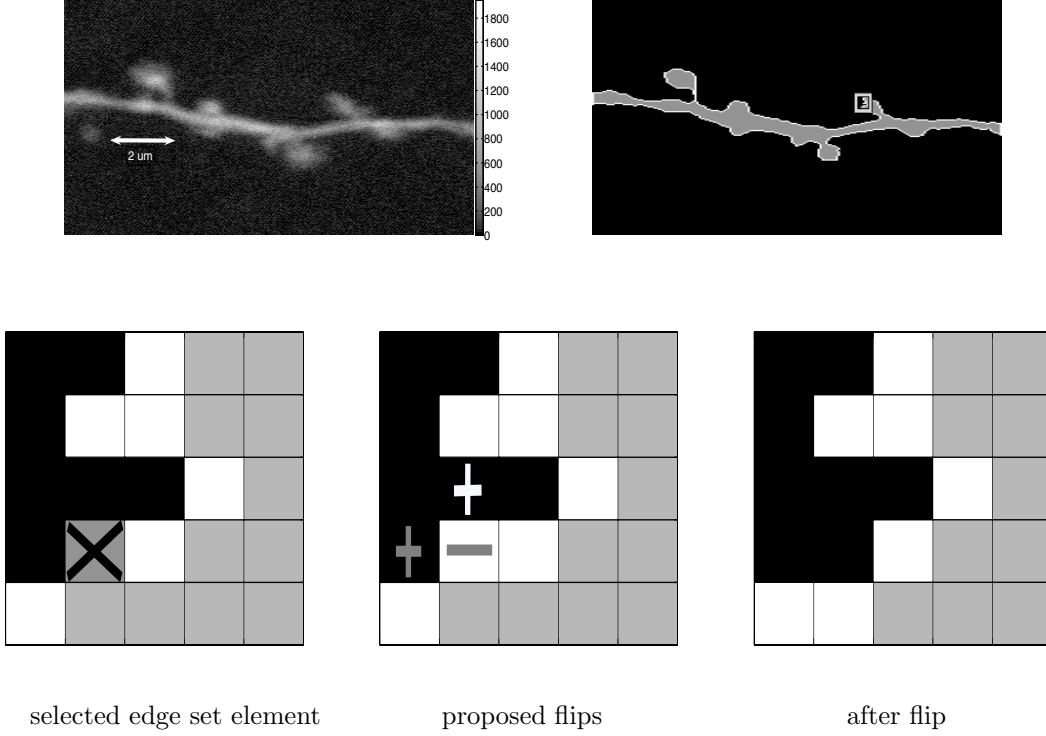


Figure 1: Illustration of iterative pixel-flipping algorithm. **Top Left:** raw sample image of a dendrite segment (courtesy of R. Araya and R. Yuste); colorbar indicates the observed counts per pixel. Image taken at 30 pixels per micron. **Top Right** a sample S , showing interior (gray and white), exterior (black), and boundary set (white) pixels. The rectangle encloses the area around the pixels displayed in the lower three panels. **Lower Panels:** these three panels illustrate a local area of the neuron before and after flipping a selected edge set pixel from inside to outside the neuron. **Left:** We begin by randomly choosing a pixel in the boundary set (dark grey pixel marked with black X). **Middle:** Negative (-) signs mark possible removals in the neighborhood of the selected pixel; positive (+) signs show possible flips which would add exterior neighbors. Gray symbols are the allowed flips (these do not violate simple connectivity), while the white symbol marks a forbidden flip: adding this pixel would create a hole in S , violating simple connectivity. If a proposed flip is accepted, elements of the edge set must be updated. **Right:** The edge set has been updated upon removing the allowed gray (-) pixel.

create a disconnected shape (the current S plus a disconnected pixel), and flipping a pixel which is on the interior of S but not on the boundary will form a hole, therefore breaking the simple connectivity of S . Thus we may restrict our attention to the boundary set of S : the set of all pixels which are in the interior of S but which contact at least one pixel on the exterior. For concreteness, we say that a pixel contacts another if they are direct neighbors either horizontally or vertically (i.e., diagonal contacts are not considered). An example of

an edge pixel flip is shown in Fig. 1.

After each accepted flip, we must update the boundary set, by removing any pixels in the interior which are now not touching the exterior and adding any interior pixels which are now on the boundary. As we discuss below, this is a local computation: we do not need to rescan the entire image after each accepted flip; instead, we need only check the four neighbors of the flipped pixel.

4.2 Simple connectivity constraint

Each possible flip is checked to ensure maintenance of the simple connectivity constraint. While simple connectivity is a global topological constraint, we can enforce it locally in two dimensions [Kong and Rosenfeld, 1989]. Mathematically, this is possible by forbidding topological changes: flips which either cut a region in two (tearing) or join two unconnected regions together (gluing). We may check these constraints in a computationally efficient, local manner by requiring the number of connected regions in a 3×3 neighborhood of the proposed pixel (i.e., the eight horizontal, vertical, and diagonal neighbors) to remain constant when we flip the pixel from the neuron.

This check on the connectivity may be performed using, e.g., the `bwlabel.m` function in Matlab, which serves to count all the connected regions in a given image. However, it is inefficient to perform this check every time we want to flip a pixel. Instead, we cache 3×3 patches that we have tested previously, along with whether the flip was accepted or not. Now for each new patch that we need to test, we need only check to see if we have checked this patch before, and if so, whether it led to a valid flip.

In the data we present, we assume only one neuron is labeled. This has an important correspondence with experiments, as some techniques label single neurons, while others typically label many cells. However, our algorithm maintains the topology of initialization— if two neuronal structures started out next to each other, but were separate, then the algorithm would keep them separate. Again, this behavior relies on the fact that by rejecting changes in local topology, we maintain global topology. If the algorithm was used to recover an area with two nearby neurons, and was initialized with separated shapes, it might incorrectly assign a

spine to the wrong neuron— leading to a less favorable recovery— but the two structures would remain separate.

4.3 Local updating of the loglikelihood

Once we have decided that a given pixel flip is acceptable (i.e., it maintains simple connectivity of S), we need to decide whether the flip will increase the likelihood. In a strict ascent algorithm, we will only accept a flip if it increases the likelihood function, while in a simulated annealing (randomized) algorithm, the probability of accepting a flip increases as a function of the ratio between the likelihood of S after and before the flip. In either case, we need to compare the post-flip and pre-flip likelihood.

