JOINT LOCALIZATION AND PARAMETER ESTIMATION FOR LOCALIZED CALCIUM RELEASE EVENTS IN VIDEO MICROSCOPY

Benjamin P. Olding and Patrick J. Wolfe

Division of Engineering and Applied Sciences
Department of Statistics, Harvard University
Harvard-MIT Division of Health Sciences & Technology
Oxford Street, Cambridge, MA 02138 USA
{olding, wolfe}@stat.harvard.edu

ABSTRACT

This paper addresses modeling and inference for localized calcium release events observed in cardiac muscle tissue known as sparks, a recently discovered and little-understood phenomenon thought to be central to cardiac pathophysiology and hence of much scientific interest. In contrast to earlier algorithms proposed for studying sparks observed via confocal line-scan data, we consider here a model for the spatiotemporal spark process in unpaced cardiac myocytes observed *in vitro* via fluorescence video microscopy. In particular, we focus on parametric modeling of the spark process, leading naturally to a framework for joint localization and estimation of spark events and their parameters in space and time, along with the ability to quantify the uncertainty of our resulting inference. We show that for simulated sparks derived from real micropatterned cardiac myocyte data, and for which ground truth is known, our method significantly outperforms a video-adapted version of the most popular spark localization method currently in use in the cardiac biophysics community. Importantly, the framework we propose extends naturally to other other applications of video microscopy, and also allows for subsequent refinement of both the models and fitting procedures described here.

Index Terms— Biomedical microscopy, Biomedical image processing, Biomedical signal detection.

1. INTRODUCTION

The release of Ca^{2+} plays a critical role in the functioning of striated muscle tissue. In particular, calcium release channels feature centrally in the electromechanics of cardiac muscle cells termed myocytes. So-called *calcium sparks*, or "localized discrete calcium release events" [1] are thought to underlie in some way the excitation-contraction coupling mechanism in heart muscle [2]. In turn, spontaneous calcium release has been reported to be associated with a form of cardiac arrhythmia known as atrial fibrillation [3], implying that the study of calcium sparks may shed light on unresolved issues in cardiac pathophysiology [4]. Therefore, it is of much scientific and medical interest to characterize the dynamics of the underlying spatiotemporal spark process through experimental observation.

In brief, the problem that we address here is as follows. First, a cardiac myocyte exhibiting sparks is observed through a video microscopy apparatus and recorded as a sequence of digital images. In these video data we observe a spark whose amplitude varies over the length of the sequence; within the sequence, we wish to localize the spark's peak and infer its parameters, while contending with the fact that the spark signal itself changes over space and time, as do local signal-to-noise ratios owing to noise processes inherent in the imaging modality as well as various biological phenomena.

Since the first explicit report of spark events in [2], the spark process is typically observed through confocal line-scan imaging of single cells, via relative levels of fluorescent Ca^{2+} indicators (fluo-3). On this basis, models for both sparks and the associated observation process have been constructed (see, e.g., [5, 6]) and algorithms devised and employed to detect spark events (see, e.g., [1, 7, 8]). Earlier work has focused on nonparametric methods for localization of sparks and collection of relevant statistics; for example, from a statistical perspective the popular method of [1] can be interpreted as a threshold-based likelihood ratio test for spark presence. Here, however, we consider a novel experimental scenario, and formulate a rigorous model-based approach that leads directly to inference regarding spark parameters in a manner that improves upon state-of-the-art methods in the cardiac biophysics literature.

2. PARAMETRIC SPARK MODEL

The work of [9] represents a recent study of sparks in cardiac muscle, and motivates our modeling approach. It focuses on characterizing the spark process in micropatterned myocytes (cardiac myocytes observed *in vitro* whose two-dimensional geometry has been controlled using specialized techniques). A central scientific question in the literature concerns the biomechanical cytoskeletal properties of these cells, as reflected through the spark process [10]; as such properties depend on the entirety of the cell, video microscopy techniques are employed in [9] in order to observe the entire cell over time. Cells are stained with calcium-sensitive fluorescent dye, excited, and then recorded via (for example) a video camera mounted on an inverted microscope. Hence, the generated data consist of a sequence of digital images corresponding to observed fluorescence amplitude as a function of space and time. For the purposes of the present comparison and to establish a baseline for future comparison, we work with simulated data based on these micropatterned cardiac myocyte data, for which ground truth is known.

