

Photon Shot Noise Limits on Optical Detection of Neuronal Spikes and Estimation of Spike Timing

Brian A. Wilt,[†] James E. Fitzgerald,[†] and Mark J. Schnitzer^{†*}

[†]James H. Clark Center and ^{*}Howard Hughes Medical Institute, CNC Program, Stanford University, Stanford, California

ABSTRACT Optical approaches for tracking neural dynamics are of widespread interest, but a theoretical framework quantifying the physical limits of these techniques has been lacking. We formulate such a framework by using signal detection and estimation theory to obtain physical bounds on the detection of neural spikes and the estimation of their occurrence times as set by photon counting statistics (shot noise). These bounds are succinctly expressed via a discriminability index that depends on the kinetics of the optical indicator and the relative fluxes of signal and background photons. This approach facilitates quantitative evaluations of different indicators, detector technologies, and data analyses. Our treatment also provides optimal filtering techniques for optical detection of spikes. We compare various types of Ca^{2+} indicators and show that background photons are a chief impediment to voltage sensing. Thus, voltage indicators that change color in response to membrane depolarization may offer a key advantage over those that change intensity. We also examine fluorescence resonance energy transfer indicators and identify the regimes in which the widely used ratiometric analysis of signals is substantially suboptimal. Overall, by showing how different optical factors interact to affect signal quality, our treatment offers a valuable guide to experimental design and provides measures of confidence to assess optically extracted traces of neural activity.

INTRODUCTION

Neuroscientists are keenly interested in the use of fluorescent Ca^{2+} indicators for optical studies that track concurrently the dynamics of large populations of individual neurons (1). Such studies rely on neurons' voltage-gated Ca^{2+} channels, and detection of single action potentials is sometimes feasible (2–7). Yet, all optical techniques for detecting neural spikes are fundamentally limited by the statistics of photon emission and detection. The photomultiplier tubes commonly used in laser-scanning microscopy and advanced camera technologies that reduce electronic noise regularly permit imaging studies that approach the photon shot noise limit. In the live brain, this limit is often unattainable due to other sources of signal corruption, such as brain motion, but the shot noise limit still sets a performance bound on spike detection. Thus, it is crucial to understand the physical limits of optical methods to optimize experimental design, compare and improve imaging techniques, develop data analysis methods, and assess the reliability of empirical results (1,8).

Several studies have explored the signal-to-noise ratio (SNR) of fluorescent neural activity indicators (8–10), but the quantitative links between neural spike detection and optical properties of the indicator and instrumentation remain poorly defined. Metrics used to evaluate indicators and experiments vary widely (11–16), from single optical parameters (11,13) to heuristic combinations of indicator brightness, signal amplitude, and signal duration (12,14,16). These metrics do not capture the relationships among the indicator, microscope performance, and the statistics of

spike detection. Few publications describe spike detection as the statistical comparison of signal amplitudes to fluctuations of the background photon flux (9,14,16,17) or consider the accompanying issues regarding spike timing estimation (16,18). Thus, basic but important questions regarding how many photons are needed to reliably detect spikes and estimate their timing have remained unanswered. Filling this gap in the theoretical literature will resolve issues central to experimental work, such as how much improvement is needed to render optical voltage indicators capable of reporting single spikes in live mammalian brains.

Here we establish optical limits on the detection of neural spikes and estimation of spike timing using the statistical tools of signal detection theory and estimation theory. We use signal detection theory to frame spike identification as a binary classification problem and quantify the ability to distinguish action potentials from background and noise. Signal detection theory sets the optimal performance on this classification task and yields not only the optical detection limits but also optimal filtering methods that attain these limits. To study the inference of spike timing we use estimation theory, which provides statistical tools for describing the accuracy of quantitative measurements. The Chapman-Robbins lower bound from this theory (19,20) sets the minimum (error) variance of any unbiased means of estimating a parameter from noisy data, which we apply to the estimation of spike times.

When shot-noise-limited data are available, our treatment provides metrics to assess if algorithms to extract spikes from optical data are optimal. We illustrate this for dual-color studies using fluorescence resonant energy transfer (FRET) indicators. Most spike detection routines examine the ratio of the two color channels (21,22). However, one channel

Submitted December 27, 2011, and accepted for publication July 24, 2012.

*Correspondence: mschnitz@stanford.edu

Editor: Leslie Loew.

© 2013 by the Biophysical Society
0006-3495/13/01/0051/12 \$2.00

<http://dx.doi.org/10.1016/j.bpj.2012.07.058>

of low signal quality can corrupt the ratiometric analysis despite a high quality second channel. Our theory identifies the regimes in which ratiometric analysis approaches optimality and provides superior methods when it does not.

RESULTS

Statistical analysis of neuronal spike detection

We first quantify how identification of neural spikes in optical recordings is limited by photon shot noise. Our model for a spike's optical waveform is based on fluorescent Ca^{2+} indicators, but our framework also applies to voltage indicators. The typical Ca^{2+} transient in response to an action potential is a rapid rise in $[\text{Ca}^{2+}]$ followed by an approximately exponential decay back to baseline (23). The decay time reflects the cell's Ca^{2+} extrusion and buffering processes as well as Ca^{2+} unbinding from the indicator. To simplify our description of the indicator, we approximate intracellular $[\text{Ca}^{2+}]$ and the fluorescence signal as being linearly related. This ignores indicator saturation and is a good approximation when the intracellular $[\text{Ca}^{2+}]$ is much less than the indicator's Ca^{2+} binding constant. This approach is commonly used for neurons with temporally sparse patterns of spiking (24). We further assume fluorescence signals from distinct spikes add linearly and do not depend on the cell's spiking history or absolute $[\text{Ca}^{2+}]$. Neither all neuron types nor all Ca^{2+} indicators obey these assumptions (8).

We express the mean fluorescence signal, $S(t)$, after a spike at $t = 0$ as

$$S(t) = F_0 + Ae^{-t/\tau} \theta(t). \quad (1)$$

F_0 represents a time-independent background fluorescence rate and is the sum of: the rate of autofluorescence from unlabeled cellular structures; the rate of fluorescence from improperly targeted indicator molecules such as those in the extracellular space; and the rate of fluorescence from properly targeted indicators when the neuron is at its resting $[\text{Ca}^{2+}]$. Thus, F_0 specifies the rate of detected photons in the absence of a neural spike. A and τ represent the signal waveform's amplitude and decay time constant, respectively. The unit step function $\theta(t)$ equals zero if its argument is negative and is one otherwise.

We further assume the detected photon signals are digitized at sampling rate ν , so each time bin has duration $1/\nu$. When there is no action potential, the mean background photon flux, \bar{B} , is the integral of the first term in Eq. 1 over one time bin (shown below for the n^{th} time bin):

$$\bar{B} = \int_{(n-1)/\nu}^{n/\nu} F_0 dt = \frac{F_0}{\nu}. \quad (2)$$

\bar{B} is the mean of a Poisson process that combines fluorescence excitation, emission, and detection. When there is

a spike, the mean photon flux of the signal plus background, \bar{S}_n , is the time integral of both terms in Eq. 1:

$$\bar{S}_n = \int_{(n-1)/\nu}^{n/\nu} S(t) dt = \frac{F_0}{\nu} + A\tau(e^{1/(\tau\nu)} - 1)e^{-n/(\tau\nu)}. \quad (3)$$

\bar{S}_n is also the mean of a Poisson process, which unlike \bar{B} depends on n because the signal is time-varying.

We now consider how to use N successive samples from the detector, $\mathbf{F} = (F_1, F_2, \dots, F_N)$, to detect spikes. In the shot-noise-limited regime, \mathbf{F} is distributed according to Poisson statistics. We use this distribution to express two mutually exclusive hypotheses regarding the occurrence of a spike: the null hypothesis, $H^{(0)}$, which posits the absence of a spike; and the alternative, $H^{(1)}$, which posits that a spike occurred at time zero. Given a specific set of photon measurements, $\mathbf{F} = \mathbf{f}$, we compute the probability of attaining \mathbf{f} under each hypothesis by multiplying the independent detection probabilities for each time bin,

$$p_{\mathbf{F}}(\mathbf{f}|H^{(0)}) = \prod_{n=1}^N \text{Poisson}(\bar{B}) = \prod_{n=1}^N e^{-\bar{B}} \bar{B}^{f_n} / f_n! \quad (4)$$

$$p_{\mathbf{F}}(\mathbf{f}|H^{(1)}) = \prod_{n=1}^N \text{Poisson}(\bar{S}_n) = \prod_{n=1}^N e^{-\bar{S}_n} \bar{S}_n^{f_n} / f_n!. \quad (5)$$

We now use signal detection theory to optimally classify \mathbf{f} as an instantiation of either $H^{(0)}$ or $H^{(1)}$. This rests on interpreting $p_{\mathbf{F}}(\mathbf{f}|H^{(i)})$ as the likelihood of hypothesis $H^{(i)}$ being true given observation \mathbf{f} . Dividing Eq. 5 by Eq. 4 and taking the natural logarithm, the log-likelihood ratio is

$$L(\mathbf{f}) = \log \frac{p_{\mathbf{F}}(\mathbf{f}|H^{(1)})}{p_{\mathbf{F}}(\mathbf{f}|H^{(0)})} = \sum_{n=1}^N f_n \log \frac{\bar{S}_n}{\bar{B}} - \sum_{n=1}^N (\bar{S}_n - \bar{B}). \quad (6)$$

Intuitively, if $L(\mathbf{f})$ is greater than zero, it is more likely a spike occurred; if $L(\mathbf{f})$ is less than zero, it is more likely there was no spike. Formally, one uses a decision rule,

$$\delta(\mathbf{f}) = \begin{cases} H^{(1)} & \text{if } L(\mathbf{f}) > \log C \\ H^{(0)} & \text{if } L(\mathbf{f}) < \log C \end{cases}, \quad (7)$$

to classify whether \mathbf{f} is an instantiation of a spike or not. The cutoff between the two hypotheses, $\log C$, is set by choosing a cost function, C_{ij} , which quantifies both the penalties for incorrect choices of $H^{(i)}$ when $H^{(i)}$ is true (for $i \neq j$) and the benefits of correct choices ($i = j$). For example, if one wanted to avoid false positives but could tolerate false negatives, one would choose $C_{01} < C_{10}$. One then minimizes the mean expected cost given the prior odds of each hypothesis. The cutoff between the two hypotheses that results is

$$\log C = \log \left[\frac{\pi_0}{\pi_1} \frac{(C_{10} - C_{00})}{(C_{01} - C_{11})} \right], \quad (8)$$

where π_i is the prior probability of hypothesis $H^{(i)}$ (25). Further, by using the log-likelihood ratio to classify \mathbf{f} , we can assess the statistical significance of each assignment (i.e., we can compute a p -value).