Computing the log-likelihood (equation (2)) requires that we compute λ_i (equation (1)) and then perform a sum over all pixels i . However, if the point-spread function $w(\cdot)$ has finite support, then updating the log-likelihood requires just a simple (fast) local computation. To see this, we expand the log-likelihood ratio,

$$\begin{aligned}
\log \frac{p(I_{obs}|S_{after})}{p(I_{obs}|S_{before})} &= \log p(I_{obs}|S_{after}) - \log p(I_{obs}|S_{before}) \\
&= \sum_i (n_i \log \lambda_i^{(after)} - \lambda_i^{(after)}) - \sum_i (n_i \log \lambda_i^{(before)} - \lambda_i^{(before)}) \\
&= \sum_i \left[n_i \left(\log \lambda_i^{(after)} - \log \lambda_i^{(before)} \right) - \left(\lambda_i^{(after)} - \lambda_i^{(before)} \right) \right] \\
&= \sum_j \left[n_j \left(\log \lambda_j^{(after)} - \log \lambda_j^{(before)} \right) - \left(\lambda_j^{(after)} - \lambda_j^{(before)} \right) \right] \\
&= \sum_j n_j \left(\log \lambda_j^{(after)} - \log \lambda_j^{(before)} \right) - \sum_j \left(\lambda_j^{(after)} - \lambda_j^{(before)} \right),
\end{aligned}$$

where the sum over j is only over those pixels for which $\lambda_j^{(after)} \neq \lambda_j^{(before)}$. In many applications, $w(\cdot)$ is only a few pixels wide, and, by linearity of convolution, changing one pixel in S will only affect λ over a few pixels. The sum over j will be much smaller (and faster to compute) than the sum over i . Furthermore, the second sum in the last line above is in fact a constant, again by linearity of convolution:

$$\sum_j \left(\lambda_j^{(after)} - \lambda_j^{(before)} \right) = \pm(\lambda_{in} - \lambda_{out})W,$$

where W is the space integral of the point spread function $w(\cdot)$ and the \pm is taken to be positive when the pixel is flipped in (added to the interior set) and negative when the pixel is flipped out (removed from the interior). So finally we obtain

$$\log \frac{p(I_{obs}|S_{after})}{p(I_{obs}|S_{before})} = \sum_j n_j \left(\log \lambda_j^{(after)} - \log \lambda_j^{(before)} \right) \pm (\lambda_{in} - \lambda_{out})W.$$

Finally, note that we only need to compute λ_j for pixels such that n_j is positive. We achieve this by employing sparse matrices to store the location of positive n_j for a tested pixel as we sample from S . Thus, updates for the loglikelihood may be computed quite quickly.

4.4 Direct maximum likelihood performs poorly

The performance of the maximum likelihood estimator \hat{S}_{MLE} is illustrated in Fig. 2. We started with the binary test image S_{true} on the left, then convolved this image with an isotropic Gaussian point-spread function $w(\cdot)$ (the standard deviation of this Gaussian was taken to be 3 pixels here), then sampled photon counts $n(x, y)$ from the blurred Poisson model as described above. The images of \hat{S}_{MLE} are arranged opposite data with increasing SNR below the ‘true’ shape. The top row uses $\lambda_{in} = 2\lambda_{out}$, the middle row uses $\lambda_{in} = 5\lambda_{out}$, and the bottom row uses $\lambda_{in} = 500\lambda_{out}$. For each set of data we computed \hat{S}_{MLE} by direct ascent methods, but similar results are observed when various forms of simulated annealing are employed.

The main result evident here is that directly computing \hat{S}_{MLE} leads to rather poor image recovery in the low-SNR regime. Roughly speaking, the MLE tries to include all pixels where many photons have been detected and to exclude all pixels in which no photons are detected; under the simple-connectivity constraint, this behavior leads to the undesirable “tendrils” (high perimeter-to-area ratio) structures seen in Fig. 2. This behavior is in fact evident over a fairly wide range of SNR regimes; only when the SNR is high enough does the MLE become a viable estimator for S .




Figure 2: The unpenalized MLE gives an overly “stringy” estimate of the underlying neuron shape. This figure shows how MLE performance scales with increasing SNR. The left column displays I_{obs} (simulated) with increasing SNR generated by the blurred Poisson model and the right column shows the corresponding \hat{S}_{MLE} . The top row uses $\lambda_{in} = 2\lambda_{out}$, the middle row uses $\lambda_{in} = 5\lambda_{out}$, and the bottom row uses $\lambda_{in} = 500\lambda_{out}$. With low-SNR the unpenalized MLE clearly fails to accurately recover the original image; note the undesirable “tendrill” (high perimeter per area) structures. As SNR increases, however, the data more adequately constrains recovery and \hat{S}_{MLE} approaches S_{true} .

5 Penalized maximum likelihood leads to better image reconstruction

Given our prior knowledge about the relative smoothness of neuronal edges, we know that the overly “hairy” recovered neuronal shape shown in Fig. 2 is inaccurate. We would like to

build this *a priori* information directly into our estimator. One direct and fairly classical way to do this is to maximize a penalized likelihood instead of maximizing the likelihood directly: we may obtain a smoother reconstruction by penalizing flip proposals which make the edge more jagged. Thus we maximize the penalized loglikelihood

$$\log p(I_{obs}|S) - Q(S) \tag{3}$$

instead of just the loglikelihood $\log p(I_{obs}|S)$. We want to choose the penalty function $Q(S)$ so that $Q(S)$ becomes larger as S becomes more jagged. However, we constrain ourselves to penalty functions Q which may be updated through strictly local computations, so that our iterative algorithm (which requires many pixel flips) remains computationally tractable. As usual, this penalized likelihood has a natural Bayesian interpretation: if our prior distribution on images is of the form

$$p(S) \propto e^{-Q(S)},$$

then the penalized loglikelihood is just the log-posterior under this prior, and the maximum penalized likelihood estimator is equivalent to the maximum *a posteriori* estimator.

Figure 3 illustrates the behavior of a penalized MLE on the data shown in the middle row of Fig. 2. We found the following $Q(S)$ to be a simple, effective penalizer here:

$$Q(S) = \alpha_1 Q_1(S) + \alpha_2 Q_2(S),$$

where

$$Q_1(S) = \text{number of exterior pixels neighboring a member of the boundary set of } S$$

$$Q_2(S) = \text{number of elements in the boundary set of } S,$$

and α_1 and α_2 are adjustable parameters. In this case both $Q_1(S)$ and $Q_2(S)$ serve to measure the roughness of the boundary of S , and may both be computed efficiently via simple local computations (specifically, counting the number of pixels in or near the boundary set of S);