2.1. Continuous Spark Model

We employ an eight-parameter model to describe the fluorescence amplitude of a spatiotemporal spark event. Let t represent time (in

The authors would like to thank Prof. Kit Parker and Drs. Mark Bray and Nick Geisse for many insightful discussions of the problem, and for the use of data from the Disease Biophysics Group at Harvard.

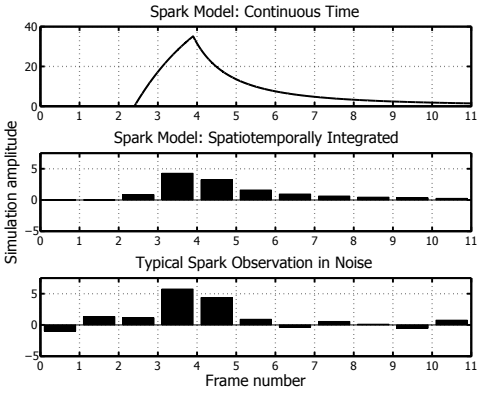

Fig. 1. Typical realization of the spark model of Section 2, with $t_0 = 3.9$, $A_0 = 35.142$, $\delta_t = 1.5$, $\tau_r = 2$, $\sigma_0 = 1.8$, and $\sigma_1 = 1$

units of video frame number), and x and y describe the spatial location of spark amplitude (in units of pixels). The work of previous researchers [1, 9] suggests that the spatial amplitude distribution of a spark is well modeled by the Gaussian form, and that this distribution is modulated by an exponential rise and decay in time. We have observed that a spatial diffusion process (as a function of time) is a better fit to collected data than the spatially homogeneous, temporal exponential decay model typically used in the literature. Following this observation, we define the spark event to have a spatially Gaussian profile that is radially symmetric; however, once the spark's exponential rise has completed, we allow the spatial extent of this profile to vary with time according to a spatial diffusion parameter $\sigma^2(t)$, with this variation parameterized as $\frac{1}{\sigma^2(t)} \exp(-\frac{x^2+y^2}{2\sigma^2(t)})$.

We model the initial growth in spark amplitude as an exponential gain process which starts from a specific point in time $t_0 - \delta_t$, and the subsequent decay in amplitude as strictly a function of the diffusion described above, which begins at time t_0; in other words, the amplitude of the spark decreases in time only as its spatial extent $\sigma(t)$ increases linearly, so that $\sigma(t)$ is defined as follows:

$$\sigma(t) = \begin{cases} \sigma_0, & \text{if } t_0 - \delta_t \le t < t_0, \\ \sigma_0 + (t - t_0)\sigma_1, & \text{if } t \ge t_0. \end{cases}$$

The complete model for a spark event $s(x, y, t)$ is thus given by

$$\begin{cases} A_0 \frac{1-e^{-\frac{t-(t_0-\delta_t)}{\tau_r}}}{1-e^{-\delta_t/\tau_r}} \exp\left(-\frac{(x-x_0)^2+(y-y_0)^2}{2\sigma_0^2}\right), & \text{if } t_0 - \delta_t \\ & \le t < t_0, \\ A_0 \frac{\sigma_0^2}{(\sigma_0+(t-t_0)\sigma_1)^2} \exp\left(-\frac{(x-x_0)^2+(y-y_0)^2}{2(\sigma_0+(t-t_0)\sigma_1)^2}\right), & \text{if } t \ge t_0, \end{cases}$$

where (x_0, y_0, t_0) is the spark peak in space and time, A_0 is its peak amplitude, δ_t is the time prior to the peak at which the spark event commences, τ_r is the growth rate of the subsequent exponential amplitude rise, σ_0 is the initial spatial extent of the spark, and σ_1 is the linear growth rate of the spark's spatial extent following its peak. The top plot in Figure 1 shows a typical realization of this continuous spark model (using parameters consistent with observed data), represented by an amplitude profile of the spatial center of the spark as a function of time.