To examine which measurements are classified as spikes, we now make two approximations to Eq. 6:

First, we assume the total number of detected photons is sufficiently large to approximate the probability distribution of the log-likelihood ratio as a Gaussian function. This is a standard approximation to Poisson counting statistics. Importantly, the Gaussian distribution differs between the set of measurements in which there was a spike and the set in which there was not.

Second, we assume the background, F_0 , is much greater than the fluorescence increment during a spike, $A/F_0 \ll 1$. This is presently true for nearly all voltage indicators and most Ca^{2+} indicators; moreover, as we show later, when this approximation is invalid, spike detection becomes straightforward from the perspective of signal detection theory. With $A/F_0 \ll 1$, the dominant contribution to the shot noise comes from the background photon flux, so the standard deviation of $L(\mathbf{f})$ is independent of whether or not there is a spike, $\sigma_L^{(1)} \approx \sigma_L^{(0)} \approx \sigma_L$, where $\sigma_L^{(i)}$ denotes the standard deviation of the distribution of the log-likelihood ratio about its mean, $\mu_L^{(i)}$, under hypothesis $H^{(i)}$ (see Section S1 in the Supporting Material).

These two approximations allow us to introduce a discriminability index, $d' = (\mu_L^{(1)} - \mu_L^{(0)})/\sigma_L$, that is often used in signal detection theory to describe systems with Gaussian fluctuations (25) (Fig. 1 A).

Varying the decision cutoff yields a curve known as the receiver operating characteristic (ROC) (25), which describes the tradeoff between detection sensitivity and susceptibility to noise (Fig. 1 B). The area under the ROC curve is a measure of detection fidelity that is independent of the cutoff (Fig. 1 C): this area equals the probability of an ideal observer correctly identifying the spike in a choice between a pair of measurements, one of which contains a spike and one of which does not (25). Importantly, d' indexes the receiver operating characteristic for spike detection (25) (Fig. 1, B and C, and see Section S2 in the Supporting Material). If the sampling rate is sufficiently high that $\tau\nu > 1$, then

$$d' \approx \frac{\Delta F}{F} \sqrt{\frac{F_0 \tau}{2}}, \quad (9)$$

where $\Delta F/F \approx A/F_0$ is the amplitude of the optical transient (Fig. 1 D, and see Section S1 in the Supporting Material). The value of d' provides a simple metric of the difficulty of spike detection. Shot noise fluctuations in the background photon flux influence d' via the dependence of $\Delta F/F$ on F_0 . Note that in our treatment $\Delta F/F$ is not an intrinsic property

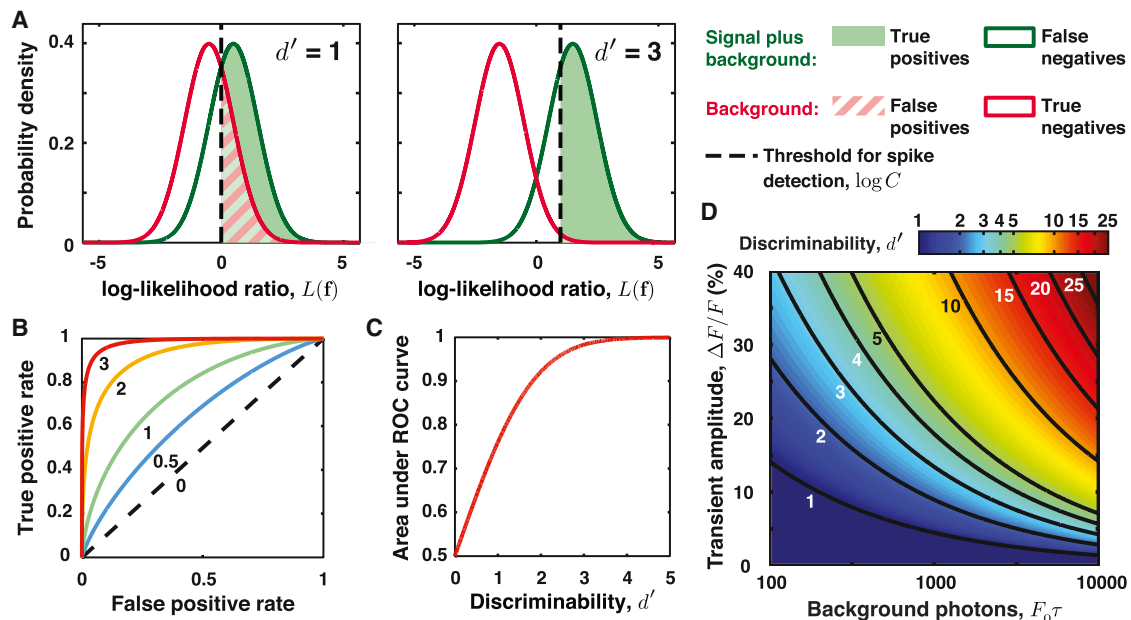


FIGURE 1 Signal detection theory. (A) Distribution of log-likelihood ratios, in cases with a spike (green curve), $p_L(\mathbf{f}|H^{(1)})$, and without a spike (red curve), $p_L(\mathbf{f}|H^{(0)})$, for $d' = 1$ and $d' = 3$ and two different detection thresholds. (B and C) The receiver operating characteristic (ROC) curve describes the trade off between the true positive rate versus the false positive rate for different values of the detection threshold. (B) In the approximation these distributions are Gaussians of equal variance, d' completely characterizes the ROC curve. For $d' \geq 5$, the ROC curve is visually indistinguishable from the panel axes. (C) The area under the ROC curve is a metric of spike detectability that is independent of the user's choice of detection threshold. This area equals the probability of an ideal observer correctly identifying the spike in a choice between a pair of instantiations of each hypothesis. (D) Contour plot showing how d' depends on the transient amplitude, $\Delta F/F$, and $F_0\tau$, the mean number of background photons collected over an interval of τ in duration. Contours are plotted for $d' = 1-5, 10, 15, 20$, and 25. Discriminability improves for increased F_0 at fixed $\Delta F/F$, due to the rise in signal photons.

of the indicator and the neuron's resting $[Ca^{2+}]$ but also depends on tissue autofluorescence and indicator labeling patterns. Any means of reducing background labeling thus directly benefits spike discriminability.

We verified Eq. 9 by simulating the detection of spikes in isolation (Fig. 2 A) and within spike trains (Fig. 2 B). We calculated the detection rate, P_D , and the false positive

rate, P_F , as a function of d' (see Section S2 in the Supporting Material) and used the equal-cost condition ($C_{10} = C_{01}$, $C_{00} = C_{11}$) to set the threshold for spike detection (Fig. 2 C). Although straightforward, this condition has some interesting consequences. For example, for $d' = 1$ we did not detect any spikes. Because of the low a priori probability of a time bin containing a spike, any drop in detection