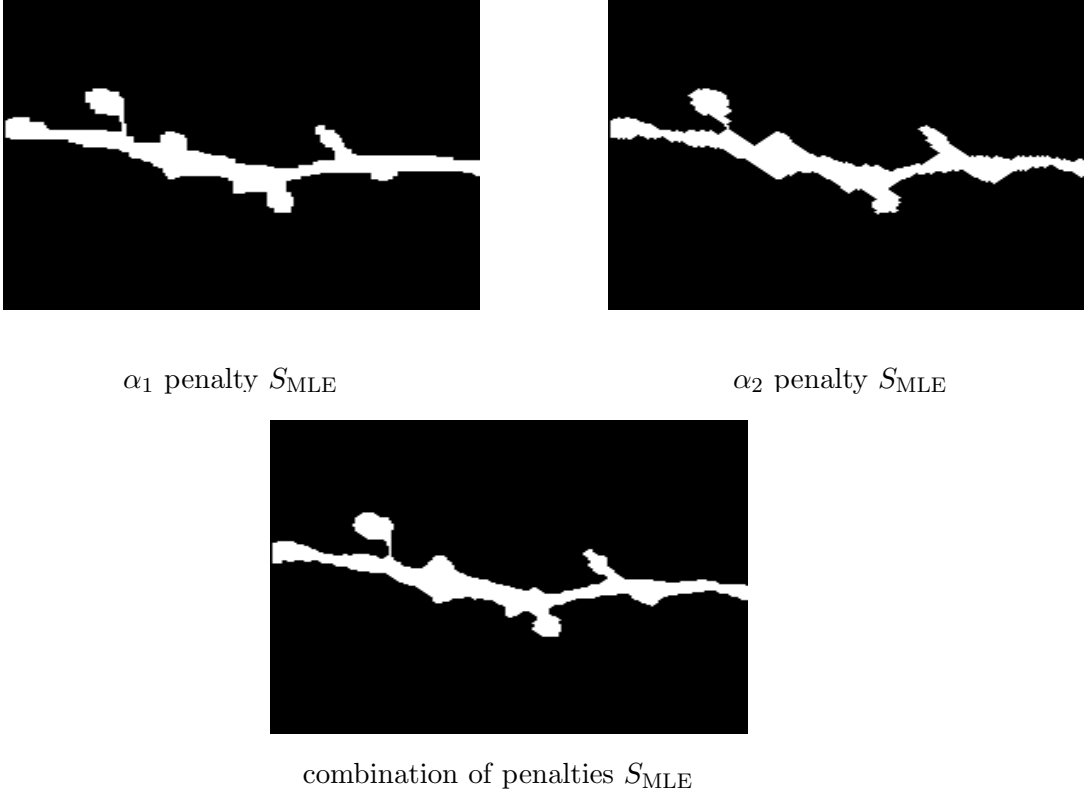


Figure 3: Illustration of the edge-length penalties $Q_1(S)$ and $Q_2(S)$. Each panel displays the penalized MLE, computed with $\alpha_2 = 0$ (left) and $\alpha_1 = 0$ (right), given the observed data shown in the middle panel ($\lambda_{in} = 5\lambda_{out}$) in Fig. 2. In each case, the penalized MLE eliminates the “tendrils” seen in the unpenalized MLE; however, employing $Q_1(S)$ or $Q_2(S)$ alone leads to somewhat blocky reconstructed images. By combining penalties (i.e., setting both α_1 and α_2 positive), we can more optimally recover a neuronal shape.

we found empirically that a combination of these two penalties leads to better recovery of the true shape S , since $Q_1(S)$ alone favors rectangular, blocky edges, while $Q_2(S)$ selects for diagonal edges.

5.1 Fitting the smoothing penalty parameters via cross-validation

The definition of the penalty function $Q(S)$ above poses the obvious, yet nontrivial, problem: how do we optimally choose the penalty weights α_1, α_2 ? A satisfactory analytic treatment of this problem has proven elusive; any definition of the “optimal” penalty value would depend upon many parameters (e.g., the signal-to-noise ratio of the observed data, the smoothness

of the true underlying image, etc.). Moreover, direct numerical minimization of the penalized loglikelihood (3) as a function of (S, α_1, α_2) clearly leads to the unpenalized maximum likelihood solution (i.e., α_1 and α_2 are just set to zero), which is clearly not the solution we are looking for (c.f. Fig. 2). Thus we took a cross-validation approach to optimizing the penalty values here.

Cross-validation refers to the technique of splitting the observed data into two components, the “training” set and the “held-out” (or “test”) set. Fitting is performed on the training data, and then the performance of the estimator (under various values of the penalization parameters) is tested on the held-out data. Using cross-validation, we then choose the penalization parameters which perform best on the held-out test data. The best-performing parameters are those for which the penalized MLE maximizes the (unpenalized) likelihood of the test data.

In the interest of minimizing imaging damage to the sample, we would like to be able to fit the penalty parameters given a single image. Thus, we randomly hold out a fraction of the observed *pixels* in a single image. This entails a straightforward change in our loglikelihood function: instead of computing the sum over all pixels i in expression (2), we compute the sum over all *observed* pixels i' instead — we simply discard the held-out test pixels from the sum.

Thus, for each candidate value of the penalization parameters (α_1, α_2) , we compute

$$\hat{S}_{\alpha_1, \alpha_2} = \arg \max_{S \in \mathcal{S}} \sum_{i'} (n_{i'} \log \lambda_{i'} - \lambda_{i'}) - \alpha_1 Q_1(S) - \alpha_2 Q_2(S),$$

where the fluorescence rate at the i' -th pixel given S , $\lambda_{i'}$, is computed as described above. Then to choose the best (α_1, α_2) , we compare the loglikelihood each $\hat{S}_{\alpha_1, \alpha_2}$ assigns to the test data:

$$\sum_{j \in \text{held-out}} (n_j \log \lambda_j - \lambda_j),$$

where the sum is over all the held-out test pixels and now λ_j is computed given the shape $\hat{S}_{\alpha_1, \alpha_2}$ computed using the training data only.

Applying this method returns non-zero values for penalty parameters (α_1, α_2) since— as

shown in Fig. 2– naïve MLE reconstruction at low-SNR is underconstrained by the data. Imposing the edge penalty constrains image recovery and increases the loglikelihood of the the held-out test data. The success of our cross-validation approach is shown in Fig. 4: we see that the cross-validated test loglikelihood does a good job of choosing penalty values which lead to the best-recovered neuronal shapes S for known $\lambda_{in} = 5\lambda_{out}$. The recovery using parameters chosen by cross-validation incorrectly classifies only a slightly higher number of pixels than the parameters which best recovered the original image; the respective percentages of erroneously classified pixels divided by number of pixels in S_{true} are 7.00% and 7.36%.

Running the algorithm on similar simulated data as in Fig. 4, we are able to simultaneously estimate the penalty parameters (α_1, α_2) and fluorescence levels $(\lambda_{in}, \lambda_{out})$. Of course, this full search covers a larger parameter space and is therefore slower. Fig. 5 shows the loglikelihood surface as a function of λ_{in} and λ_{out} at optimal values of α_1, α_2 . We see the surface is fairly well behaved; we may fit fluorescence levels in addition to penalty values for (α_1, α_2) . Cross-validation, then can robustly determine favorable image recovery parameters.

6 Markov Chain Monte Carlo methods for sampling from

$$p(S|I_{obs})$$

Now we may finally turn to the primary goal of this paper, which is to develop fully Bayesian methods for quantifying our knowledge of the neuronal shape S given the observed image data I_{obs} . Up until now we have discussed techniques for computing estimators \hat{S} which (locally) maximize the posterior $p(S|I_{obs})$. However, it is equally important to quantify our uncertainty in these estimates. One standard way to do this is to draw samples $S^{(i)}$ from the posterior $p(S|I_{obs})$, and then to quantify the variability of these samples. Directly sampling from these posteriors is infeasible, but standard Markov Chain Monte Carlo [Robert and Casella, 2005] methods may be employed easily.