2.2. Spark Model Observed via Digital Video

A digital imaging device sums light energy over time and space at each pixel's spatial location. Thus, our model for the spark event as observed via digital video is not a spatiotemporal *sampling* of the continuous spark model, but rather an *integration* defined as follows:

$$\begin{cases} A_0 \frac{t_2-t_1+\tau_r\left(e^{-\frac{t_2-(t_0-\delta_t)}{\tau_r}} - e^{-\frac{t_1-(t_0-\delta_t)}{\tau_r}}\right)}{1-e^{-\delta_t/\tau_r}} & \\ \quad \cdot \int_{y_1}^{y_2}\int_{x_1}^{x_2} \exp\left(-\frac{(x-x_0)^2+(y-y_0)^2}{2\sigma_0^2}\right) dx\, dy, & \text{if } t_0 - \delta_t \\ & \le t_1, t_2 < t_0, \\ A_0 \int_{t_1}^{t_2}\int_{y_1}^{y_2}\int_{x_1}^{x_2} \left[\frac{\sigma_0^2}{(\sigma_0+(t-t_0)\sigma_1)^2}\right. & \\ \quad \left. \cdot \exp\left(-\frac{(x-x_0)^2+(y-y_0)^2}{2(\sigma_0+(t-t_0)\sigma_1)^2}\right)\right] dx\, dy\, dt, & \text{if } t_1 \ge t_0, \end{cases}$$

where t_2 marks the end of an exposure frame and t_1 its beginning. Note that frames which include t_0 will be a combination of the above two equations, and frames which include $t_0 - \delta_t$ will obey the constraint $t_1 = t_0 - \delta_t$. With respect to the integration process, $\{x_1, x_2\}$ and $\{y_1, y_2\}$ describe the edges of the observed pixel in question.

As an example, the middle plot of Figure 1 shows the continuous spark event of the top plot, but following integration in both time and space as described above. The bottom plot shows a typical example of an observation of this event after white Gaussian noise of zero mean and unit variance has been added to the integrated spark signal.

3. SPARK LOCALIZATION AND ESTIMATION

To isolate the tasks of spark localization and parameter estimation for careful initial study, we reduce the problem at hand to that of estimating the eight model parameters $\{x_0, y_0, t_0, A_0, \delta_t, \tau_r, \sigma_0, \sigma_1\}$ describing a spark event from a video sequence whereupon (1) a spark is known to exist and (2) the temporal peak of the spark has been localized to within one exposure frame. We identify the null hypothesis H_0 as a noise model with mean zero, and the alternate hypothesis H_1 for a given set of parameters as the same noise model, but with a mean given by the amplitude under the integrated spark model. The localized maximum-likelihood estimation (MLE) algorithm we propose here starts by finding several likely spark candidates by selecting the least likely pixels under H_0. It improves these estimates by maximizing the likelihood of these candidates under H_1 using a hill-climbing algorithm (as in the maximization step of the well-known Expectation-Maximization algorithm). Finally, it selects the most likely candidate under H_1 as the localized spark event, on which it bases its estimates of the spark parameters:

Algorithm 1 Proposed MLE Spark Localization and Estimation

1. Given a model for the additive observation noise, calculate a one-sided p-value under H_0 for each observed pixel.
2. Create a "neighborhood" p-value by multiplying together the p-values of pixels neighboring each other in space and time (as pixel values are independent under H_0).
3. Locate the n smallest neighborhood p-values whose spatial and temporal extents do not overlap. Set these to be the n starting points for a spark peak in space and time.
4. For each such starting point of peak location, create conditional starting-point estimates of A_0, δ_t, τ_r, σ_0, and σ_1.
5. Apply a hill-climbing algorithm to find n local modes (i.e., the most likely set of parameters under H_1).
6. Compare the likelihood under H_1 of the n modes and declare the most likely result as the localized spark event.