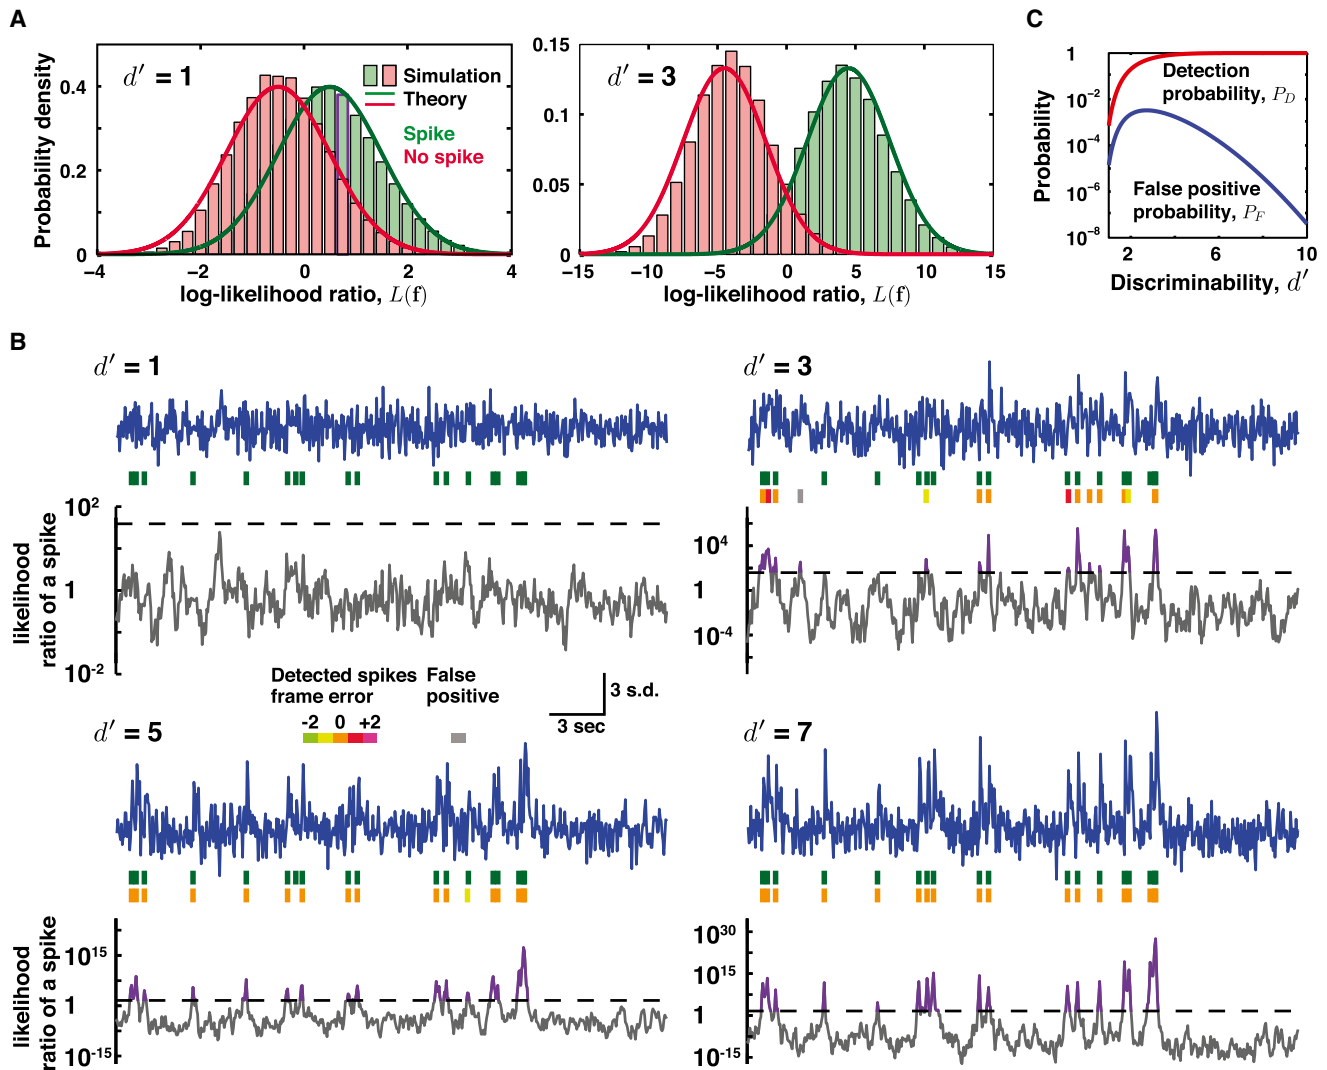


FIGURE 2 Simulations of spike detection. (A) Probability distributions of log-likelihood ratios, Eq. 6, taken over all possible photon measurements in the presence (green) and absence (red) of a neural spike, based on simulations of neural spiking with shot-noise-limited optical detection at $d' = 1$ or $d' = 3$. (Solid lines) Theoretically predicted Gaussian distributions. (Histograms) measured distributions from the simulation. The slight disagreements between the two are due to the use of leading-order approximations in the theory and can be remedied by including higher-order terms (see Section S1 in the Supporting Material). (B) Simulated optical traces and detected spikes for several values of d' . (Blue lines) Optical traces shown in units of the standard deviation from the mean photon count. (Green spikes) The true spike train. (Orange spikes) correctly estimated spikes. (Spikes in non orange hues) Spikes estimated with errors in frame timing. (Gray spikes) False positives. (Gray trace) For a moving window of nine time bins, the log-likelihood ratio for the hypotheses that there is ($H^{(1)}$) and is not ($H^{(0)}$) a spike. (Dashed black line) Spike detection threshold, $\log C$, given equal costs for false positives and false negatives (see Eq. 8). (Purple) Threshold crossings. Spikes were detected using an iterative, greedy algorithm that assigned a spike to the instance of the log-likelihood ratio's maximum in each iteration. At low d' , few spikes are detected with this choice of threshold (see panel C). Simulation parameters: $\Delta F/F = 0.05$, $\tau = 0.15$ s, $\nu = 20$ Hz, spike rate $\lambda = 0.5$ Hz. The cutoff is given by $\log C = \log((\nu/\lambda) - 1)$. For $d' = 1, 3, 5,$ and 7 , the probability of detection is $7.8 \times 10^{-4}, 0.61, 0.96,$ and 0.999 , respectively, whereas the expected number of false positives is $0.0092, 1.9, 0.36,$ and 0.017 , respectively. (C) Detection probability, P_D , and false positive probability, P_F , as a function of d' calculated using Eqs. S17 and S18 in the Supporting Material and the same equal-cost condition as in panel B. For low values of d' , the false positive probability rises with d' , because at $d' = 0$ there are very few spikes detected and thus very few false positives. At large d' values nearly all spikes are detected, and the false positive rate nearly vanishes.

threshold to improve sensitivity would be more than offset by a rise in false positives.

To estimate the spike trains in Fig. 2 B we used an iterative approach. We first initialized the estimated spike train with zero spikes. For each time point in the entire trace, we then evaluated the log-likelihood ratio for the occurrence of a spike. If this ratio ever exceeded $\log C$, we added a spike to the estimated train at the time for which the log-likelihood ratio was greatest. We then iterated this process, one spike at a time, by calculating the log-likelihood ratio between the estimated spike train (the null hypothesis that no more spikes are present) and the estimated train plus one additional spike. As in the first iteration, we considered the candidate additional spike at the time when the log-likelihood ratio was greatest. The algorithm stopped as soon as the log-likelihood ratio failed to exceed $\log C$ across the entire trace. In the terminology of optimization theory, this algorithm is called “greedy” because it makes the locally optimal decision at each iteration (26). This approach is simple and fast, but methods that can estimate globally optimal spike trains would clearly be superior (26).

Equation 6 defines a linear filter that can be used to estimate spike trains from optical data when spiking patterns are temporally sparse. In this regime, spike train estimation is well approximated by the detection of a series of isolated spikes (27). Thus, given the waveform of Eq. 1 subject to Poisson noise, Eq. 6 provides the optimal filter for spike train estimation. When the background photon flux is much greater than the signal flux, both Eq. 6 and the noted

Wiener filter reduce to a scaled version of what is termed a matched filter in the signal detection literature (28).

Robustness of the discriminability index

Having characterized spike detection using d' , we studied this metric's robustness to deviations from the simplest model. We first considered spikes that initiate incommensurate with the start of a time bin, as occurs in most experiments. The value of d' decreased by $<14\%$ across all sampling rates and by $<5\%$ at sampling rates $\nu > 2/\tau$ (see Section S3 in the Supporting Material). We also explored scenarios in which the neuron is examined during only a fraction of the data acquisition time bin. This arises in laser-scanning microscopy, as a single cell usually occupies a subset of a scanned line or frame. This can often be modeled by scaling the emission rates A and F_0 by the fraction of the time bin that the laser dwells on the cell (see Section S3 in the Supporting Material). We further studied how parameter misestimation affects spike detection. Misestimating τ by up to 50% only has an $\sim 11\%$ effect on d' (see Section S4 in the Supporting Material). Another study has also reported minimal effects from misestimating τ (24).

We also explored the low photon limit, for which the derivation of d' is invalid (Fig. 3). We simulated optical traces of spikes and then detected them using our greedy approach to spike detection. Though the log-likelihood ratios were not Gaussian-distributed, the area under the ROC curve was still well predicted by the Gaussian approximation (Fig. 3, A and

