ELSEVIER

Journal of Neuroscience Methods 159 (2007) 325–336

JOURNAL OF NEUROSCIENCE METHODS

www.elsevier.com/locate/jneumeth

Automated classification of evoked quantal events

Mark Lancaster [a], Kert Viele [a], A.F.M. Johnstone [b], Robin L. Cooper [b,*]

[a] *Department of Statistics, University of Kentucky, Lexington, KY 40506-0027, United States*
[b] *Department of Biology, University of Kentucky, Lexington, KY 40506-0225, United States*

Received 21 November 2005; received in revised form 8 July 2006; accepted 10 July 2006

Abstract

We provide both theoretical and computational improvements to the analysis of synaptic transmission data. Theoretically, we demonstrate the correlation structure of observations within evoked postsynaptic potentials (EPSP) are consistent with multiple random draws from a common autoregressive moving-average (ARMA) process of order (2, 2). We use this observation and standard time series results to construct a statistical hypothesis testing procedure for determining whether a given trace is an EPSP. Computationally, we implement this method in R, a freeware statistical language, which reduces the amount of time required for the investigator to classify traces into EPSPs or non-EPSPs and eliminates investigator subjectivity from this classification. In addition, we provide a computational method for calculating common functionals of EPSPs (peak amplitude, decay rate, etc.). The methodology is freely available over the internet. The automated procedure to index the quantal characteristics greatly facilitates determining if any one or multiple parameters are changing due to experimental conditions. In our experience, the software reduces the time required to perform these analyses from hours to minutes.

Keywords: Synapse; Quantal; Analysis; Computational

1. Introduction

Chemical synaptic transmission occurs by the release of transmitter in "packets" from the presynaptic nerve terminal which then gives rise to a postsynaptic current that may produce a synaptic potential depending on the electrochemical driving gradients of the ions. The synaptic potentials are incremental in relation to the numbers of packets of transmitter released (Del Castillo and Katz, 1954a). For the most part, the incremental increase in the excitatory synaptic potentials during an evoked release from the nerve terminal matches the average unitary size of spontaneous excitatory synaptic potentials that are recorded in the absence of stimulating the presynaptic nerve terminal. This phenomenon occurs for postsynaptic cells that are non-spiking or for subthreshold conditions to the induction of generating action potentials. Such observations lead to the "quantal hypothesis" proposed by Del Castillo and Katz (1954a). This is generally accepted as a packet of neurotransmitter within a clear core vesicle that when released from the presynaptic nerve terminal into the synaptic cleft will produce a quantal postsynaptic current that is manifested as a quantal potential across the postsynaptic membrane. There is some variation to quantal events that can arise due to presynaptic as well as postsynaptic factors. A number of investigations into the cause of quantal variation and reasons for non-linear addition of evoked synaptic potentials have been undertaken (Del Castillo and Katz, 1954a,b; Ginsborg, 1973; Martin, 1955 and reviews by Faber et al., 1998; McLachlan, 1978). In addition to understanding mechanisms in quantal variations, it is also of interest to be able to index synaptic strength and to determine the characteristics of synapses such as the number of places vesicles fuse with the presynaptic membrane and with what probability vesicles will fuse at each site. Thus, synaptic preparations that allow analysis of the shapes of quantal events and distributions in the numbers of evoked quantal occurrences with a minimum degree of complexity provide insight in better understanding the principles of synaptic transmission.

Many previous analyses (Del Castillo and Katz, 1954b and see review article of Faber et al., 1998) have focused on estimating the number of release sites n and the probability of release p solely from the number of events in each current trace. Thus, traces were classified as "failures", "singles", "doublets", etc., and n and p were estimating using Binomial or Poisson sampling methods. These methods are well known in the statistics

* Corresponding author. Tel.: +1 859 257 5950; fax: +1 859 257 1717.
E-mail address: rlcoop1@uky.edu (R.L. Cooper).

0165-0270/$ – see front matter © 2006 Elsevier B.V. All rights reserved.
doi:10.1016/j.jneumeth.2006.07.014

literature to be unstable (Olkin et al., 1981; Viele et al., 2003) meaning that while an accurate estimate of the mean quantal content $m = np$ may be accurately acquired, the separate estimates of n and p are highly variable, meaning that even for large sample sizes the estimates of n and p may separately be far from their true values. Olkin et al. (1981), in particular, presents an example where changing a single data point by one unit changes the estimate of n from 99 to 190.

Fortunately, there is a large amount of information in the current trace in addition to the number of events. Differences in sizes and shapes of the individual EPSPs may indicate multiple sites at work. Thus, for example, Viele et al. (2003) cluster functionals of the EPSPs such as peak amplitude and area under the curve (AUC) and can produce precise estimates of n and p where counting methods would fail (if one uses counting method to compute, for example, 95% confidence intervals, then the resulting confidence intervals will contain a very large number of possible values of n).

Unfortunately, large amounts of investigator preprocessing have been required to acquire this functionals for cluster analysis. In this paper we provide the following theoretical and computational advances.

1. Previous analyses (Silver, 2003) have noted the voltage observations within an EPSP are strongly correlated. However, these previous analyses have not investigated whether the correlation structure is fixed over time. We demonstrate that a single autoregressive, moving-average (ARMA) process of order (2, 2) describes the variation seen in the correlation structure over time, indicating that any variation over time is simply due to chance, not temporal changes.
2. We use the ARMA process develop a statistical hypothesis test for classifying traces into EPSPs or failures. Demonstrating the ARMA process has no temporal shifts is fundamental to this procedure, since the parameters of the ARMA process can then be estimated by aggregating all the traces to produce a single estimate. This single estimate, since it is estimated from a very large sample size, is quite precise.
3. We provide an implementation of the statistical hypothesis test in R R Development Core Team (2004), a freeware statistical language. This implementation is freely available from the internet. The result is that the investigator is freed from having to manually classify traces as either EPSPs or failures.
4. Finally, we provide an automated method for computing several commonly used EPSP functionals. These are the peak amplitude, the latency, the decay time τ, the rise time, and the area under the curve (AUC). These methods are entirely automated, unlike current software which require the user to manually locate begin and end points for each EPSP.

Overall, in our experience the software reduces the task of processing 1000 traces from hours to minutes by eliminating the tedious viewing of all traces, even the failures, as well as not having to place cursors for measures of evoked events.