We implemented a simple Gibbs sampler [Geman and Geman, 1984] here. The idea is that we iteratively choose a pixel i randomly and sample from $p(S_i|I_{obs}, \alpha_1, \alpha_2, \{S_j\}_{j \neq i})$, the distribution of the states of the pixel S_i given the observed data I_{obs} , the param-

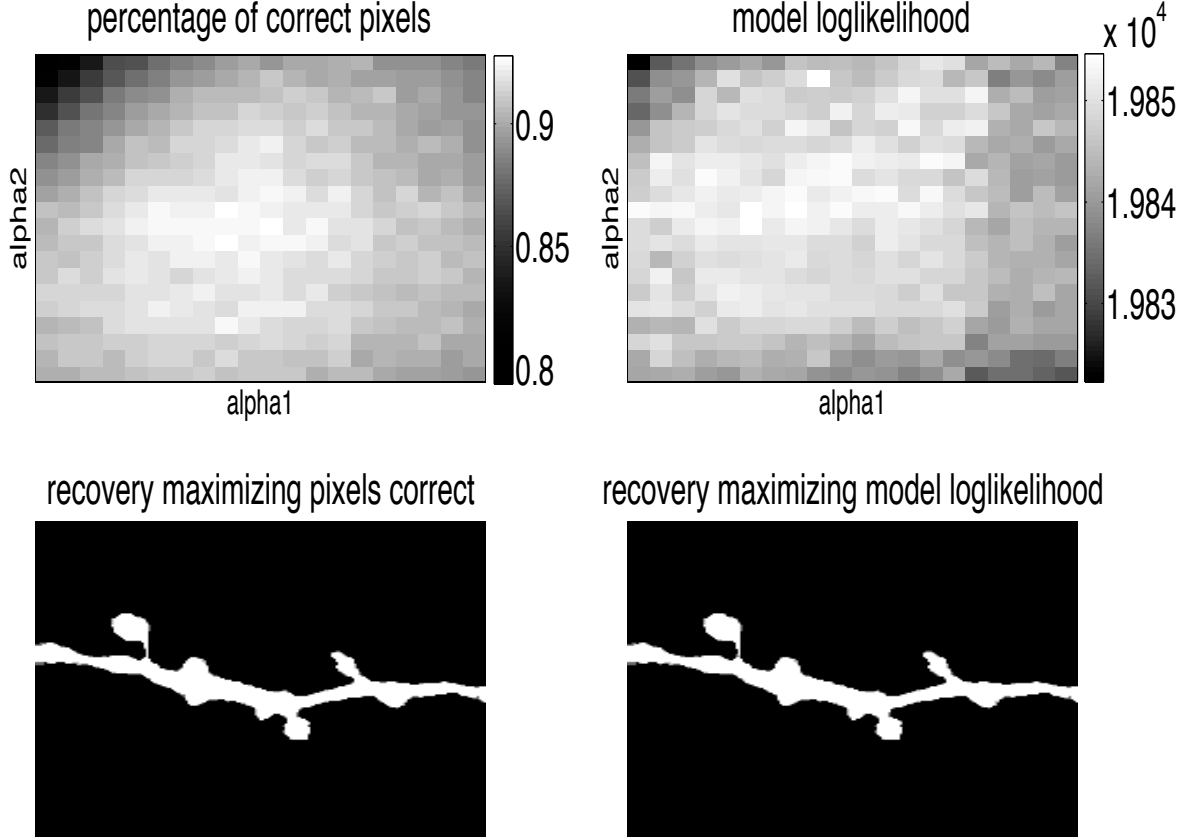


Figure 4: Illustration of the performance of cross-validation in selecting the optimal penalization parameters (α_1, α_2) for data shown in middle panel of Fig. 2 ($\lambda_{in} = 5\lambda_{out}$). **Top:** comparison of cross-validated test loglikelihood surface (right) vs. percentage of pixels which were correctly classified as being inside or outside of the true neuron S_{true} (left), as a function of (α_1, α_2) . In this case, the percent correctly classified is calculated as $\left(1 - \frac{\text{erroneously classified pixels}}{\text{number of pixels in } S_{true}}\right)$. The two surfaces match each other qualitatively, indicating that cross-validation is an effective technique for selecting (α_1, α_2) . **Bottom Left:** recovered shape $\hat{S}_{\alpha_1, \alpha_2}$ which maximizes the number of recovered pixels. White and black indicate pixels where the two agree, light grey denotes misses, and dark grey shows false positives. **Bottom Right:** $\hat{S}_{\alpha_1, \alpha_2}$ selected by using cross-validation to fit (α_1, α_2) . The two estimated shapes are comparable recoveries of the true shape, the respective percentages of erroneously classified pixels divided by number of pixels in S_{true} are 7.00% and 7.36%. The comparable fidelity of these recoveries again show that cross-validation is an effective technique for selecting (α_1, α_2) .

ters (α_1, α_2) which specify our prior $p(S) \propto e^{-Q(S)}$, and the states $\{S_j\}_{i \neq j}$ of every other pixel in S . It is well-known that this iterative sampling strategy leads to samples from

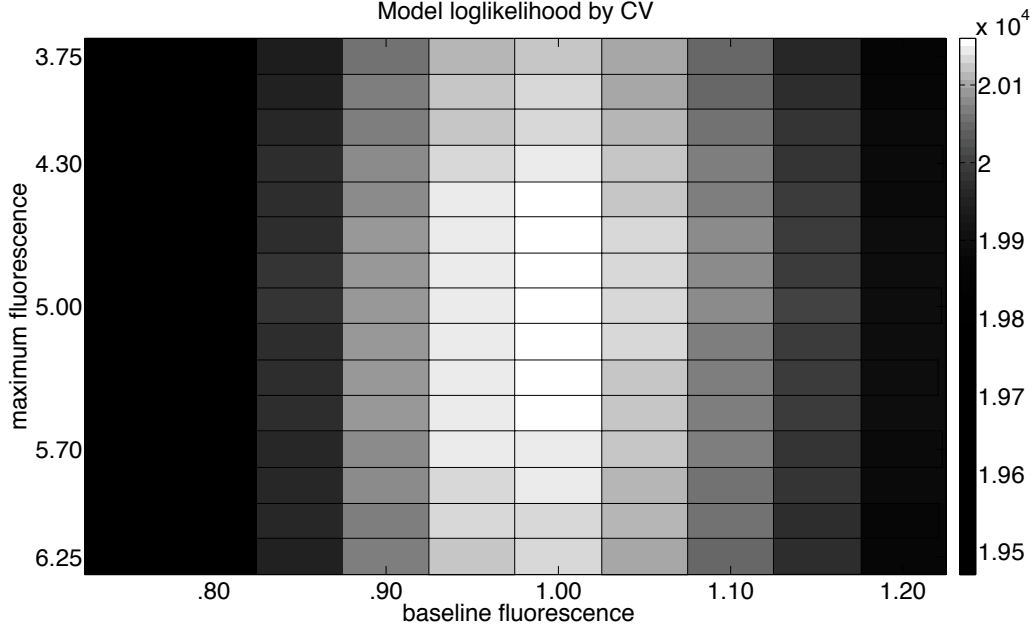


Figure 5: Cross validation can be used to simultaneously determine optimal penalty parameters (α_1, α_2) and fluorescence levels $(\lambda_{in}, \lambda_{out})$; here we show the loglikelihood surface as a function of λ_{in} and λ_{out} at optimal values of α_1, α_2 . As with Fig. 4, this figure illustrates performance using same-SNR data ($\lambda_{in} = 5\lambda_{out}$) as in the middle panel of Figure 2. λ_{in} vary respectively from .75 to 1.25 times and .75 to 1.20 times the true values used to generate the data. The maximal loglikelihood value exactly coincides with fluorescence levels used for data generation. Note that our estimate for λ_{out} is more strongly constrained by the data, since there are many more pixels (i.e., more information) outside the dendrite than inside.

the full posterior $p(S|I_{obs}, \alpha_1, \alpha_2)$, although these samples will not be independent; see, e.g., [Robert and Casella, 2005] for further details and background on Gibbs sampling.