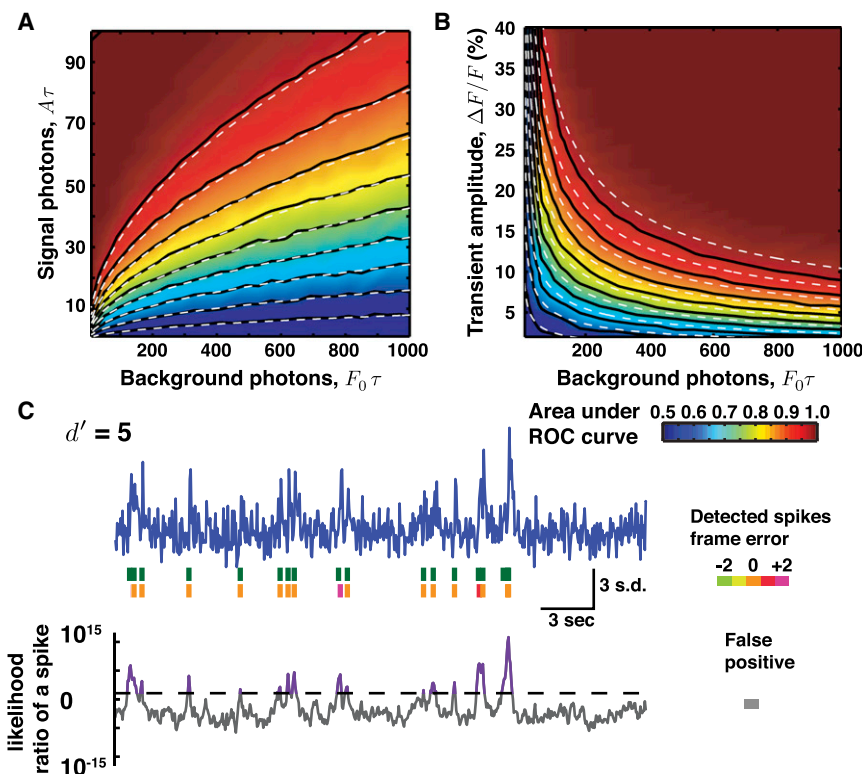


FIGURE 3 Spike detection with modest photon counts. When the number of photons collected from signal and background sources is insufficient to justify Gaussian approximations, the treatment involving d' is not guaranteed to be valid. To explore this regime, we randomly sampled the log-likelihood ratio distributions of Eq. 6, computed the ROC curves, and calculated the underlying areas. The area under the ROC curve is plotted as a joint function of background and signal photons, (A), or of background photons and $\Delta F/F$ (B). (Dashed lines) Areas under ROC curves determined using the d' calculation, demonstrating the robustness of the calculation result despite the invalidity of the Gaussian approximation. The contours in both plots are spaced in 0.05 increments of area under the ROC curve. (Note that unlike in Fig. 1 D, the x-axes are on a linear scale.) As expected, the agreement between the direct calculations and the d' approximation degrades as the condition $A/F_0 \sim \Delta F/F \ll 1$ is relaxed. This leads to a maximum disagreement of ~ 0.025 in the area under the ROC curve. (C) Simulated trace plotted using the same conventions as in Fig. 2. Simulation parameters: $\Delta F/F = 0.75$, $\tau = 0.15$ s, $\nu = 20$ Hz, spike rate $\lambda = 0.5$ Hz. Probability of detection is 0.96 and the expected number of false positives is 0.36.

B). Further, some studies with Ca^{2+} indicators may achieve signal amplitudes much greater than the background photon flux. These cases do not fit into the d' framework, which requires $A/F_0 \ll 1$. However, the simulations showed that if the signal amplitude is comparable to or greater than background fluorescence, the area under the ROC curve can attain values >0.99 (e.g., Fig. 3).

We also examined the issue that real indicators do not respond instantly to an action potential and have a finite on-time, τ_{on} . This reduces the signal amplitude and broadens its waveform (see Fig. S1, A and B, in the Supporting Material). For Ca^{2+} indicators, the on-time is typically brisker than the decay, $\tau_{on}/\tau < 0.1$, because the former is usually limited by Ca^{2+} binding to the indicator but the latter is limited by Ca^{2+} buffering, extrusion, and unbinding from the indicator (23). In this case, d' is reduced by at most $\sim 10\%$ (see Fig. S1 B). For voltage indicators the on- and off-times can be similar, necessitating a more careful treatment (see Section S5 in the Supporting Material).

Spike detection with FRET indicators

Most studies that use FRET indicators analyze the ratio of the two emission channels. This helps remove correlated noise, such as from motion artifacts or illumination fluctuations (21). But with these factors increasingly under experimental control, nonratiometric analyses may yield spike trains of higher fidelity (29).

We assume both channels are shot-noise-limited and obey Eq. 1 but with different signal amplitudes and F_0 values. Both channels share the same τ , reflecting the time course by which a Ca^{2+} transient alters the indicator's physical conformation and thus the FRET efficiency. We derive the noise covariance between the two channels to be zero in Section S6 in the Supporting Material. From these starting points, we calculated the discriminability of the log-likelihood ratio for a ratiometric analysis (d'_V) and for a direct analysis that treats the two channels separately (d'_W). A Gaussian approximation to the ratio of the two Poisson-distributed channels allowed us to determine d'_V (see Section S7 in the Supporting Material). We found d'_W by adding the d' for the two channels in quadrature.

This comparison revealed that although a ratiometric analysis cannot outperform the direct analysis, in many situations a ratiometric analysis is adequate. However, asymmetries between the two channels, such as different collection efficiencies or donor and acceptor fluorors of different brightness, exacerbate the deficits of ratiometric analyses (Fig. 4, A and B). We verified this using simulated traces (Fig. 4 C). Notably, our analysis reveals the ratiometric approach attains the optimum discriminability if the two channels have matching signal amplitudes, irrespective of the background photon fluxes (see Section S7 in the Supporting Material). There is also a regime in which it is better to discard a channel completely than to use a ratiometric analysis.

Our framework reveals the expected performance decrease associated with ratiometric analysis. For those struggling in the regimes highlighted above, a direct approach may facilitate extraction of higher quality spike trains. In particular, one can process each channel individually using Eq. 6 and sum the log-likelihood ratios from each channel. We did this to compare the traces for the direct approach to those based on a ratiometric analysis, and the superiority of the direct approach is apparent in Fig. 4 C.

Spike time estimation

For many neurons, spike times may encode information such as in the retina (30), hippocampus (31–33), visual cortex (34), and auditory system (35–37). Thus, studies of neural coding often benefit from precise determinations of spike timing.

We use estimation theory to compute a lower bound on the (error) variance of an unbiased estimate of a spike's occurrence time, t_0 , based on properties of the optical indicator and instrumentation. We use the Chapman-Robbins lower bound (19,20) to describe the optimal performance of any unbiased estimator of the spike time. Unbiased estimators are those that, on average, yield the true value. Although biased estimators may sometimes be preferable, we do not consider them here. Some readers may be familiar with the Cramer-Rao lower bound (19). The Chapman-Robbins bound is as small or smaller than the Cramer-Rao bound when they both exist (19,20) and is applicable to a wider class of models including the discretized measurements considered here.

We model the mean photon flux of the signal plus background for a spike at time t_0 as

$$\bar{S}_n^{(t_0)} = \begin{cases} F_0/\nu & ; n \leq t_0\nu \\ F_0/\nu + A\tau(1 - e^{-(t_0\nu - [t_0\nu])/(\tau\nu)}) & ; n = [t_0\nu] \\ F_0/\nu + A\tau(e^{1/(\tau\nu)} - 1)e^{-(n/\nu - t_0)/\tau} & ; n > [t_0\nu] \end{cases} \quad (10)$$

Here $[x]$ denotes the smallest integer not less than x and identifies the first frame containing signal photons. The Chapman-Robbins lower bound assures the variance of the estimated spike time, $\sigma_{i_0}^2$, satisfies

$$\sigma_{i_0}^2 \geq \max_{\Delta} \frac{1}{\mathbb{E}\{J|t_0\}}, \quad (11)$$

$$J = J(t_0, \Delta) = \frac{1}{\Delta^2} \left\{ \left[\frac{p_{\mathbf{F}}(\mathbf{f}; t_0 + \Delta)}{p_{\mathbf{F}}(\mathbf{f}; t_0)} \right]^2 - 1 \right\}, \quad (12)$$

where \hat{t}_0 denotes the estimated spike time, \mathbb{E} is the expectation operator over the distribution of \mathbf{F} , and \max_{Δ} represents the maximization over the dummy variable Δ . We restrict Δ

to a few τ , because spikes found beyond the duration of the signal transient are better interpreted as false positives than mistimed spikes. Similarly to Eqs. 4 and 5, the probability of a set of measurements $\mathbf{F} = \mathbf{f}$ is

$$p_{\mathbf{F}}(\mathbf{f}; t_0) = \prod_{n=1}^N \text{Poisson}(\bar{S}_n^{(t_0)}) = \prod_{n=1}^N e^{-\bar{S}_n^{(t_0)}} \left(\bar{S}_n^{(t_0)}\right)^{f_n} / f_n!. \quad (13)$$

For d' between 2 and 10, τ between 0.15 and 2 s, and ν between 8 and 40 Hz, the Chapman-Robbins lower bound varies by two orders of magnitude. Contour plots of this lower bound are shown in Fig. 5 A–C, for fixed values of τ , ν , or d' . Even experiments with high discriminability and sampling rates can suffer from poor spike timing. Importantly, the precision of spike time estimation can be poorer or finer than the time bin $1/\nu$ and depends most sharply on d' and τ . Super-resolution of spike times — i.e., timing of spike occurrences to better than the $2/\nu$ resolution suggested by the Nyquist sampling criterion — may be possible even when spike discriminability is low (Fig. 5, A–C). Intuitively, such super-resolution may be attainable by use of the prior knowledge about the shape of the spike's waveform (Eq. 1) for the estimation of its timing. However, estimation theory does not guarantee the existence of an estimator attaining the Chapman-Robbins lower bound.