In this study a sample data set is used to illustrate the detailed procedures implemented by the software for analysis of the evoked quantal events. This analysis method and related software can be used for a wide variety of synaptic preparations and experimental procedures. The software is readily modifiable by a user familiar with the common statistical "R" language. The standard export format allows easy use of the obtained quantal measures for statistical analysis in any commercial available software package.

2. Methods

2.1. The preparation

All experiments were performed using the first walking leg of crayfish, Procambarus clarkii, measuring 4–6 cm in body length (Atchafalaya Biological Supply Co., Raceland, LA). The opener muscle of the first walking legs was prepared by the standard dissection (Cooper et al., 1995). The tissue was pinned out in a Slygard dish for viewing with a Nikon, Optiphot-2 upright fluorescent microscope using a 40× (0.55 NA) Nikon water immersion objective. Dissected preparations were maintained in crayfish saline (modified Van Harreveld's solution: 205 mM NaCl; 5.3 mM KCl; 13.5 mM $CaCl_2$; 2.45 mM $MgCl_2$; 0.5 mM Hepes/NaOH, pH 7.4) at 14 °C.

2.2. Physiology: field excitatory postsynaptic potentials (fEPSPs)

Synaptic potentials were obtained with a focal macropatch electrode Dudel (1981) by lightly placing a 10–15 μm diameter, fire polished, glass electrode directly over a spatially isolated varicosity. The lumen of the patch electrode was filled with the same solution as the bathing medium. The seal resistance was in the range of 100 KΩ to 1 MΩ. All events were obtained with an Axoclamp 2b (Axon Instruments) 0.1× LU head stage acquired at 20 kHz without additional filtering. The varicosities on the living terminals were visualized by the use of the vital fluorescent dye 4-Di-2-ASP (Molecular Probes) (Cooper et al., 1995; Magrassi et al., 1987). The evoked field excitatory postsynaptic potentials (fEPSPs) and field miniature excitatory postsynaptic potentials (fmEPSPs) were recorded. The data sample shown for description of the methodology was stimulated at 1 Hz.

3. Statistical analysis of quantal events

3.1. Motivation for examining functionals of current trace

A common classical method for obtaining the quantal parameters (n) and (p) is based on direct counts of the number of failures, single events, double events, etc. As discussed previously (Viele et al., 2003) the direct count method may be used to estimate the mean quantal content (the mean number of events per stimulus, usually denoted as m), or they may be used to estimate (n) and (p) by fitting various discrete distributions such as a Poisson or some type of binomial to the count data, thus predicting the distribution of the number of events resulting from each stimulus. It is much easier to estimate (m) than simultaneously estimate the parameters n and p. Olkin et al. (1981) described the issue that if both n and p are unknown in a binomial distri-

bution, then the mean np may be accurately estimated (the mean quantal content m) but estimates of the individual parameters n and p are unstable.

Essentially, the difficulty lies with the fact the different values of n and p which have the same mean quantal content np are almost indistinguishable on the basis of count data alone. Assuming a binomial likelihood, If $n = 2$ and $p = 0.1$, the probability of getting a failure is 0.810, the probability of a single is 0.180, and the probability of a double is 0.010. In contrast, if $n = 10$ and $p = 0.02$, the probability of a failure is 0.817, the probability of a single is 0.167, and the probability of a doublet is 0.015 (there is limited probability of getting a triple or more). These probabilities are so close that one requires extremely large sample sizes to distinguish them, and thus to distinguish the underlying n and p. Note that both distribution have the same value of $np = 0.2$.

The purpose of the current research is to examine information within each trace to refine the coarse classifications such as "single release event" with information from the underlying trace. Previously, we used a method to determine if clusters of evoked quantal potentials occurred by only using the entire area of the fEPSPs from single evoked events (Viele et al., 2003). In this current report, we have automated the computation of the peak amplitude and provided an automated method for computing other functionals such as the latency time, the rise time, and the decay time of the release events.

3.2. Classification of traces into release events and failures

Fig. 1 shows the raw data obtained from an experiment in which the preparation was only exposed to saline and stimulated at 1 Hz. Fig. 2 contains a closeup of the area "A" in Fig. 1 corresponding to the area where EPSPs occur. In this preparation, there are two artifacts common to all the traces. The first one is the stimulus artifact that is recorded as the stimulus travels in the saline bath. This first artifact dominates the voltage axis of the plot and thus we have "clipped" the voltage axis of the plot. The second artifact is the extracellular recording of the action potential monitored over the nerve terminal, also referred to as a spike. This artifact is more easily viewed in Fig. 2, the closeup. The portion of the graph that holds our interest is immediately after the second artifact. The nerve terminal may or may not respond by vesicle fusion to the voltage stimulus. If a release event occurs, it will occur within a 20–30 ms window after the nerve terminal is depolarized.

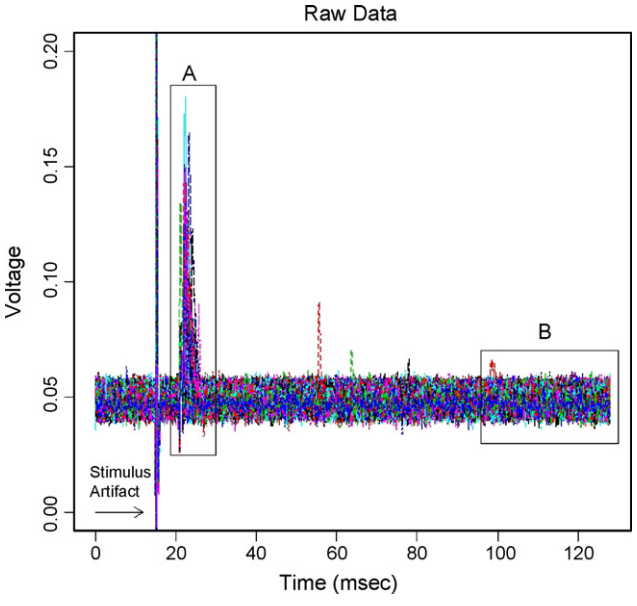

Fig. 1. Raw data. The y-axis (voltage) has been clipped to remove a stimulation artifact just before 20 ms. The area marked "A" corresponds to the area where EPSPs occur. The area marked "B" corresponds to the area used to estimate the error structure of the data.