For comparison purposes, we adapted a version of the threshold-based algorithm proposed in [1], the current *de facto* standard in the cardiac biophysics literature on calcium spark phenomena:

Algorithm 2 Threshold-Based Spark Localization [1]

1. For a given frame, identify all pixels exceeding a threshold T.
2. Merge identified pixels by considering all pixels within a distance k of each other to be members of the same "island."
3. Estimate the spark center as the centroid of the largest island.
4. Using the maximum value of this island, estimate spark spatial extent based on the recorded full-width half-maximum.

The authors of [1] recommend applying a threshold of $T = 3.5\sigma_n$ to a single frame of image data, where σ_n represents the standard deviation of the observation noise. When the noise model is identical for every pixel, this algorithm can be viewed as applying a likelihood ratio test to each pixel and selecting pixels for which the null is rejected under this test. These pixels are then merged into candidate "islands" of pixels, and the islands used to estimate spark location and spatial extent. Under the assumptions forming our basis for comparison, the correct frame to which to apply this test is known *a priori*; moreover, a spark is known to exist. Therefore, if no pixel is found that exceeds T, then the pixel with largest value is selected as the spark center.

4. SIMULATION AND RESULTS

To compare Algorithms 1 and 2 over the range of local signal-to-noise ratios (SNR) typically seen in practice, we simulated multiple trials consisting of sets of eleven video frames per trial, each having a resolution of 128×128 pixels. We used zero-mean white Gaussian noise of unit variance as our noise model, and simulated instances of the spark parameters depicted in Figure 1 while varying peak spark amplitude A_0. In keeping with [1], we took relative SNR to be the ratio of the peak integrated value of the spark model to that of the standard deviation of the noise. (Thus, for example, the SNR of Figure 1 is $20\log_{10}(4/1) = 12.0$ dB.) In this manner we varied A_0 while holding all other parameters constant, and simulated 16290 trials for each relative SNR in the set $\{-6, -5, \ldots, 12\}$.

For each trial, the peak location of each simulated spark was chosen uniformly at random in space (x, y), and the peak location in time t was chosen uniformly at random between frames three and four. Algorithm 1 (Local MLE) was applied to each set of frames, with the number of candidate neighborhoods set to $n = 5$ and a spatiotemporal neighborhood size of $4 \times 4 \times 4$ pixels; the resultant estimates of the eight model parameters $\{x_0, y_0, t_0, A_0, \delta_t, \tau_r, \sigma_0, \sigma_1\}$ were then recorded. Algorithm 2 (Threshold) was preliminarily applied using a variety of values for threshold T and merge distance k, based on which we observed $T = 3.5$ and $k = 1$ to give the best results (consistent with the recommendations of [1]).

For the purposes of evaluation, we define a spark to have been successfully localized if the estimated spark spatial peak location lies within σ_0 of the true spatial location (recalling that for these simulations, $\sigma_0 = 1.8$). Figure 2(a) compares the rates of successful spark localization for Algorithms 1 and 2 as a function of relative SNR, and Figure 2(b) compares the corresponding root mean squared error (RMSE) in estimated spatial peak location. Note that this value is a combination of variance in estimation and failure of localization, as incorrectly identified sparks will by definition have a large RMSE. Therefore, it is of interest to examine the RMSE of the estimated spatial position for just those sparks that were correctly localized, bearing in mind the rates of correct localization indicated in Figure 2(a) as a function of relative SNR for each of the two algorithms. In this vein, Figure 2(c) demonstrates that the location estimates of Algorithm 1 quickly improve to sub-pixel precision as the relative SNR increases. Finally, conditional upon a spark being correctly localized, Figure 2(d) compares estimates of spark spatial extent (with Algorithm 2's determination of full-width half-maximum converted into an estimate of σ_0 for comparison purposes). For each quantity of interest, an approximate 95% confidence interval was calculated based on the number of simulated observations in question; these intervals have been added to each subplot of Figure 2 in the form of error bars. The significantly larger error bars for low relative SNR levels in Figures 2(c) and 2(d) are a result of the limited number of successful spark localizations observed in these conditions.