Using an exhaustive likelihood maximization routine, we estimated the timing of simulated spikes across a range of d' and τ values. The estimated spike times were approximately distributed in a Laplace distribution centered on the true spike time,

$$p(\hat{t}_0 | t_0) = \frac{e^{-|\hat{t}_0 - t_0|/b}}{(2b)}.$$

To determine the errors in estimated spike times, we removed the contributions to the variance from the uniform background of the false positives and then performed a curve to fit a Laplace distribution (Fig. 5, D and E). Although the spike time estimates obeyed Eq. 11, for longer transients and lower d' values exhaustive likelihood maximization was unable to attain the Chapman-Robbins lower bound. The difficulty or impossibility of finding an estimator that attains this bound implies it should perhaps be considered the best-case for the estimation variance. Still, the Chapman-Robbins lower bound provides useful limits on the statistical confidence for spike times. When estimating the time interval between two spikes, the minimum possible variance is the sum of the minimum variances associated with each spike.

To improve spike timing resolution, researchers are pushing optical microscopy to higher sampling rates, for example by using acousto-optic laser-scanning methods (15,38–40). To examine the efficacy of such strategies, we tested how rapid sampling improves spike timing resolution. To take an example, the theoretically calculated Chapman-Robbins lower bound is 20.7 ms for $d' = 5$, $\tau = 0.15$ s, and

$\nu = 20$ Hz. Increasing the sampling rate to $\nu = 2000$ Hz, which also slightly increased d' , improved the lower bound to 2.8 ms. As above, the simulations did not attain the Chapman-Robbins lower bound and revealed only modest improvement in timing resolution from 19.9 ± 1.7 ms to 16.0 ± 1.3 ms at the faster sampling rate. However, more substantial gains in spike timing resolution are attainable using random-access laser-scanning methods to increase the laser dwell time per cell and hence d' (see Section S3 in the Supporting Material) (38,39).

DISCUSSION

Microscopists are familiar with how the numerical aperture, signal intensity, field of view, integration time, and background labeling affect photon counts, but knowledge of how photon detection relates to spike detectability and timing estimation has been missing. The metric of discriminability, d' , enables a quantitative treatment of these issues. Our derivation of d' is valid when background photon counts are larger than signal amplitudes and is consistent with the SNR heuristic proposed but not derived in Yasuda et al. (10). By relating the photon collection capacity to the parameters governing d' , experimentalists can compare specific instruments, fluorescent indicators, or experimental configurations for studies of neural spiking. Our framework quantifies how methods that improve fluorescence excitation (41,42) or detection (43) impact spike detection and timing estimation, even though the physical principles underlying these methods vary considerably.

Researchers use a variety of computational methods to extract neural spikes from optical traces (5,16–18,24,44,45). Provided the data are shot-noise-limited and that our model of $[\text{Ca}^{2+}]$ -related fluorescence and its noise fluctuations remains valid, our bounds on detection performance are applicable. However, our treatment assumes signals from individual cells have already been isolated from the raw data. Extracting individual cells' dynamics from within raw movie data can often be challenging in its own right (7). Some cell sorting methods produce dynamical traces that are weighted sums of the signals from different pixels (7), and these traces will generally not have Poisson-distributed noise fluctuations. Further, poor optical signals may limit the sorting process, causing an incorrect assignment of photons to cells or failure to remove optical crosstalk between cells. Such factors could prevent subsequent spike detection from achieving the limit expressed via d' .

When estimating trains of temporally sparse spikes, we used a filtering approach based on log-likelihood ratios, which has several advantages. Equation 6 follows directly from Eqs. 1–5 and does not require further approximations regarding the nature of the signal. As Eq. 6 does not assume Gaussian noise, it applies even in the low background limit in which use of d' is inappropriate. Equation 6 readily

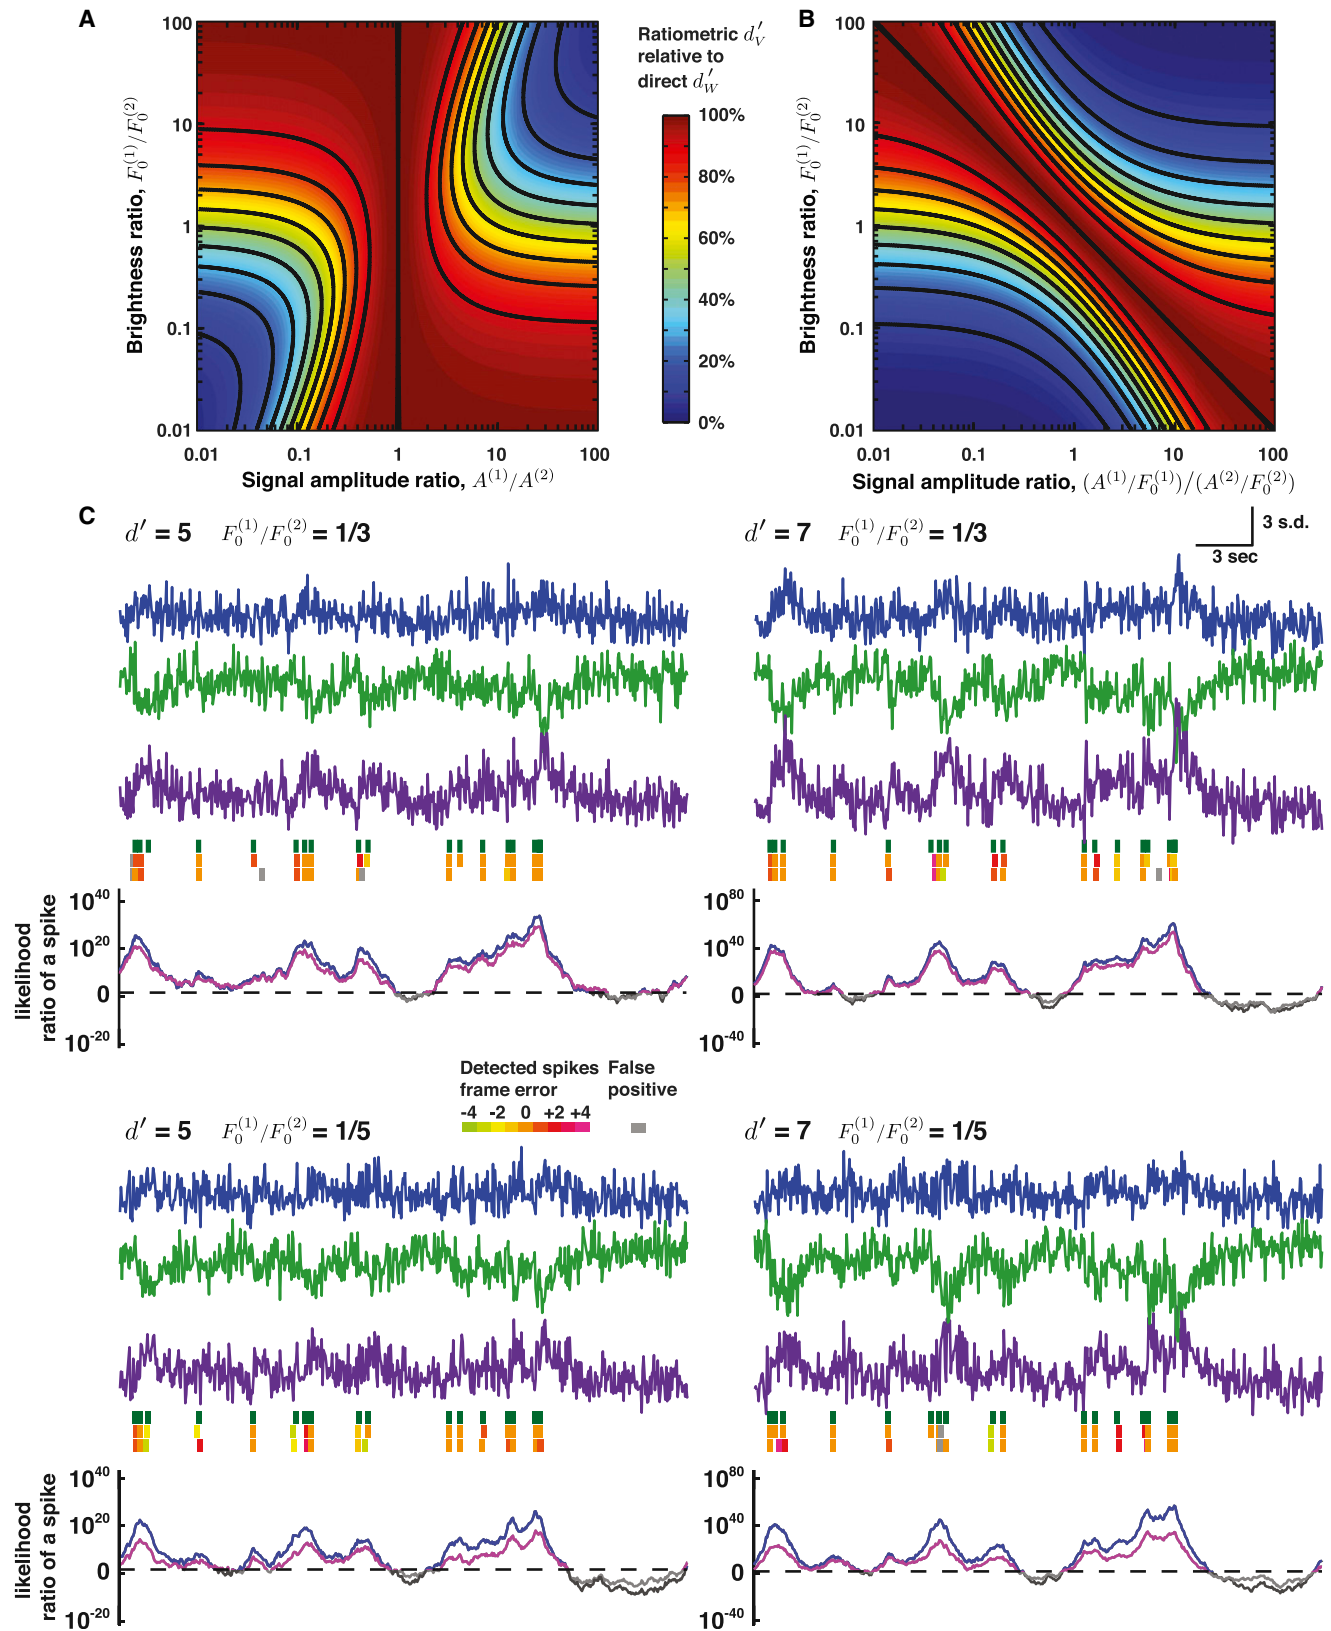


FIGURE 4 Spike detection using FRET indicators. In the shot noise limit, ratiometric treatment of FRET indicators generally makes suboptimal use of the photon statistics for spike detection (A and B). (A) The decline in spike detectability is plotted as a joint function of the ratios of signal amplitude $A^{(1)}/A^{(2)}$ and the background brightness $F_0^{(1)}/F_0^{(2)}$. When the fluorescence transients are equal in amplitudes, a direct statistical analysis does not yield any benefit.