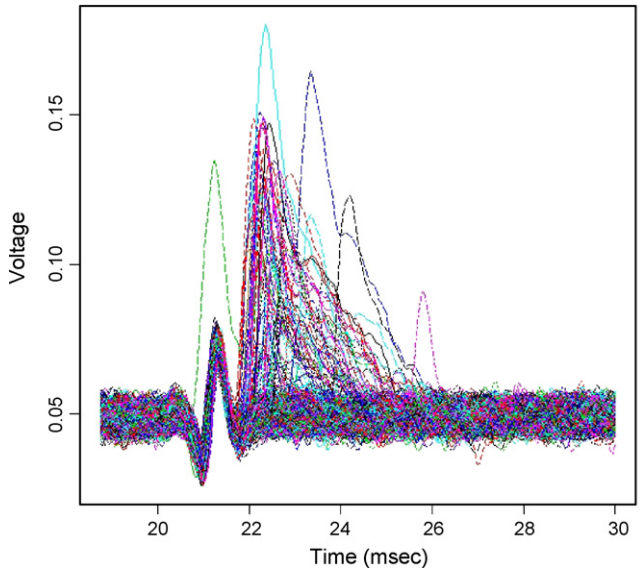

Fig. 2. Closeup of the region "A" from Fig. 1 where the EPSPs occur. The EPSPs are preceded by artifacts which are removed automatically as part of the analysis.

Some traces are obviously release events and some are not. Fig. 3 shows three traces within the zoomed in area from Fig. 2. Fig. 3A shows an obvious release event (the artifact is followed by a large increase), while Fig. 3C shows an obvious failure (except for the artifact there is no obvious change in voltage besides noise). Unfortunately, there is a continuum of magnitudes for the release events, which results in some traces that are not obviously either release events or failures. Fig. 3B shows a trace that is ambiguous in that the largest voltage values of the trace occur just after the artifact. It is unclear whether this increase is due to noise (e.g. an increase in voltage appeared by chance) or due to a "small" release. It is possible that some may consider the particular trace in Fig. 3B to be an obvious failure, but for everyone some of the traces will be ambiguous. One of the purposes of this paper is to recognize that different investigators can have different subjective definitions of what is questionable and what is not, and thus a well justified method for automating the classification into release or non-release events is valuable. In what follows, our automated procedure classifies the trace in Fig. 3B as a failure, but it is just on the borderline of being declared a release.

To classify the release events, there are a few computational issues that must be considered, in order.

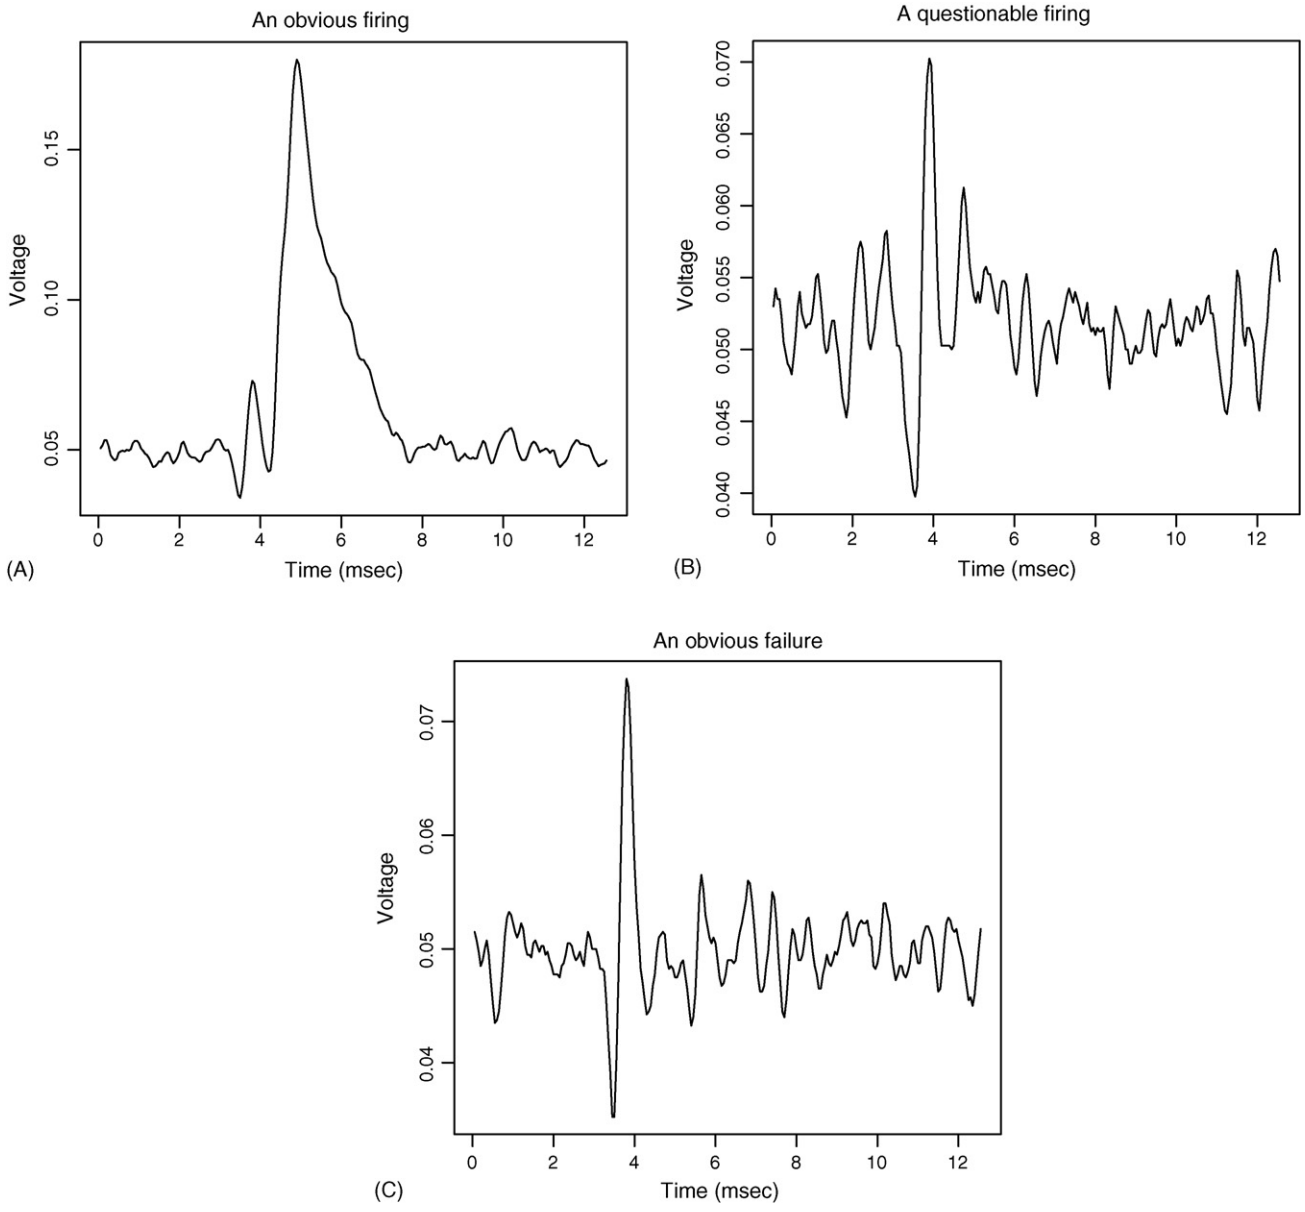

Fig. 3. The classification problem—which are release events? Some traces are obvious release events and some are obvious failures. However, there are always traces "in between". We propose here an objective classification mechanism based on time series results.