Thus we only need to discuss the problem of sampling from $p(S_i|I_{obs}, \alpha_1, \alpha_2, \{S_j\}_{i \neq j})$; this turns out to be quite easy given the methods we have already developed. We simply write out

$$\begin{aligned}
 p(S_i|I_{obs}, \alpha_1, \alpha_2, \{S_j\}_{i \neq j}) &\propto p(S_i, I_{obs}, \{S_j\}_{i \neq j} | \alpha_1, \alpha_2) \\
 &= p(S_i, \{S_j\}_{i \neq j} | \alpha_1, \alpha_2) p(I_{obs} | S_i, \{S_j\}_{i \neq j}) \\
 &= 1(S \text{ is simply connected}) e^{-Q(S_i, \{S_j\}_{i \neq j})} p(I_{obs} | S_i, \{S_j\}_{i \neq j}),
 \end{aligned}$$

where the first term in the last line is one or zero depending on whether the shape S made

of the pixels $(S_i, \{S_j\}_{i \neq j})$ is simply connected. Luckily, we have already developed methods for efficiently checking simple connectivity and for computing and updating the penalty $Q(S)$ and the likelihood $p(I_{obs}|S)$, one pixel S_i at a time. Thus we see that sampling from the posterior $p(S|I_{obs})$ via the Gibbs approach requires no more code than what we have already written to maximize $p(S|I_{obs})$.

The behavior of the Gibbs sampler is illustrated in Figs. 6 and 7 (code to be made publicly available at <http://www.stat.columbia.edu/~liam>). With very low-SNR, image recovery is not particularly constrained by the data. Figure 6 shows the results of unpenalized sampling for increasing SNR (this is equivalent to sampling from the posterior given a “flat” prior, i.e., $p(S) = 1$ for all simply connected S). At the lowest SNR $\log p(I_{obs}|S)$ displayed (top row), samples from the unpenalized Markov chain shows high variability; in fact the Gibbs chain does not equilibrate around S_{true} . This analysis therefore indicates that the chain is relatively unconstrained by the data. If we increase the strength of the penalty, by raising α_1 and α_2 , we can constrain image recovery more effectively (Fig. 7). The top row shows what happens when we consider penalized Gibbs sampling for the same lowest-SNR data. In this case, $\log p(I_{obs}|S)$ equilibrates fairly quickly. Furthermore, sample shapes show the effect of the constraint and are much less variable than without a penalty. Thus, given appropriate prior information on dendritic shape, we can decrease the uncertainty in our recovered image, as expected. (In both figures, top rows have a ratio $\lambda_{in}:\lambda_{out}$ of 2:1; the penalized Gibbs chain equilibrates in about $5e3$ steps, or 45 seconds. All computations performed on a Pentium 4 CPU 3.40GHz laptop with 1GB RAM).

If we increase SNR instead of altering the penalty, the Gibbs sampler also covers nicely to a stationary state and $\log p(I_{obs}|S)$ stabilizes quickly. The bottom row of Fig. 6 shows the lower sample variability and the better approximation of the true shape, indicating that the data adequately constrains recovery. Combined with the penalized sampling results, the results of increased SNR indicate that the Gibbs sampler can be used as a valuable tool to determine whether a dendritic image is in fact adequately constrained by the observed data. (The higher SNR samples with ratios $\lambda_{in}:\lambda_{out}$ of 5:1 and 500:1 equilibrate after around $7.5e3$ steps, or 30 seconds, with or without the penalty; as in the case of lowest SNR we considered,

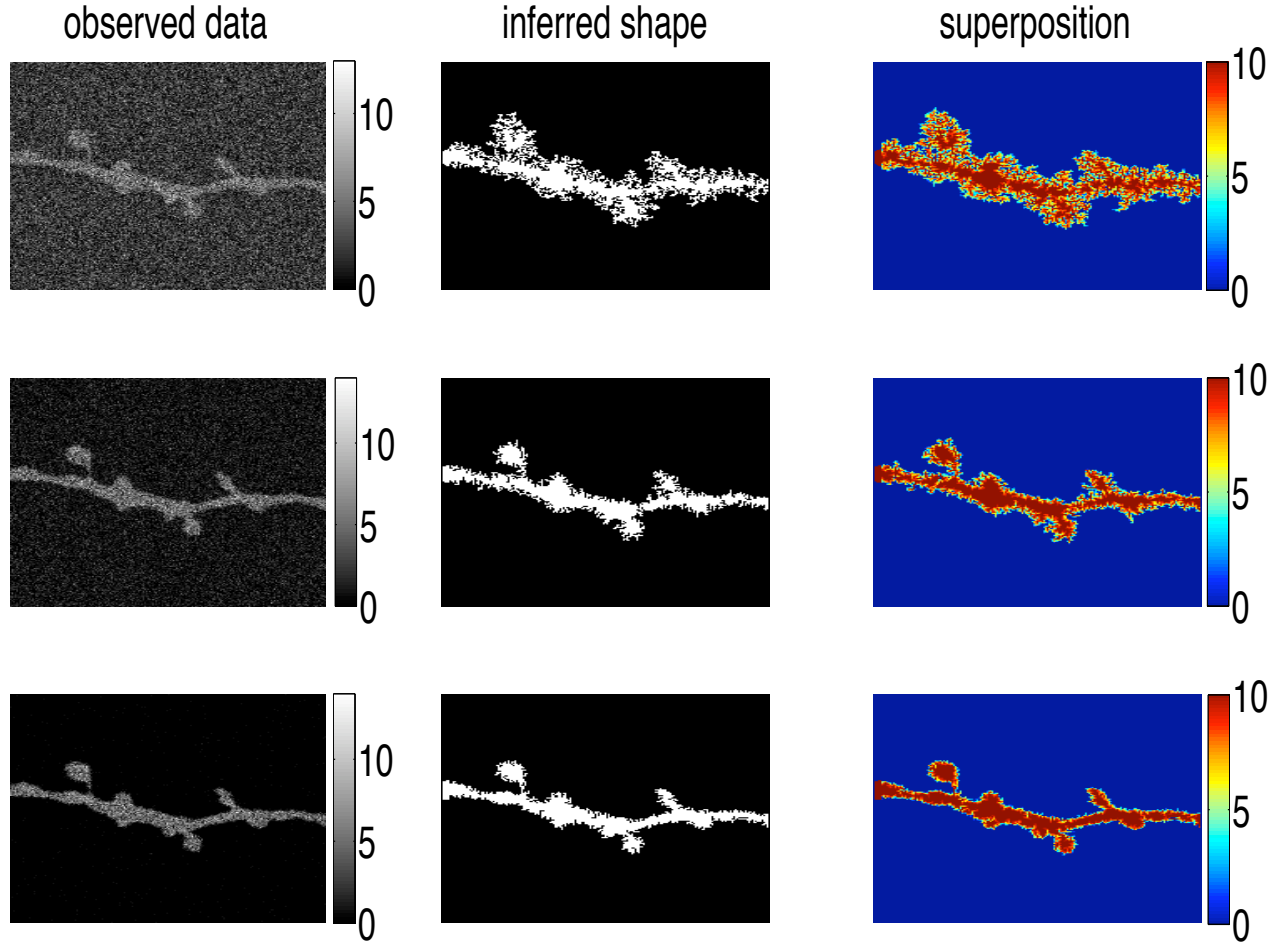


Figure 6: This figure displays how unpenalized MCMC sampling scales with increasing SNR. From the center and right columns, we gain a sense of the variability of samples, which decreases with increasing SNR. **Left Column:** I_{obs} (simulated) generated by the blurred Poisson model (same data as in fig 2). **Center Column:** a random sample S drawn during Gibbs sampling. **Right Column** pictures which show average shape of samples drawn using Gibbs sampling. Each image represents the square root of the superposition of 100 samples. Colorbar represents the percentage of samples that have a given pixel on the interior, red(10) means the pixel was on the neuronal interior on every sample, blue(0) means the pixel was on the exterior in every sample. Samples were drawn once every 200 proposals from the Markov chain after an initialization (burn-in) period of $5e3$ proposals.

sample shape and variability are visibly affected by the constraint).