generalizes to FRET indicators by summing the log-likelihood ratios of both color channels (Fig. 4 C) and to spike train estimation by adjusting the null hypothesis to include previously detected spikes. Finally, the filter of Eq. 6 is simple to implement and truer to the statistics of optical recordings than other forms of deconvolution. More sophisticated methods exist for spike train estimation involving supervised machine learning (16), particle filtering (17), or nonnegative deconvolution (18), but these are more challenging to implement and not yet widely adopted. Although Eq. 6 does not prescribe how to estimate the optical waveform parameters, our simulations show that successful spike extraction does not usually hinge on accurate parameter estimates. Pseudo-expectation maximization approaches that alternate estimation of the spike train with estimation of the optical parameters may facilitate spike train estimation (17,18).

Notable effort has focused on creating brighter fluorescent indicators (3,11–13). Yet, d' improves only as the square-root of the brightness for fixed $\Delta F/F \approx A/F_0$ (Eq. 9). The linear dependence of d' on $\Delta F/F$ suggests greater gains may be possible by increasing $\Delta F/F$, such as by reducing background emissions and increasing the signal's dynamic range in response to Ca^{2+} binding. Although $\Delta F/F$ is bounded above by the indicator response in the labeled cell type, background flux from mistargeted indicator and poor rejection of out-of-focus emissions reduce the $\Delta F/F$ of the data. This is especially pertinent with voltage indicators, due to the challenge of selectively labeling neural membranes (46).

We illustrate these considerations using published in vivo recordings of dendritic Ca^{2+} spikes from cerebellar Purkinje neurons (6,7). Using the Ca^{2+} indicator Oregon Green BAPTA-1, two-photon microscopy attains $\Delta F/F$ values of $\sim 30\%$ (7,47), but one-photon microscopy attains $\Delta F/F$ values of only $\sim 0.5\text{--}1.5\%$ due to lack of optical sectioning (6,48). Yet, one-photon microscopy captures far more photons per frame, because the pixels are sampled in parallel rather than by laser-scanning. We compared d' values for the two modalities and determined that the improved photon collection of one-photon microscopy adequately compensates for the increased background, yielding comparable spike discriminability for both forms of microscopy. Approaches that furnish one-photon microscopy with back-

ground rejection mechanisms, such as planar illumination (45), have the potential to improve spike discriminability, at least at shallow tissue depths where scattering does not substantially diminish signal quality.

As our calculations show, another key influence on the fidelity of spike detection is the duration of the indicator's response. Several genetically encoded Ca^{2+} indicators, such as TN-XXL (13), d3cpv (3), and members of the Cameleon-Nano family (11), exhibit prolonged optical transients (> 1 s) in response to one spike, allowing these indicators to emit far more signal photons than an indicator with comparable brightness and dynamic range but brisker kinetics. Especially when monitoring neurons with temporally sparse firing patterns, these indicators may offer superior d' values than faster indicators such as Oregon Green BAPTA-1 (OGB-1) ($> 100\text{--}200$ ms). However, indicators with high Ca^{2+} affinity may also perturb a cell's native Ca^{2+} -dependent processes (2,23). We found that d' was relatively insensitive to the fine temporal structure of the transient waveform, depending most sensitively on the rise time and total number of signal photons collected above background. To retain single spike resolution, the indicator's transient duration can be prolonged up to roughly the reciprocal of the spike rate.

Although our model was developed with Ca^{2+} indicators in mind, our formalism also applies to membrane voltage indicators. Existing voltage indicators have been successfully applied to individual cells or axons (49–52), invertebrate neurons (49–51,53), tissue slices and cultured cells (49–52,54,55), and cells exhibiting prolonged subthreshold voltage transients (56). However, the available voltage indicators generally fail to reliably report single action potentials in the live mammalian brain (46). Voltage depolarizations associated with Na^+ spikes are typically ~ 1 ms or less, and the associated optical signals lack the prolonged decay times seen with Ca^{2+} indicators. Thus, for voltage indicators to achieve spike discriminability comparable to Ca^{2+} indicators, they must be brighter (> 100 times), generate a larger signal response (> 10 times), or be used in a way that affords each neuron a far higher proportion of the frame acquisition time. This suggests possible advantages of voltage indicators based on changes in emission color, provided the background could be minimized in the emission color band used to signal a spike occurrence (Fig. 3) (57).

However, when the ratio of signal amplitudes differs from unity, a ratiometric analysis is inferior. (B) The decline in spike detectability is plotted as a joint function of the ratios of the background-normalized signal amplitudes $(A^{(1)}/F_0^{(1)})/(A^{(2)}/F_0^{(2)})$ and the background brightness. (C) Simulated pairs of traces (blue and green traces) for the donor and acceptor channels of a FRET indicator. Here, we calculate the log-likelihood ratio of a spike with a moving window of 60 observations using Eq. 6 (the sum of the log-likelihood ratios for the donor and acceptor traces) and a Gaussian approximation (Eqs. S49–S51 in the Supporting Material) to the ratio of the donor and acceptor traces (purple trace). The true spike train is shown (green, top spike trains) above the spike trains estimated via direct analysis of the channels (middle spike trains) or a ratiometric analysis (bottom spike trains). (Spikes in non orange hues) Spikes estimated with errors in frame timing. (Gray spikes) False positive spikes. (Light- and dark-gray traces) Log-likelihood ratio of a spike for the ratiometric and direct channel analyses, respectively. Threshold crossing events for these two traces (highlighted in purple and blue, respectively) were analyzed using the greedy algorithm approach to detect spikes, applying the same equal-cost condition to set the detection threshold for both methods. Note the superior log-likelihood ratio values attained using the direct analysis. The strong correlations between the ratiometric and direct channel calculations occur because the same simulation trial was used for both calculations. In these simulations, $\Delta F/F = 0.05$ for both channels, $\tau = 1.0$ s, $\nu = 20$ Hz, spike rate $\lambda = 0.5$ Hz.