1. *dc drift*. Often, as the data from the nerve is recorded, the image received will drift vertically on the oscilloscope. Thus, we first align the curves in the y axis by using trimmed means.
2. *Artifact removal*. The stimulation artifact within the saline bath and spike of the extracellularly recorded action potential need to removed from the trace.
3. *Release event classification*. Each curve is determined to be a release event or failure on the basis of a significance test on the maximum value. This divides the data set into "release events" and "failures".
4. *Removal of doublets*, *triplets*, *etc.* This is still done manually by the investigator from the traces that are classified as "release events". Often these are obvious because of multiple non-artifact peaks in the trace. In the data used here, only one release event was deemed a "doublet" by the investigator (AFMJ). This trace is shown in Fig. 6. While the method for classification to "release/failure" still works if there are many multiple release events in a single trace, the computation of the functionals currently only works for single release events. We do not currently have a computational method for separating a multiple release event into its constituent pieces; however various approaches have been develop to use deconvolution techniques of multi-quantal evoked events at NMJs (Bykhovskaia et al., 1996).

In the remainder of this section we describe each of these tasks in order.

3.2.1. Removing dc shift

The first task is to align all traces vertically to a common baseline of 0. Since there is no standard for reference (i.e., there

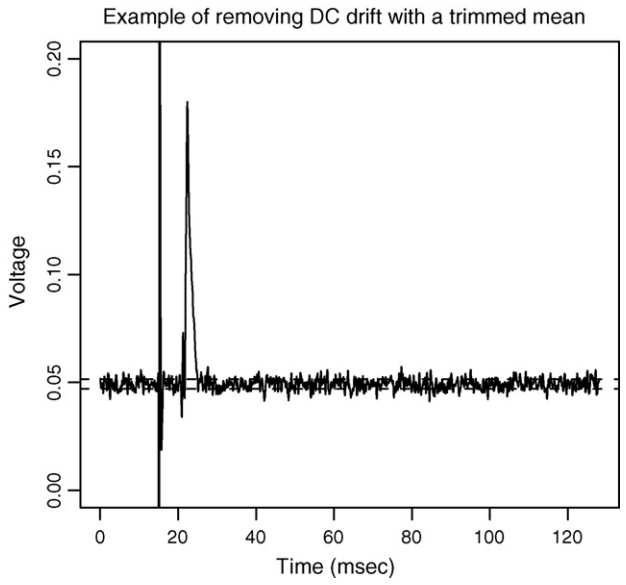

Fig. 4. Example of removing dc drift with a 40% trimmed mean. The trimmed means removes the highest 20% and lowest 20% observations from the trace and averages the remaining observations. For this trace, the observations between the dashed lines are averaged to produce a baseline value for the trace. As with Fig. 1, the stimulus artifact has been clipped for better viewing of the EPSP. The average of the data between the dashed lines is 0.0493, which is then subtracted from every voltage value to give the trace baseline value 0.

is no way to record the dc shift over time during the experiment), we used features inherent to the curve itself in order to obtain the baseline. Within each trace, there is very little movement vertically (only two artifacts and the occasional release event). Since most of the graph is noise, we take a 40% trimmed mean and use that value in order to shift that curve to a baseline value of 0.

In order to calculate a 40% trimmed mean for a single trace, one first takes all 2560 observations in the trace and then sorts them in increasing order. From that, the bottom 20% and top 20% of those observations are removed, and the remaining middle 60% are averaged. Based on 2560 observations, we would only use the sorted values between the 512th and 2048th to calculate the mean.

Fig. 4 shows how this works graphically. For this particular trace, the 512th largest voltage value is 0.047 and the 2048th largest voltage value is 0.0515. These voltage values are shown with dashed lines (red in web version). The data whose voltage values are between the two horizontal dashed lines are then averaged, resulting in a trimmed mean voltage of 0.0493. The trimmed mean is then subtracted from every voltage value in the trace, thereby aligning the trace to a baseline voltage of 0. These lines are difficult to identify on the graph because they are so close to the data, but this is the point. The trimmed mean avoids the areas where there are deviations from baseline and averages the rest. This process is repeated separately for each trace (thus a different value is subtracted from each curve, accounting for the dc drift).

Our accompanying software allows for a different percentage (other than 40%) for the trimmed mean. The central issues to consider in resetting this parameter are: (1) all systematic deviations (firings, artifacts, minis, etc.) should take less than the trim factor of the trace, and (2) ideally several hundred observations should remain after trimming, so the baseline is accurately estimated.

3.2.2. Removing the artifacts

One of the assumptions of our current methodology is that the traces contain one artifact near the region where the EPSPs occur. The data in Fig. 1 clearly contain two artifacts, the artifact near the EPSPs (most visible in Fig. 2) and the large stimulus artifact that has been clipped in Fig. 1. The large stimulus artifact, fortunately, is well separated from the release events. Thus, without adding subjectivity into the analysis, we can simply remove the time values that include this artifact and focus on the area shown in detail in Fig. 2. Note that we removed dc drift before focusing on the area around the release events because the "deleted" area of the traces provided information on dc drift. By focusing on the area emphasized in Fig. 2, we have reduced the data to contain a single artifact. In the code described later, the user may select the time values for the rectangular region. Thus, the user must identify the *x*-axis bounds of the window "A" in Fig. 1. This task is typically not difficult. The left edge occurs between the two artifacts, and the right edge is selected so that "A" contains all the EPSPs.