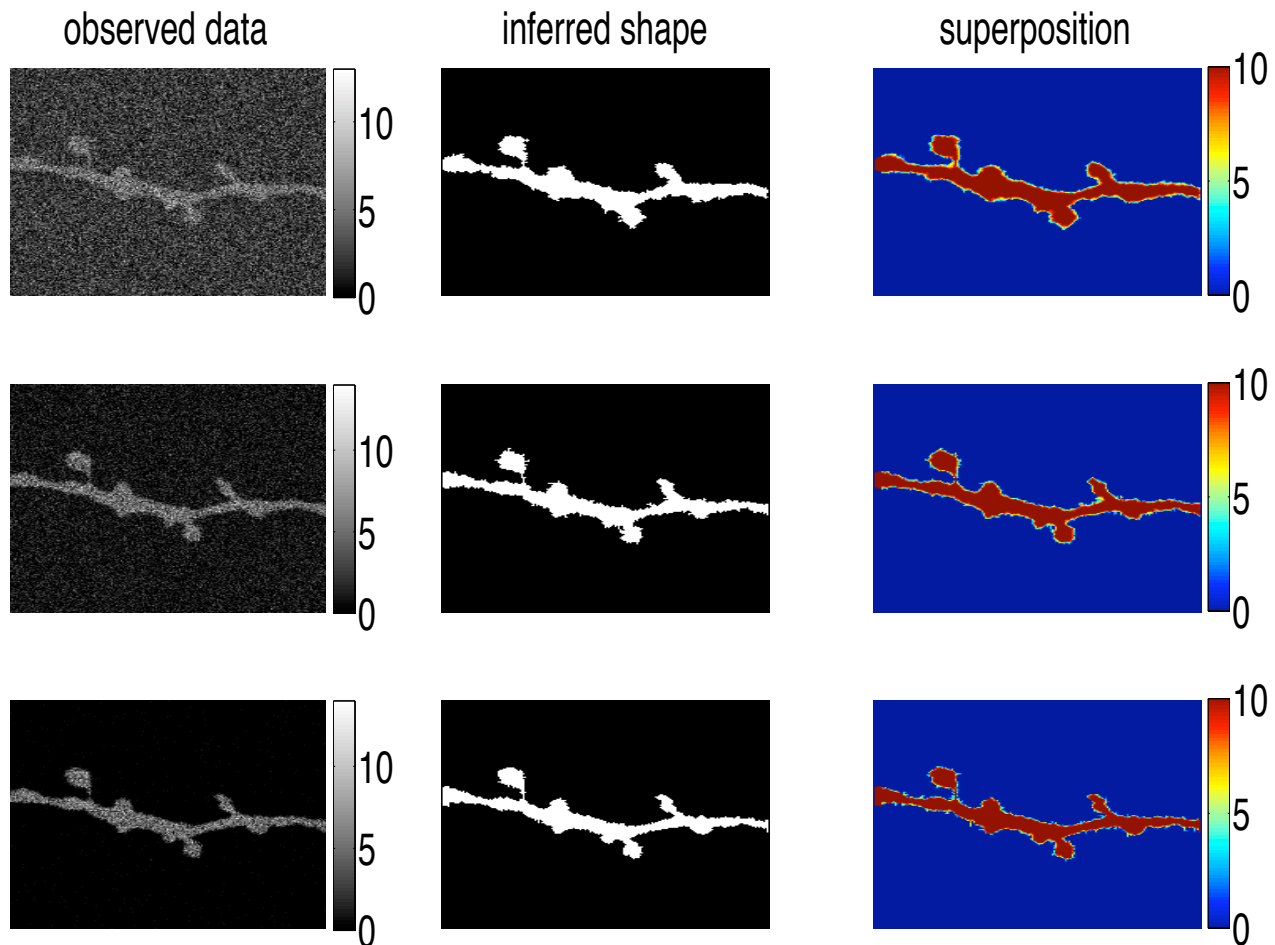


Figure 7: This figure shows how penalized MCMC sampling scales with increasing SNR. From the center and right columns, we see that imposing a penalty term ($\alpha_1=.2, \alpha_2=2$ here) constrains the variability of samples drawn using the Gibbs sampler. **Left Column:** I_{obs} (simulated) generated by the blurred Poisson model (same data as in fig 2). **Center Column:** a random sample S drawn during Gibbs sampling. **Right Column** pictures which show average shape of samples drawn using Gibbs sampling. Conventions as in Fig. 6. Note that penalization leads to much less variability- especially at low SNR- than is evident in Fig. 6.

7 Discussion

We have introduced Bayesian methods for quantifying neuronal shape given low-SNR image observations. Our key insights are that: 1) neurons are simply connected topological structures, and this geometric constraint may be easily incorporated in algorithms for determining neural shape; 2) direct maximum likelihood estimation leads to poor recovery of

the neuronal shape, while simple penalization methods perform much more effectively; and 3) Markov chain Monte Carlo techniques allow us to quantify the uncertainty in our estimates of the neuronal shape and can be used to determine whether this shape is actually constrained by the data. These aspects of our method make it a potentially useful front end for software whose goals include identifying spines or tracking neuronal branching, such as NeuronStudio (Computational Neurobiology and Imaging Center, Mount Sinai School of Medicine; <http://www.mssm.edu/cnic/tools.html>) or IMARIS (Bitplane AG, Zurich, Switzerland; <http://www.bitplane.com/go/products/imaris>).

Other groups have examined priors on roughness for imaging work. Good’s roughness penalty [Good and Gaskins, 1971] has received attention for use in biological imaging applications including tomography [Miller and Roysam, 1991], confocal microscopy [Verveer and Jovin, 1998] and optical sectioning microscopy [Joshi and Miller, 1993]. Good’s penalty is a kind of weighted Laplacian penalty; there are close connections between the edge-based penalty introduced here and these Laplacian penalties, since the Laplacian serves to effectively detect edges in the case of the binary shapes S of interest here.

A major direction for future work is to incorporate more “realism” into our priors. One attractive possibility has recently been developed by [Ascoli et al., 2001]: the idea is to develop priors which are truly “neuronal,” instead of the simple edge-based priors we have used here. This could potentially lead to much more accurate recovery of the underlying neuronal shapes. In particular, we would like to use different priors for different types of neuronal types, which might differ in the roughness of the dendritic membrane, the number, size, and shape of dendritic spines, etc. The challenge here will be to incorporate this more detailed prior knowledge of neuronal shape while maintaining the local nature of the computations described here. Future work could also attempt to relax the assumption that neuronal fluorescence can be modeled as a step function, and account for subtleties related to dye distribution in spines and dendrites. Another important track involves translating our 2-dimensional work into 3-dimensions. Most critically, we must adapt methods from digital topology to efficiently enforce the more complicated 3-d simple connectivity constraint [Lohou and Bertrand, 2004, Kong and Rosenfeld, 1989].