As Algorithm 2 does not produce estimates of A_0, δ_t, σ_1, or τ_r, we have no reference for comparison. With the exception of τ_r, we have observed all estimates from Algorithm 1 to improve as a function of increasing relative SNR. The RMSE of δ_t decreased from 0.6 frames to 0.5 frames, that of A_0 held nearly constant at about 0.5 times the true value of A_0, and that of σ_1 decreased from around 1 to below 0.2. In contrast, the RMSE associated with τ_r actually increased. We believe that, given the short time period over which the simulated spark rose, there was not enough information available to correctly estimate τ_r. As the simulated spark parameters were consistent with our observations of experimental data relative to typical imaging rates, we suspect that Algorithm 1 may only be able to suggest a rough bound on τ_r in practice.

5. DISCUSSION

The simulation results described above clearly indicate the potential of a rigorous approach to spark localization. However, several important modeling and algorithmic extensions must be employed before arriving at an accurate, widely applicable model for video microscopy data of this kind. Most fundamentally, we intend next to relax the assumption that a spark exists and that its temporal peak can be localized to within one frame. This can be accomplished by extending the algorithm to fit sparks over time, and by incorporating formal model selection criteria to test for the existence of a spark. Useful extensions will also include enhanced robustness to nuisance processes, and more accurate noise modeling.

Additionally, observation models for actual imaging modalities are more complex than simple additive white Gaussian noise. At each pixel, the output of a typical imaging device is most accurately modeled as a Poisson process with a time-varying rate. We may then take the observed counts as the sum of the resultant process and an additive noise term, whose (time-invariant but spatially varying) mean models the so-called "fixed-pattern noise" due to variations in the baseline voltage at each photodiode location, and whose variance accounts for thermal noise in the device. Fixed-pattern noise may be estimated directly by taking data in the absence of any light input; based on a preliminary analysis we expect to be able to successfully apply the normal approximation to the Poisson distribution in this case. Doing so yields a model in which the time-varying rate is observed in the presence of additive noise whose variance depends in turn upon this rate.

Even before considering the extensions described above, we are able to draw two clear conclusions from the simulations presented in this paper: First, the superior performance of our proposed localization and estimation methodology can be explained by the fact that we formally "borrow strength" across regions of pixels to inform our

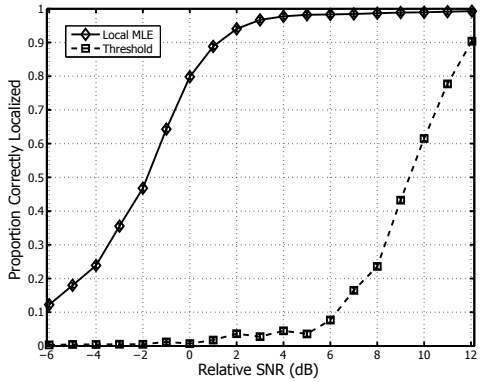

(a) Rate of successful spark localization (spatial peak estimation to within σ_0 of true location) vs. input SNR

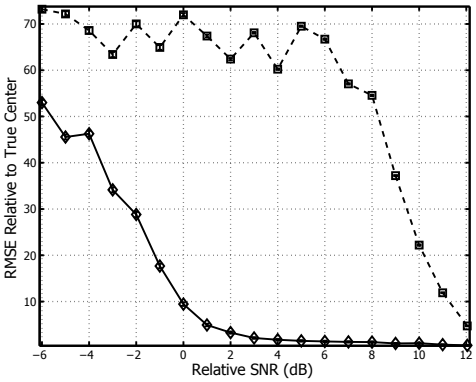

(b) Root mean squared error in estimated spatial peak location, also shown as a function of input SNR

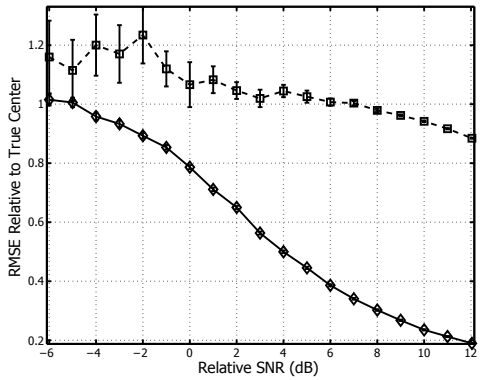

(c) Spark location error as in (b), but shown conditioned upon successful spark localization as defined in (a)