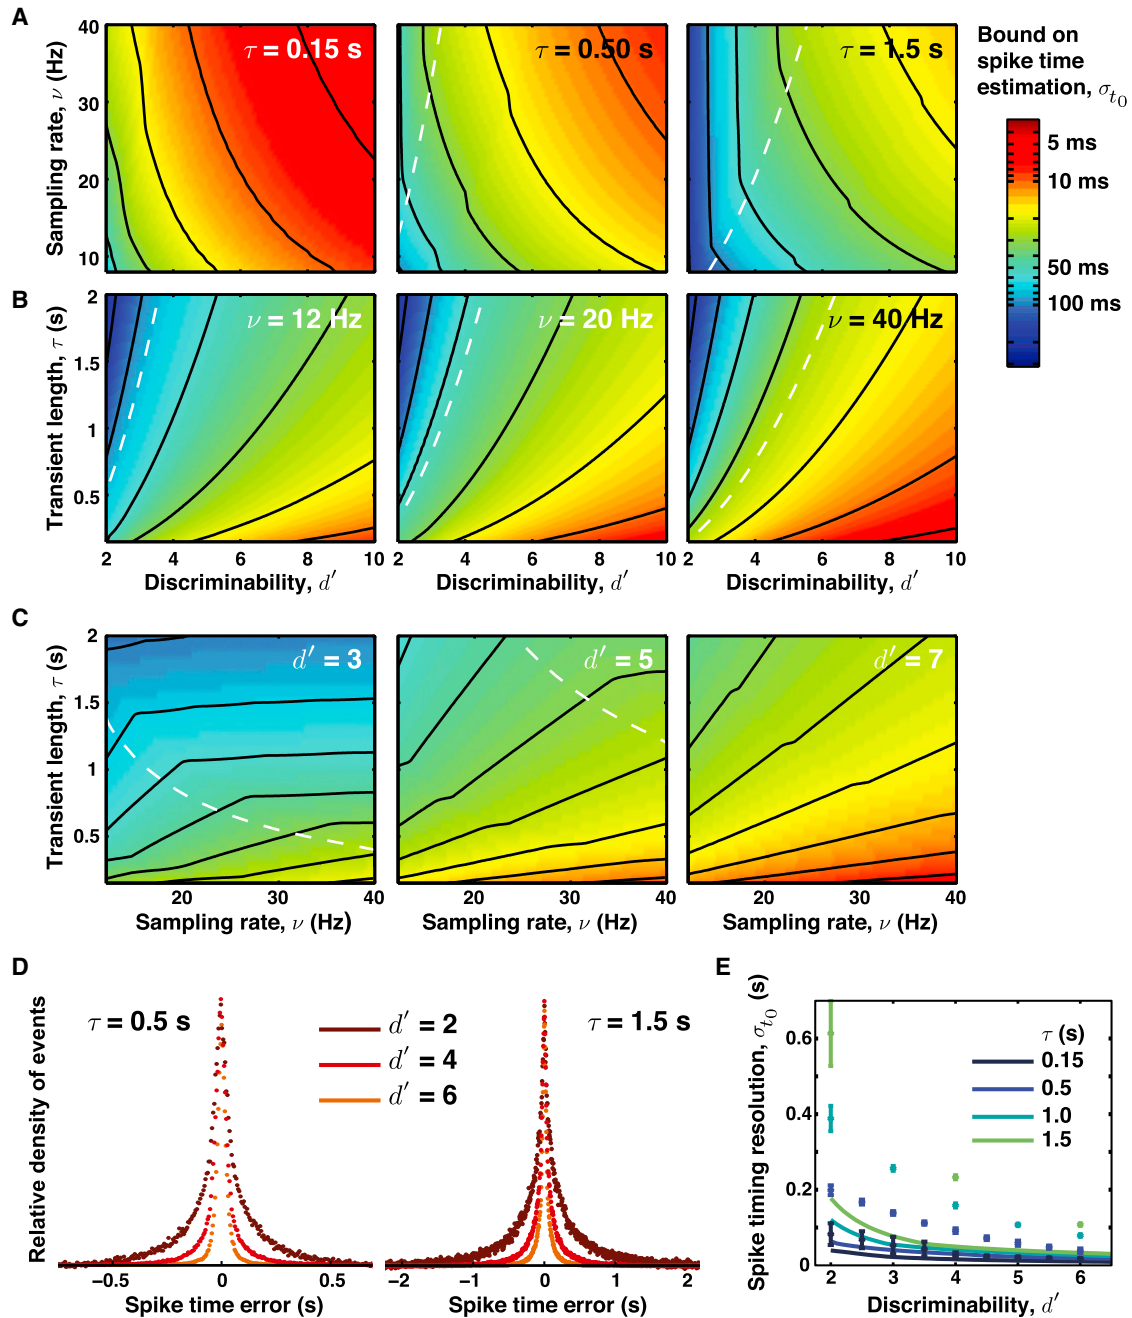


FIGURE 5 Chapman-Robbins lower bounds on the estimation of spike times. The ability to accurately estimate a spike's time is governed by the discriminability, d' , the sampling rate, ν , and the duration of the signal transient, τ . For fixed d' , spike timing precision improves with shorter optical transients and higher sampling rates. The Chapman-Robbins lower bound can be either lower or higher than the sampling rate and exhibits strong dependence on discriminability d' and transient duration. (A) Comparison of three optical indicators with different decay kinetics. A fast optical indicator ($\tau \sim 0.15$ s) such as Oregon Green BAPTA-1 (OGB-1) can localize spikes to ~ 30 ms, even in regimes with modest SNR regimes ($d' \sim 3$). Slower indicators with time constants resembling those of GCaMP3 (~ 0.5 s) exhibit poorer spike timing characteristics. In regimes with modest SNR, even slower indicators with kinetics resembling those of d3cpv or TN-XXL (1.5 s) allow spike localization with ~ 100 ms resolution. (B) Comparison of estimation bounds at three different sampling rates. Increasing the sampling rate modestly improves spike timing resolution. (C) Increasing discriminability also significantly improves spike timing resolution. (A–C, white dashed line) Boundary of the regime of temporal super-resolution. (D) Simulations of spike timing resolution. Using a brute-force maximum likelihood method for determining the spike time, we obtained histograms of the spike time error for two indicators with distinct temporal dynamics. Note the different time scales on the two panels. For visual clarity, histograms are shown normalized to a common peak value. (E) Plots of simulated spike timing resolution and the theoretically calculated Chapman-Robbins lower bound. The simulations generally do not attain the Chapman-Robbins lower bound, especially for situations with low SNR and slow temporal dynamics. Simulations in panels D and E were done using $\nu = 20$ Hz.

When reporting the properties of fluorescent activity indicators, researchers can facilitate quantitative comparisons by carefully describing the imaging context. Assessments of indicator performance using laser line scans over a single cell or detection at fixed points in space inflate d' by the square-root of the laser dwell time. For example, if a cell occupies 1% of an image frame, d' can be inflated by a factor of 10 by fixing the laser beam over the cell. Similarly, pipette loading of the dye to high levels or constructing unusually low background levels also overstate indicator performance as compared to state-of-the-art studies of large numbers of individual neurons. Unfortunately, the brightness contributions from extinction coefficients, expression patterns, and labeling densities are sometimes reported ambiguously, making it difficult to predict d' in novel experimental situations. For voltage indicators this ambiguity is compounded by the challenges of reliable membrane targeting (46).

In addition to examining spike detection, our study introduces limits on spike time estimation based on optically recorded neural activity. Computational methods for spike train inference should be evaluated jointly on their capabilities for spike detection and timing estimation. The Chapman-Robbins lower bound on timing errors is distinct from d' and strongly depends on transient duration for fixed d' . Intuitively, for a given d' value, a brisk transient makes timing estimation easier than a prolonged one. More surprisingly, faster sampling rates in our simulations yielded only modestly improved spike timing precision. The substantially improved timing estimates achieved using acousto-optic random-access laser-scanning microscopy (15) seem to stem from the increased dwell times over the chosen subset of cells, which increases spike discriminability, rather than faster sampling rates.

Even with modest SNR and sampling rates, one can resolve spike times with greater precision than suggested by the naïve Nyquist criterion. Many spike train estimation routines assign spikes at discrete times, but it may be preferable to treat spike times as continuous variables. Our calculations bound the statistical confidence of spike times, and researchers should ideally report confidence intervals alongside spike trains. Importantly, the Chapman-Robbins lower bound is only valid for unbiased estimators, and we did not examine how parameter misestimation impacts spike time estimation.

Future work might extend our results in several ways. Our treatment considered only isolated spikes. Spike detection with temporally overlapping spike waveforms should be rigorously examined, including with indicator saturation effects that reduce signals from a rapid succession of action potentials. This is not easily described in the framework presented here, but could be analyzed by inclusion of additional parameters. Our reliance on d' is only suitable for binary classification problems, so spikes that may have been detected at times not coincident with, but within time τ , of an

actual spike are not considered detected in our classification. A more complete treatment might ameliorate this deficiency. Finally, the ideas presented here may be applicable to the inference of time-varying spike rates from fluorescence data.

SUPPORTING MATERIAL

Seven sections, one figure, and 58 equations are available at [http://www.biophysj.org/biophysj/supplemental/S0006-3495\(12\)00969-1](http://www.biophysj.org/biophysj/supplemental/S0006-3495(12)00969-1).

The authors thank Limor Freifeld, Eran Mukamel, Damon Clark, and Jesse Marshall for helpful conversations.

B.A.W. was supported by a Bio-X Research Fellowship. J.E.F. was supported by a predoctoral fellowship from the National Science Foundation and a National Science Foundation Integrative Graduate Education and Research Traineeship under Award No. 0801700. M.J.S.'s work on imaging was supported by the Packard, Paul G. Allen Family, and W.M. Keck foundations, the National Institute of Neurological Disorders and Stroke, the National Institute on Drug Abuse, the National Science Foundation, the Office of Naval Research, and a National Institutes of Health Director's Pioneer award.

REFERENCES

1. Wilt, B. A., L. D. Burns, ..., M. J. Schnitzer. 2009. Advances in light microscopy for neuroscience. *Annu. Rev. Neurosci.* 32:435–506.
2. Helmchen, F., J. G. G. Borst, and B. Sakmann. 1997. Calcium dynamics associated with a single action potential in a CNS presynaptic terminal. *Biophys. J.* 72:1458–1471.
3. Wallace, D. J., S. Meyer zum Alten Borgloh, ..., M. T. Hasan. 2008. Single-spike detection in vitro and in vivo with a genetic Ca^{2+} sensor. *Nat. Methods.* 5:797–804.
4. Mao, T., D. H. O'Connor, ..., K. Svoboda. 2008. Characterization and subcellular targeting of GCaMP-type genetically-encoded calcium indicators. *PLoS ONE.* 3:e1796.
5. Kerr, J. N. D., D. S. Greenberg, and F. Helmchen. 2005. Imaging input and output of neocortical networks in vivo. *Proc. Natl. Acad. Sci. USA.* 102:14063–14068.
6. Flusberg, B. A., A. Nimmerjahn, ..., M. J. Schnitzer. 2008. High-speed, miniaturized fluorescence microscopy in freely moving mice. *Nat. Methods.* 5:935–938.
7. Mukamel, E. A., A. Nimmerjahn, and M. J. Schnitzer. 2009. Automated analysis of cellular signals from large-scale calcium imaging data. *Neuron.* 63:747–760.
8. Hires, S. A., L. Tian, and L. L. Looger. 2008. Reporting neural activity with genetically encoded calcium indicators. *Brain Cell Biol.* 36: 69–86.
9. Sjulson, L., and G. Miesenböck. 2007. Optical recording of action potentials and other discrete physiological events: a perspective from signal detection theory. *Physiology (Bethesda).* 22:47–55.
10. Yasuda, R., E. A. Nimchinsky, ..., K. Svoboda. 2004. Imaging calcium concentration dynamics in small neuronal compartments. *Sci. STKE.* 2004:pl5.
11. Horikawa, K., Y. Yamada, ..., T. Nagai. 2010. Spontaneous network activity visualized by ultrasensitive Ca^{2+} indicators, yellow Camelion-Nano. *Nat. Methods.* 7:729–732.
12. Tian, L., S. A. Hires, ..., L. L. Looger. 2009. Imaging neural activity in worms, flies and mice with improved GCaMP calcium indicators. *Nat. Methods.* 6:875–881.
13. Mank, M., A. F. Santos, ..., O. Griesbeck. 2008. A genetically encoded calcium indicator for chronic in vivo two-photon imaging. *Nat. Methods.* 5:805–811.