Finally, we should emphasize that the Gibbs sampling methods described here, while simple to code and understand, are far from optimal in terms of computational speed. In fact, the problem of efficiently sampling from a binary Markov random field remains a topic of ongoing interest in applied statistics, computer vision, and statistical physics. This body of work has a great potential for adaptation to our particular circumstances. Instead of flipping a single graph vertex, ideas like graph clustering and graph relabeling can be used to split, merge, and regroup chunks of the graph [Boykov et al., 2001]. This drastically speeds convergence when adjacent graph vertices are strongly coupled [Barbu and Zhu, 2006], as they are for neuronal structures. We hope in the future to apply these advanced sampling techniques to the analysis of spine geometry in real neuronal data.

Appendix: initialization of S

In the main text we chose not to give full details of the initialization of our estimate of the neuronal shape S , since our maximization and sampling methods may be initialized with any algorithm that extracts starting shapes with the correct topology from a noisy image. At low-SNR, generating topologically faithful initializations is itself a difficult problem which this particular paper does not aim to solve in full generality. Other authors consider this difficulty; for example, [Cheng et al., 2007] presents a method employing adaptive thresholding which efficiently determines local dendritic morphology. Furthermore, better initializations for S can reduce computation time and increase the accuracy of image recovery; thus, for completeness, we describe in this appendix the method we used to extract an initial S from the data shown in this paper (Fig. 8).

To obtain the ‘true’ neuronal shape analyzed here, we first thresholded the high-SNR empirical data shown in the top left of Fig. 1, then extracted the largest connected shape, filled in the resulting small observed holes, and smoothed the edges. To extract largest connected components, we first use Matlab’s `bwlabel.m` function to determine connectivity, and then `regionprops.m` to label connected areas; to fill in holes, we used Matlab’s `imfill.m` function. After using this shape to generate observed data I_{obs} according to our model, our procedure for initialization relies upon thresholding a smoothed version of this observed




Figure 8: This figure shows how we obtain a starting S from the observed data I_{obs} . **Left:** observed data as in fig. 2 (with $\lambda_{in} = 5\lambda_{out}$). **Center:** smoothed data, η_{data} . **Right:** $S_{threshold}$, obtained by thresholding η_{data} .

data. To choose the threshold value, we maximize the Poisson likelihood over a small set of candidates $S^{(k)}$. We generate candidate $S^{(k)}$ as follows: we first convolve the data, $n(x, y)$, with the psf to obtain a smoothed image η_{data} . Next, we take a number of thresholds of η_{data} , i.e. $\eta_{data} > \gamma * \text{mean}(\eta_{data})$ where we vary γ . Each threshold produces a binary image (with all pixels greater than the threshold set to one). We then extract the largest simply connected component, $S_{threshold}$. We choose the $S_{threshold}$ maximizing the Poisson loglikelihood according to our image degradation model, $\log p(I_{obs}|S)$, as our initial neuronal shape.

If the observed pixel counts are particularly sparse, and the neuron is particularly thin, this simple procedure can choose a poor starting point for our optimization. Sometimes this even leads to an initial guess with an incorrectly truncated neuron (figure 9, top right). As the average number of observed photons falling within an area of the neuron the size of the point-spread function drops to zero, the optimization easily becomes stuck in a local minima: the iterative algorithm cannot “see” enough data to flip the pixels that would move towards a globally optimal solution, and therefore cannot recover from a faulty truncation. In this very low-SNR case, we have found that a coarse-to-fine strategy is effective: we perform the above thresholding procedure after smoothing with an unphysically wide psf. This leads to much more reasonable initializations for the algorithm in these very low-intensity cases (fig. 9, bottom right).

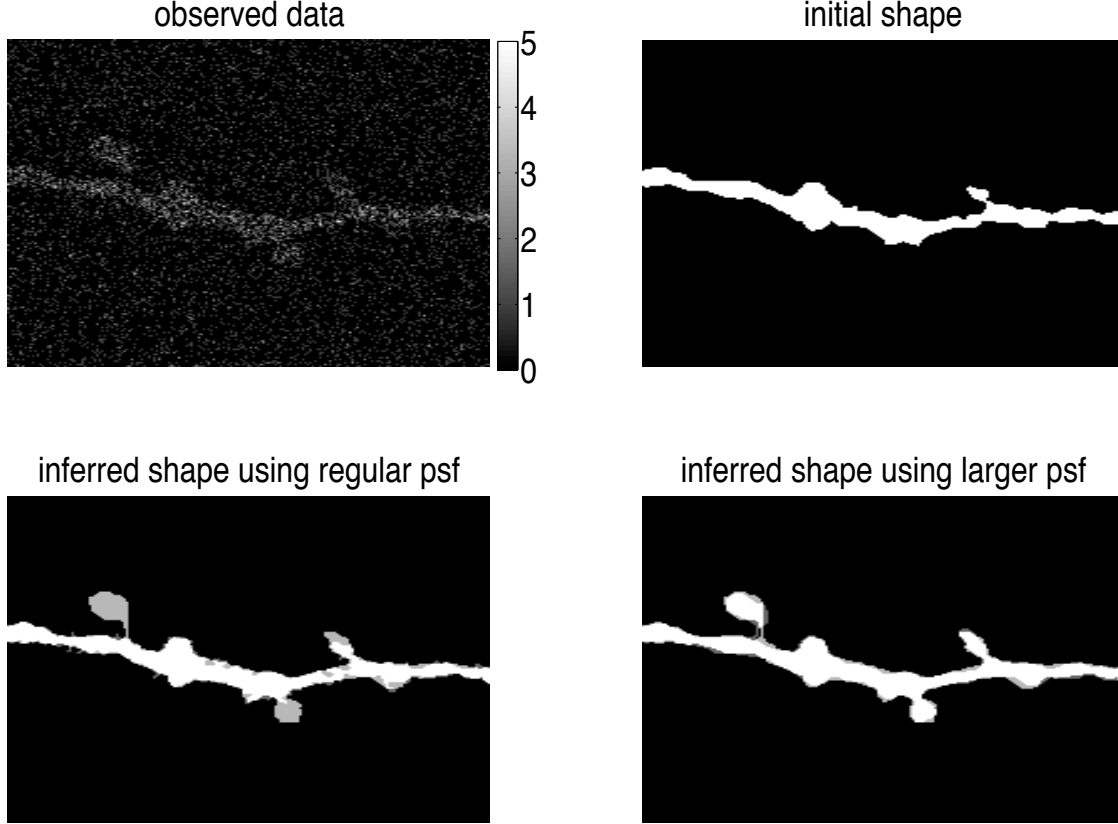


Figure 9: Demonstration of the utility of prior optimization with an unphysically large psf for recovering images with even lower SNR. **Top Left:** observed data using ($\lambda_{in} = 5\lambda_{out}$) but with λ_{in} at $\frac{1}{6}$ of its value in fig. 2 etc. **Top Right:** starting shape from thresholding. Poor starting points, such as this one, become increasingly common as SNR decreases and the data becomes more sparse. Optimization using an unphysically large psf can improve image recovery. **Bottom Left:** Comparison between \hat{S}_{MLE} and S_{true} after using the ‘physical’ psf, w , used to generate data. White and black indicate pixels where the two agree, light grey denotes misses, and dark grey shows false positives. **Bottom Right:** comparison between \hat{S}_{MLE} and S_{true} after prior optimization with an unphysically large psf (4x spread). In the second case, the recovery is clearly more successful and \hat{S}_{MLE} much better approximates S_{true} . In particular, note how the top left spine, which recovery using the ‘physical’ psf misses entirely, is present after optimization with the larger psf.