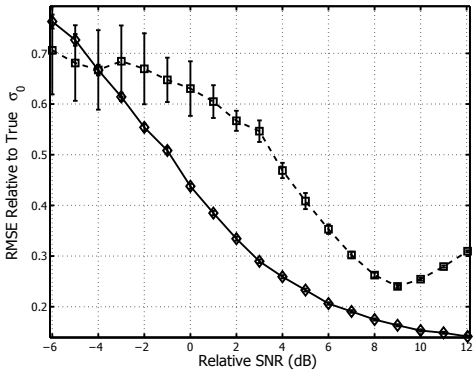

(d) Error in estimation of spark spatial extent $\sigma_0 = 1.8$, also shown conditioned upon successful localization

Fig. 2. Comparison of Algorithms 1 (Local MLE, solid line) and 2 (Threshold, dashed line) for spark localization and parameter estimation. Error bars (which often lie within the plotted squares and diamonds) represent approximate 95% confidence intervals for the values shown.

inference, in contrast to the threshold-based algorithm of [1]. Second, our framework of modeling and simulation not only enables us to quantitatively compare the performance of various algorithms, but will also allow for full uncertainty quantification in data analysis, through means such as confidence intervals for spark parameter estimates and formal hypothesis tests on spark process statistics.

6. REFERENCES

[1] H. Cheng, L.-S. Song, N. Shirokova, A. González, E. G. Lakatta, E. Ríos, and M. D. Stern, "Amplitude distribution of calcium sparks in confocal images: Theory and studies with an automatic detection method," *Biophys. J.*, vol. 76, pp. 606–617, 1999.

[2] H. Cheng, W. J. Lederer, and M. B. Cannell, "Calcium sparks: Elementary events underlying excitation-contraction coupling in heart muscle," *Science*, vol. 262, pp. 740–744, 1993.

[3] L. Hove-Madsen, A. Llach, A. Bayes-Gens, S. Roura, E. Rodriguez Font, A. Ars, and J. Cinca, "Atrial fibrillation is associated with increased spontaneous calcium release from the sarcoplasmic reticulum in human atrial myocytes," *Circulation*, pp. 1358–1363, 2004.

[4] M. Lindner, M. C. Brandt, H. Sauer, J. Hescheler, T. Böhle, and D. J. Beuckelmann, "Calcium sparks in human ventricular cardiomyocytes from patients with terminal heart failure," *Cell Calcium*, vol. 31, pp. 175–182, 2002.

[5] V. R. Pratusevich and C. W. Blake, "Factors shaping the confocal image of the calcium spark in cardiac muscle cells," *Biophys. J.*, vol. 71, pp. 2942–2957, 1996.

[6] G. D. Smith, J. E. Keizer, M. D. Stern, W. J. Lederer, and H. Cheng, "A simple numerical model of calcium spark formation and detection in cardiac myocytes," *Biophys. J.*, vol. 75, pp. 15–32, 1998.

[7] S. Sebille, A. Cantereau, C. Vandebrouck, H. Balghi, B. Constantin, G. Raymond, and C. Cognard, "Calcium sparks in muscle cells: Interactive procedures for automatic detection and measurements on line-scan confocal images series," *Comput. Programs Biomed.*, vol. 77, pp. 57–70, 2005.

[8] F. von Wegner, M. Both, and R. H. A. Fink, "Automated detection of elementary calcium release events using the *á trous* wavelet transform," *Biophys. J.*, vol. 90, pp. 2151–2163, 2006.

[9] M.-A. Bray, N. A. Geisse, and K. K. Parker, "Multidimensional detection and analysis of Ca^{2+} sparks in cardiac myocytes," *Biophys. J.*, vol. 92, pp. 4433–4443, 2007.

[10] S. C. Calaghana, J.-Y. Le Guennec, and E. White, "Cytoskeletal modulation of electrical and mechanical activity in cardiac myocytes," *Progr. Biophys. Mol. Biol.*, vol. 84, pp. 29–59, 2004.