The remaining artifact must then be removed. Fortunately, in the datasets we are considering the minimum of the curve occurs at the beginning of the artifact, and is then succeeded with an upward and then downward oscillation before we resume either noise or a release event. Thus, to remove the artifact, we find the minimum of the curve and then the next subsequent maximum and minimum. That second minimum is then declared the beginning of the trace for further calculation. The minimum of the curve and the succeeding maximum and minimum are marked in Fig. 5. The area before the "cut before here" is then removed before further calculations.

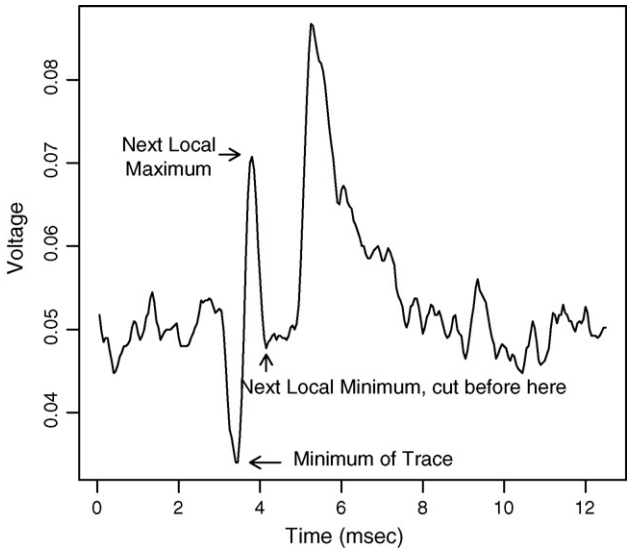

Fig. 5. Demonstration of points used to remove artifacts. The oscillation immediately following the minimum of the trace is removed.

In the code that accompanies this paper, certain other artifact structures are possible. In particular, it is possible to analyze data with no lead artifact.

3.2.3. Determining a cutoff for release versus failure

With the data vertically centered and time-trimmed, we then begin the process of classifying the data into release events and failures. Thus, we construct a statistical hypothesis test of:

H_0. The trace is a failure versus.

H_1. The trace is a release event.

We construct this test based on whether the maximal value of the trace exceeds a threshold T. To find the appropriate value of T, we need the sampling distribution of the maximal value of the trace.

As has been noted by other others in different context (Sacchi et al., 1998), the current traces are composed of dependent data that can be modelled by an autoregressive, moving-average (ARMA) process (Brockwell and Davis, 1991). Specifically, if V_{ij} is the jth voltage value from the ith trace, then the V_{ij} values may be described by the equation:

$$V_{ij} = \sum_{k=1}^{p} \phi_k V_{i,j-k} + Z_{ij} + \sum_{k=1}^{q} \theta_q Z_{i,j-k} \quad (1)$$

where the ϕ_k values are called the autoregressive coefficients, the θ_k values are the moving-average coefficients, and the Z_{ij} values are all jointly independent and each distributed $N(0, \sigma^2)$ (i.e., are white noise). Typically $p = q = 2$ is sufficient for a wide variety of purposes (Brockwell and Davis, 1991). Note in the results section we investigate fitting separate ϕ and θ parameters for each trace, and demonstrate a single set of coefficients simultaneously fits all the traces well.

Thus, we fit the model:

$$V_{ij} = \phi_1 V_{i,j-1} + \phi_2 V_{i,j-2} + Z_{ij} + \theta_1 Z_{i,j-1} + \theta_2 Z_{i,j-2} \quad (2)$$

We use the data in Fig. 1 contained in the rectangle marked "B". This region contains the last 25% of the observations (640 for our sample data) in each trace. This is a sufficient number of observations per trace to acquire accurate estimation of the time series coefficients, and in addition there are few systematic deviations from baseline such as miniature EPSPs (mEPSPs). In the software provided, the user has control over the size of the rectangle in "B".

Estimating the coefficients ϕ_1, ϕ_2, θ_1, θ_2, and innovation variance σ^2 requires numerical methods available in many statistical software packages. We use R, a freeware package, which estimates the coefficients by maximum likelihood (Casella and Berger, 2001). After estimating the coefficients (note we typically have several hundred thousand observations available, so these coefficients can be estimated quite accurately).

For each "artifact-removed" trace isolated in Section 3.2.2, let n be the number of observations in each trace. Given the trace is not an EPSP, the n observations in each trace arise from Eq. (2). To determine the sampling distribution of the maximum, we repeatedly sample from Eq. (2) and record the maximum for each simulation. The resulting maxima provide an estimate of the underlying distribution. After choosing a size (α-level) of the test, our threshold T is the value such that the proportion of simulated maxima who exceed T is α.

Because we are performing a hypothesis test on each trace, we need to address multiple hypothesis testing issues (see for example Benjamini and Hochberg, 1995). We emphasize here, however, that we are not interested in controlling the family-wise error rates, but simply are using the hypothesis tests as a classifier. Thus, our concern is about the rates of correct classification (true release events should be classified as release events, and failures should be classified as failures). Our choice of α is motivated by classification. If we used $\alpha = 0.05$, for example, we would expect 50 out of ever 1000 failures ($\alpha = 0.05$ multipled by 1000 traces) to be incorrectly classified as a release event, which would typically imply a large proportion of the identified release events are actually incorrectly identified. Note α is just the rate of falsely identified traces. Thus, if you want to expect only 1 failure out of 1000 to be falsely identified as a release event, choose $\alpha = 1/1000 = 0.001$. This is our standard choice in practice. Thus, we can be reasonably confident few failures are incorrectly classified as release events.

3.3. Removing multiple events

One aspect that is not automated in our procedure is the removal of doublets, triples, etc., from the traces classified as release events. These are determined by visual inspection, where a trace was classified as a multiple release if it contains multiple non-artifact peaks in voltage. In this dataset, only one doublets was found. It is shown in Fig. 6. While this process is not automated, the investigator needs to only look at the traces classified as release events, rather than the entire

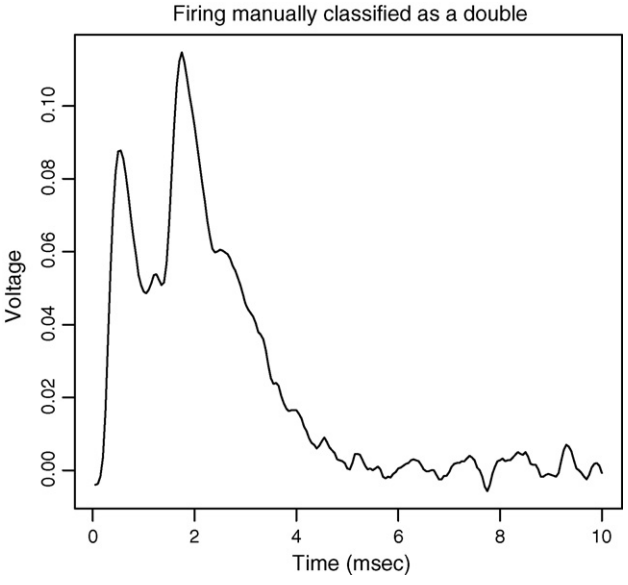

Fig. 6. This trace was classified as a firing, but was manually determined to be a multiple firing and not included in the subsequent calculation of the functionals.

set of traces as was previously required. Typically the number of release events is substantially smaller than the number of traces.