Acknowledgments

We thank R. Araya and R. Yuste for providing the example neural image used here, and for helpful conversations about fluorescence imaging. GF is supported by a Rabi Scholarship and

a Goldwater Scholarship; LP is supported by an NSF CAREER award and by an Alfred P. Sloan Research Fellowship.

References

- [Araya et al., 2006] Araya, R., Jiang, J., Eiselthal, K. B., and Yuste, R. (2006). The spine neck filters membrane potentials. *PNAS*, 103(47):17961–17966.
- [Ascoli et al., 2001] Ascoli, G., Krichmar, J., Nasuto, S., and Senft, S. (2001). Generation, description and storage of dendritic morphology data. *Philosophical Transactions: Biological Sciences*, 356:1131–1145.
- [Barbu and Zhu, 2006] Barbu, A. and Zhu, S.-C. (2006). Generalizing Swendsen-Wang to sampling arbitrary posterior probabilities. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 27(8):1239–1253.
- [Boykov et al., 2001] Boykov, Y., Veksler, O., and Zabih, R. (2001). Fast approximate energy minimization via graph cuts. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 23(11):1222–1239.
- [Cheng et al., 2007] Cheng, J., Zhou, X., Miller, E., Witt, R., Zhu, J., Sabatini, B., and Wong, S. (2007). A novel computational approach for automatic dendrite spines detection in two-photon laser scan microscopy. *Journal of Neuroscience Methods*, 165:122–134.
- [Conchello and McNally, 1996] Conchello, J. and McNally, J. (1996). Fast regularization technique for expectation maximization algorithm for optical sectioning microscopy. *SPIE Proceedings*, 2655:199–208.
- [Denk et al., 1994] Denk, W., Delaney, K., Gelperin, A., Kleinfeld, D., Strowbridge, B., Tank, D., and Yuste, R. (1994). Anatomical and functional imaging of neurons using 2-photon laser scanning microscopy. *Neuroscience Methods*, 54(2):151–162.
- [Engert and Bonhoeffer, 1999] Engert, F. and Bonhoeffer, T. (1999). Dendritic spine changes associated with hippocampal long-term synaptic plasticity. *Nature*, 399:66–70.

- [Farivar, 2005] Farivar, S. (2005). *Cytoarchitecture of the locust olfactory system*. PhD thesis, Caltech.
- [Geman and Geman, 1984] Geman, S. and Geman, D. (1984). Stochastic relaxation, Gibbs distribution, and the Bayesian restoration of images. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 6:721–741.
- [Good and Gaskins, 1971] Good, I. J. and Gaskins, R. A. (1971). Nonparametric Roughness Penalties for Probability Densities. *Biometrika*, 58(2):255–277.
- [Joshi and Miller, 1993] Joshi, S. and Miller, M. I. (1993). Maximum a Posteriori estimation with Good’s Roughness for Optical Sectioning Microscopy. *Journal of the Optical Society of America A*, 10:1078–1085.
- [Koh et al., 2002] Koh, I. Y. Y., Lindquist, W. B., Zito, K., Nimchinsky, E. A., and Svoboda, K. (2002). An image analysis algorithm for dendritic spines. *Neural Comput.*, 14(6):1283–1310.
- [Kong and Rosenfeld, 1989] Kong, T. and Rosenfeld, A. (1989). Digital topology: introduction and survey. *Comput. Vision Graph. Image Process.*, 48(3):357–393.
- [Kontaxakis et al., 1999] Kontaxakis, G., Strauss, L., and Tzanakos, G. (1999). An efficient implementation of the iterative MLEM image reconstruction algorithm for PET on a Pentium PC platform. *Journal of Computing and Information Technology*, 7(2):153–163.
- [Lohou and Bertrand, 2004] Lohou, C. and Bertrand, G. (2004). A 3d 12-subiteration thinning algorithm based on p-simple points. *Discrete Appl. Math.*, 139(1-3):171–195.
- [Lucy, 1974] Lucy, L. B. (1974). An iterative technique for the rectification of observed distributions. *Astrophysical Journal*, 79:745–753.
- [Markham and Conchello, 2001] Markham, J. and Conchello, J. (2001). Fast maximum-likelihood image-restoration algorithms for three- dimensional fluorescence microscopy. *Journal of the Optical Society of America A*, 18:1062–1071.

- [Miller and Roysam, 1991] Miller, M. and Roysam, B. (1991). Bayesian Image Reconstruction for Emission Tomography Incorporating Good’s Roughness Prior on Massively Parallel Processors. *Proceedings of the National Academy of Sciences*, 88:3223–3227.
- [Molina, 1994] Molina, R. (1994). On the hierarchical bayesian approach to image restoration: Applications to astronomical images. *IEEE Trans. Pattern Anal. Mach. Intell.*, 16(11):1122–1128.
- [Nimchinsky et al., 2004] Nimchinsky, E. A., Yasuda, R., Oertner, T. G., and Svoboda, K. (2004). The number of glutamate receptors opened by synaptic stimulation in single hippocampal spines. *J Neurosci*, 24(8):2054–2064.
- [Richardson, 1972] Richardson, W. H. (1972). Bayesian-Based Iterative Method of Image Restoration. *Journal of the Optical Society of America*, 62(1):55–59.
- [Robert and Casella, 2005] Robert, C. and Casella, G. (2005). *Monte Carlo Statistical Methods*. Springer.
- [Rodriguez et al., 2008] Rodriguez, A., Ehlenberger, D. B., Dickstein, D. L., Hof, P. R., and Wearne, S. L. (2008). Automated three-dimensional detection and shape classification of dendritic spines from fluorescence microscopy images. *PLoS ONE*, 3:e1997.
- [Simoni et al., 2004] Simoni, A. D., Fernandes, F., and Edwards, F. A. (2004). Spines and dendrites cannot be assumed to distribute dye evenly. *Trends in Neurosciences*, 27:15–16.
- [Svoboda, 2004] Svoboda, K. (2004). Do spines and dendrites distribute dye evenly? *Trends in Neurosciences*, 27:445–446.
- [van Kempen et al., 1997] van Kempen, G., van Vliet, L., Verveer, P., and van der Voort, H. (1997). A quantitative comparison of image restoration methods for confocal microscopy. *Journal of Microscopy*, 185(3):354–365.
- [Vardi et al., 1985] Vardi, Y., Shepp, L. A., and Kaufman, L. (1985). A Statistical Model for Positron Emission Tomography. *Journal of the American Statistical Association*, 80(389):8–20.

- [Verveer and Jovin, 1998] Verveer, P. J. and Jovin, T. M. (1998). Image restoration based on Good’s roughness penalty with application to fluorescence microscopy. *Journal of the Optical Society of America A*, 15:1077–1083.
- [Xu and Wong, 2006] Xu, X. and Wong, S. (2006). Optical microscopic image processing of dendritic spines morphology. *IEEE Signal and Processing Magazine*, 23(4):132–135.