14. Hendel, T., M. Mank, ..., D. F. Reiff. 2008. Fluorescence changes of genetic calcium indicators and OGB-1 correlated with neural activity and calcium in vivo and in vitro. *J. Neurosci.* 28:7399–7411.
15. Grewe, B. F., D. Langer, ..., F. Helmchen. 2010. High-speed in vivo calcium imaging reveals neuronal network activity with near-millisecond precision. *Nat. Methods.* 7:399–405.
16. Sasaki, T., N. Takahashi, ..., Y. Ikegaya. 2008. Fast and accurate detection of action potentials from somatic calcium fluctuations. *J. Neurophysiol.* 100:1668–1676.
17. Vogelstein, J. T., B. O. Watson, ..., L. Paninski. 2009. Spike inference from calcium imaging using sequential Monte Carlo methods. *Biophys. J.* 97:636–655.
18. Vogelstein, J. T., A. M. Packer, ..., L. Paninski. 2010. Fast nonnegative deconvolution for spike train inference from population calcium imaging. *J. Neurophysiol.* 104:3691–3704.
19. Chapman, D. G., and H. Robbins. 1951. Minimum variance estimation without regularity assumptions. *Ann. Math. Stat.* 22:581–586.
20. Hammersley, J. M. 1950. On estimating restricted parameters. *J. R. Stat. Soc. B.* 12:192–240.
21. Grynkiewicz, G., M. Poenie, and R. Y. Tsien. 1985. A new generation of Ca^{2+} indicators with greatly improved fluorescence properties. *J. Biol. Chem.* 260:3440–3450.
22. Miyawaki, A., J. Llopis, ..., R. Y. Tsien. 1997. Fluorescent indicators for Ca^{2+} based on green fluorescent proteins and calmodulin. *Nature.* 388:882–887.
23. Helmchen, F., K. Imoto, and B. Sakmann. 1996. Ca^{2+} buffering and action potential-evoked Ca^{2+} signaling in dendrites of pyramidal neurons. *Biophys. J.* 70:1069–1081.
24. Yaksi, E., and R. W. Friedrich. 2006. Reconstruction of firing rate changes across neuronal populations by temporally deconvolved Ca^{2+} imaging. *Nat. Methods.* 3:377–383.
25. Poor, H. V. 1998. An Introduction to Signal Detection and Estimation. Springer, New York.
26. Cormen, T. H., C. E. Leiserson, ..., C. Stein. 2009. Introduction to Algorithms. The MIT Press, Cambridge, MA.
27. Kay, S. M. 1993. Fundamentals of Statistical Signal Processing. Prentice-Hall, Englewood Cliffs, NJ.
28. Reiffen, B., and H. Sherman. 1963. An optimum demodulator for Poisson processes: photon source detectors. *Proc. IEEE.* 51:1316–1320.
29. Joucla, S., A. Pippow, ..., C. Pouzat. 2010. Quantitative estimation of calcium dynamics from ratiometric measurements: a direct, nonratiometric method. *J. Neurophysiol.* 103:1130–1144.
30. Rieke, F., D. Warland, ..., W. Bialek. 1999. Spikes: Exploring the Neural Code. MIT Press, Cambridge, MA.
31. Skaggs, W. E., B. L. McNaughton, ..., C. A. Barnes. 1996. Theta phase precession in hippocampal neuronal populations and the compression of temporal sequences. *Hippocampus.* 6:149–172.
32. O’Keefe, J., and M. L. Recce. 1993. Phase relationship between hippocampal place units and the EEG θ -rhythm. *Hippocampus.* 3:317–330.
33. Buzsáki, G. 2002. Theta oscillations in the hippocampus. *Neuron.* 33:325–340.
34. Usrey, W. M., and R. C. Reid. 1999. Synchronous activity in the visual system. *Annu. Rev. Physiol.* 61:435–456.
35. Covey, E., and J. H. Casseday. 1999. Timing in the auditory system of the bat. *Annu. Rev. Physiol.* 61:457–476.
36. Oertel, D. 1999. The role of timing in the brain stem auditory nuclei of vertebrates. *Annu. Rev. Physiol.* 61:497–519.
37. Trussell, L. O. 1999. Synaptic mechanisms for coding timing in auditory neurons. *Annu. Rev. Physiol.* 61:477–496.
38. Duemani Reddy, G., K. Kelleher, ..., P. Saggau. 2008. Three-dimensional random access multiphoton microscopy for functional imaging of neuronal activity. *Nat. Neurosci.* 11:713–720.
39. Reddy, G. D., and P. Saggau. 2005. Fast three-dimensional laser scanning scheme using acousto-optic deflectors. *J. Biomed. Opt.* 10:064038.
40. Vucinić, D., and T. J. Sejnowski. 2007. A compact multiphoton 3D imaging system for recording fast neuronal activity. *PLoS ONE.* 2:e699.
41. Ji, N., J. C. Magee, and E. Betzig. 2008. High-speed, low-photodamage nonlinear imaging using passive pulse splitters. *Nat. Methods.* 5:197–202.
42. Ji, N., D. E. Milkie, and E. Betzig. 2010. Adaptive optics via pupil segmentation for high-resolution imaging in biological tissues. *Nat. Methods.* 7:141–147.
43. Engelbrecht, C. J., W. Göbel, and F. Helmchen. 2009. Enhanced fluorescence signal in nonlinear microscopy through supplementary fiber-optic light collection. *Opt. Express.* 17:6421–6435.
44. Greenberg, D. S., A. R. Houweling, and J. N. D. Kerr. 2008. Population imaging of ongoing neuronal activity in the visual cortex of awake rats. *Nat. Neurosci.* 11:749–751.
45. Holekamp, T. F., D. Turaga, and T. E. Holy. 2008. Fast three-dimensional fluorescence imaging of activity in neural populations by objective-coupled planar illumination microscopy. *Neuron.* 57:661–672.
46. Peterka, D. S., H. Takahashi, and R. Yuste. 2011. Imaging voltage in neurons. *Neuron.* 69:9–21.
47. Ozden, I., H. M. Lee, ..., S. S. Wang. 2008. Identification and clustering of event patterns from in vivo multiphoton optical recordings of neuronal ensembles. *J. Neurophysiol.* 100:495–503.
48. Ghosh, K. K., L. D. Burns, ..., M. J. Schnitzer. 2011. Miniaturized integration of a fluorescence microscope. *Nat. Methods.* 8:871–878.
49. Salzberg, B. M., H. V. Davila, and L. B. Cohen. 1973. Optical recording of impulses in individual neurones of an invertebrate central nervous system. *Nature.* 246:508–509.
50. Cohen, L. B., B. M. Salzberg, ..., C. H. Wang. 1974. Changes in axon fluorescence during activity: molecular probes of membrane potential. *J. Membr. Biol.* 19:1–36.
51. Sacconi, L., D. A. Dombeck, and W. W. Webb. 2006. Overcoming photodamage in second-harmonic generation microscopy: real-time optical recording of neuronal action potentials. *Proc. Natl. Acad. Sci. USA.* 103:3124–3129.
52. Acker, C. D., P. Yan, and L. M. Loew. 2011. Single-voxel recording of voltage transients in dendritic spines. *Biophys. J.* 101:L11–L13.
53. Zecević, D., J.-Y. Wu, ..., C. X. Falk. 1989. Hundreds of neurons in the *Aplysia* abdominal ganglion are active during the gill-withdrawal reflex. *J. Neurosci.* 9:3681–3689.
54. Sacconi, L., J. Mapelli, ..., F. S. Pavone. 2008. Optical recording of electrical activity in intact neuronal networks with random access second-harmonic generation microscopy. *Opt. Express.* 16:14910–14921.
55. Holthoff, K. P., D. Zecević, and A. Konnerth. 2010. Rapid time course of action potentials in spines and remote dendrites of mouse visual cortex neurons. *J. Physiol.* 588:1085–1096.
56. Akemann, W., H. Mutoh, ..., T. Knöpfel. 2010. Imaging brain electric signals with genetically targeted voltage-sensitive fluorescent proteins. *Nat. Methods.* 7:643–649.
57. Kralj, J. M., D. R. Hochbaum, ..., A. E. Cohen. 2011. Electrical spiking in *Escherichia coli* probed with a fluorescent voltage-indicating protein. *Science.* 333:345–348.