3.4. Computing functions (area, rise time, etc.) for each release event

Computation of the area under the curve, the time to peak, the decay rate, and other functionals is very time consuming to determine by hand even with current software packages marketed since the cursors still have to be placed by hand at the beginning and end of each event. The automated procedure described herein eliminates the placement of cursors for every quantal occurrence.

Once the traces are classified, we calculate several functionals for each of the traces classified as release events. These are area under the curve, peak of the curve, time to curve peak (we provide two versions, either latency based or rise based), and the decay time. This section describes the computational details involved in each functional. For each, assume that the second artifact has been removed, so the first time point represents the minimum shown in Fig. 5. Let the voltage values observed, in order, be $v_1, \ldots, v_n$ (the code allows n to be specified by the user).

1. *Area under the curve* (*AUC*). We use the trapezoidal rule, where
$$\text{AUC} = \frac{1}{2\Delta}(v_1 + 2v_2 + 2v_3 + \cdots + 2v_{n-2} + 2v_{n-1} + v_n)$$
where Δ is the common width between two successive time points. The trapezoidal rule essentially "connects the dots" of the voltage values and computes the resulting area. Note that when values are observed with error, there is little to be gained by using Simpson's rule (as opposed to when function values are known exactly, when Simpson's rule is superior). Determining the start and end of the firing cannot be done exactly because of the electrical noise. We provide a user specifiable parameter auc.threshold which determines how many standard deviations above baseline are required to determine these points. The beginning and end of the event are defined as the points, moving away from the peak in each direction, where the adjusted (versus baseline) voltage first falls beneath auc.threshold standard deviations. The default value of auc.threshold is 0, indicating the event begins and ends when the voltage returns to baseline in each direction.
2. *Peak of the curve*. This is simply the highest voltage value observed in $v_1, \ldots, v_n$.
3. *Time to curve peak*. This is the time between the "start of the curve" and the time the maximum is achieved for the first time (typically the maximum is achieved only once, but in some cases the same maximum appears in successive time points, for example).

   We implemented two different definitions of the "start of the curve". This is due to the fact that while many release events occur immediately after the second artifact, some release events have a "latency" period of a few milliseconds. The two definitions are based on whether or not to include the latency as part of the "time to curve peak".

   Unfortunately, it is not possible to identify exactly where a release event begins, because of the random noise associated with the recording. We estimate the beginning of the release event by working backward from the peak of the release event until the first time point where the trace falls below $2s$, where s is the sample standard deviation of the data in region "B" of Fig. 1. We call the time period between the estimated start of the release event and the peak as the "rise time". It is shown as a dotted line in Fig. 7 (note we are measuring time here, so only the x-axis difference is measured). Horizontal lines in Fig. 7 show the maximum and two standard deviations cutoffs. We call the period before the estimated start of the release event, beginning at the end of the second artifact, the "latency time". The latency time is shown as a dashed line in Fig. 7.

   Our functional "time to curve peak" is thus calculated in two ways. The first way is to include both the latency and rise periods together, producing the time from the end of the second artifact to the peak of the release event. The second way only includes the rise time.
4. *Decay time*. This is the time elapsed between the maximum (again, the first maximum in case of ties), and where the descent first falls below $p \times \max$, where p is the percent decay and 'max' is the maximum value. The user may choose p, by default we choose $p = 0.37$ as used in prior studies (Hille, 1992). Fig. 7 contains a horizontal line at 37% of the release event maximal voltage and shows the decay period as a dashed and dotted line (green in web version). The time elapsed in the decay period is the decay time, commonly referred to as τ.

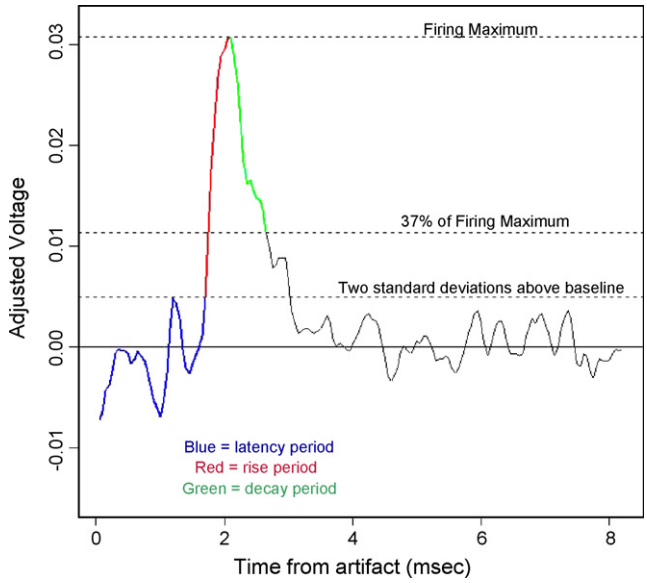

Fig. 7. Components of a release event. Here, we define latency as the time between the beginning of the trace and the first point where the trace exceed two standard above baseline. The rise time is the time following the latency period until the maximum, and the decay time is the time from the maximum until the trace falls below 37% of the maximum.

4. Results

4.1. Time series results

In Section 3.2.3 we use only one set of ϕ and θ values, rather than separate parameters for each trace. In this section we illustrate that this is consistent with the data.

We make this argument on two grounds, first through a formal hypothesis test and next through an exploratory study indicating the sampling distribution of individual time series coefficients follows the pattern that would be expected from a common value.

First, expand Eq. (2) to allow for the possibility of each trace having its own set of time series coefficients ϕ_{i1}, ϕ_{i2}, $\theta_{i1}\theta_{i2}$, so the model is now

$$V_{ij} = \phi_{i1}V_{i,j-1} + \phi_{i2}V_{i,j-2} + Z_{ij} + \theta_{i1}Z_{i,j-1} + \theta_{i2}Z_{i,j-2}$$

Letting s be the number of traces, a formal hypothesis test can then be conducted by testing:

H_0.

$$\phi_{11} = \cdots = \phi_{s1}, \qquad \phi_{12} = \cdots = \phi_{s2},$$
$$\theta_{11} = \cdots = \theta_{s1}, \qquad \theta_{12} = \cdots = \theta_{s2}.$$

H_1. All coefficients are separate.

Because of the large number of observations, we use Bayesian information criteria (BIC) for this test (Kass and Raftery, 1995). Maximizing the likelihood under each hypothesis and computing the BIC, we find the approximate log Bayes factor to be −17,900, which is overwhelming evidence in favor of H_0. Thus, the hypothesis test concludes that a single set of coefficients is strongly preferred over separate coefficients.

Second, suppose $H_0$is true and a common set of coefficients exists. Under this assumption, we estimated the coefficients using all the data and the associated information matrix (Casella and Berger, 2001). Using these, we can derive out the sampling distribution of estimates derived from each individual trace. In Fig. 8, these are shown as the elliptical contours. The left pane shows the autoregressive (AR) coefficients, while the right pane shows the moving-average (MA) coefficients. There are three ellipses in each pane, showing 90%, 95%, and 99% contours.

The points in each pane are the individual estimates arrived at by fitting a separate set of coefficients for each trace. While there is certainly variation about the individual estimates (estimates differ from trace to trace), the pattern of variation is quite consistent with the expected sampling variability from the common estimate. Thus, we conclude there is a single set of coefficients, and the differences observed from trace to trace are normal sampling variability.

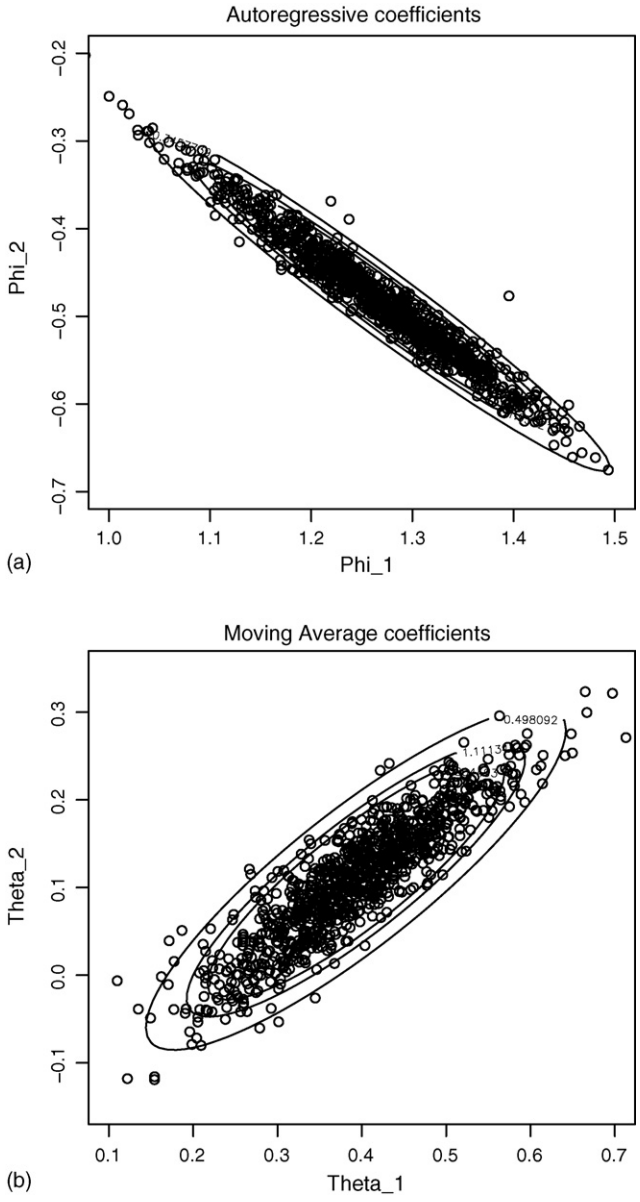

Fig. 8. Demonstration that individual estimates for each trace follow the sampling distribution expected from a common estimate. The ellipses correspond to the sampling distribution expected from the common estimate. The points correspond to estimates of the time series component from each individual trace.

4.2. Sample run through the software

4.2.1. Getting started

The routines herein are implemented in a statistical language called "S" using the freeware software "R". "R" is maintained by a consortium of statisticians and others, and may be downloaded at: http://www.r-project.org/. "R" is available in all common operating systems. We avoid a detailed "user guide" here, instead emphasizing features and the effect of user controlled parameters. The code, with instructions both specific to the methods here and on R in general, may be found at: http://www.ms.uky.edu/~viele/epsps/epsps.html and consists of the following files:

1. traces.RData: a shortcut that allows you to start R with all the required functions preloaded.
2. traces.txt: contains the functions used to process the traces. This file is not necessary unless you intend on modifying the code. All the code in this file is preloaded into traces.RData.
3. output05g.txt: contains the sample dataset used in this paper.

Data may be loaded from either a single column file, where the voltages for each traces follow one after the other (some software refers to this as a "scope" file), or from a matrix file. The software provides functions which plot the raw data either together, as in Fig. 1, or separately. The software provides options for zooming in on particular areas of the traces, as in Fig. 2 or 4.

4.2.2. An "all in one" method

The software provides an all-inclusive function *analyze.traces* which allows the user to remove artifacts, adjust for dc drift, find a threshold for classification into firing/failure, classify traces, and compute functionals in one step. Thus, with one command one can perform all the functions described below, and still retain the ability to change any of the user specifiable parameters described below.

Similarly, while separate functions are provided for each step described in Section 3, we have provided "combination" functions that simultaneously adjust for baseline drift, remove the artifact, and so on. The intent of the separate functions is to allow for future enhancements and give a user wishing to perform their own analysis freedom to do so, while the combined commands save time for the user who wants a fully automated analysis. In what follows we describe the separate functions which allow us to focus on each of the user specifiable options.

4.2.3. Baseline adjustment

The function *set.baseline* performs the trimmed mean adjustment discussed in Section 3.2.1. The user is given the option of changing the "trim" parameter (the default is 0.2, which removes a total of 40% of the data). The result of the function is an adjusted trace matrix which removes the dc drift.

For our sample data in Fig. 1, there is a fairly small amount of dc drift throughout the experiment. Fig. 9 shows the before and after results of the setbaseline function. The central result of using the function is that the average voltage has been shifted to 0.

4.2.4. Setting a threshold for classification

The function *get.threshold* determines the threshold value for the classification of firing versus failure as described in Section 3.2.3. This function is by far the largest portion of the processing time. The function estimates the appropriate time series coefficients, uses them to estimate the sampling distribution of the maximal voltage value, and then determines the appropriate threshold. In addition, the function estimates the variance of the noise to be used in computing the functionals.

The user has control over the options (among others, such as how long to run simulations):

1. *Noise*. Determines how much of the tail of each trace is assumed to be noise and therefore used to estimate the time series coefficients. The default option is 0.25, which uses 25% of the data for each trace. Ideally, this option should be set as high as possible, while still being sure few firings appear in the "noise" area, to guarantee accurate estimates.

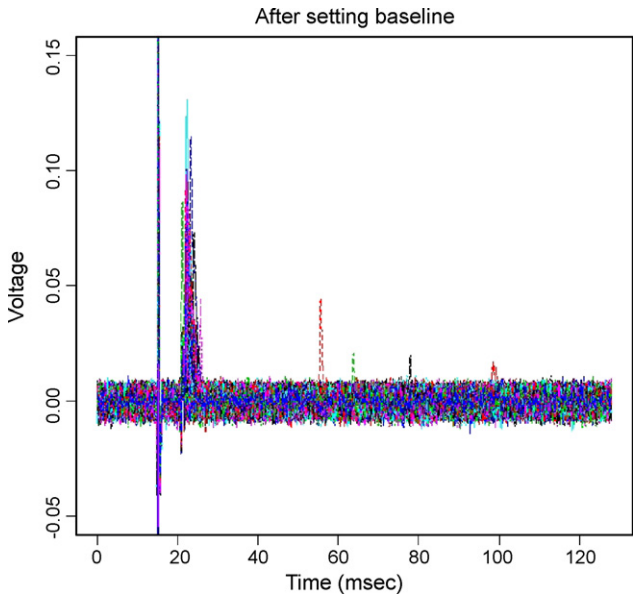

Fig. 9. The results of the baseline adjustment on the sample data. In this example there was little dc drift, so the traces have not been shifted much relative to each other. The main effect of setting the baseline has been to center the traces on a baseline voltage of 0.

2. *Alpha*. The probability of type I error for the classification of any particular trace. May be set anywhere between 0 and 1, although the default is the reciprocal of the number of traces as described in Section 3.2.3.
3. *Keep*. How many observations will be contained in each of the traces after removing lead artifacts. It is crucial this parameter agrees with what is used in the *remove.artifact* function below. The reason is that the distribution of the maximal voltage depends on how many voltage values are present in each traces (more observations results in a larger probability of a "chance" high voltage). In the "all in one" functions, the agreement in the arguments of the *get.threshold* and *remove.artifact* functions is handled automatically.

This function should be used with the data after the baseline adjustment but before removing artifacts. The point is that the baseline adjustment removes variation in the voltage axis, but we want to use the extreme tails of the distribution (the area away from the events) to estimate the noise parameters. These tails are removed by the *remove.artifact* function.

For our sample dataset, the threshold for distinguishing firing from failure is an adjusted (above baseline) voltage of 0.0105. We include this value on the plot in the next section.

4.2.5. Removing artifacts

The function *remove.artifact* removes artifacts from the beginning of the trace. The default options involve the artifact being removed as described in Section 3.2.2 and Fig. 5. The user may select the following options:

1. *Remove.data*. Allows the data to remove a number of observations from the beginning of each trace before searching for artifacts (we use this to remove the large spike from Fig. 1).

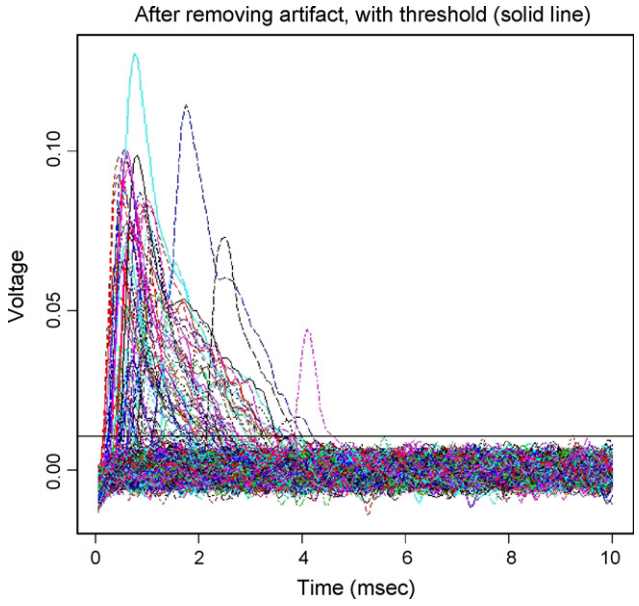

Fig. 10. The traces after the removal of the artifact and the threshold for classification. This graph is similar to Fig. 2 with the exception of the removal of the artifact. Should there be variation in the timing of the artifact, the methods described here will also provide alignment in the time axis. The dashed line indicates the threshold failure for classification into a failure or firing. Traces whose maximal value exceeds the dashed line are classified as firings.

2. *Keep*. The number of observations to include in each processed trace. This allows for experiments where events have differing lengths or account for the acquisition rate.
3. *Artifact.form*. Specifies the form of the artifact. The user may specify no lead artifact, in which case the function simply supplies the first "keep" observations in each trace. The default is "hl", which removes an artifact as in Fig. 5.

The function returns two pieces, the artifact-removed data and a vector indicating any problems encountered. The software notes whenever it has trouble removing an artifact, whenever an insufficient amount of data is available, or when no pulse is detected (indicating a recording problem).

The results of using *remove.artifact* on our sample data (after baseline adjustment) are shown in Fig. 10. We used remove.data = 350 (a value past the large spike but before the events), keep = 200, and the default artifact.form. The result isolates the area where the events occur. The results are shown in Fig. 10.

4.2.6. Classifying traces

The *classify.traces* function simply looks at each of the preprocessed traces and determines if the maximal value of the trace exceeds the threshold for classifying the trace as a firing. The function also produces the plot in Fig. 11 that plots the set of traces classified as failures and the set of traces classified as firings.

4.2.7. Computing functionals

The function *get.functions* computes the functionals described in Section 3.4 for all the traces classified as firings. The user can/must supply:

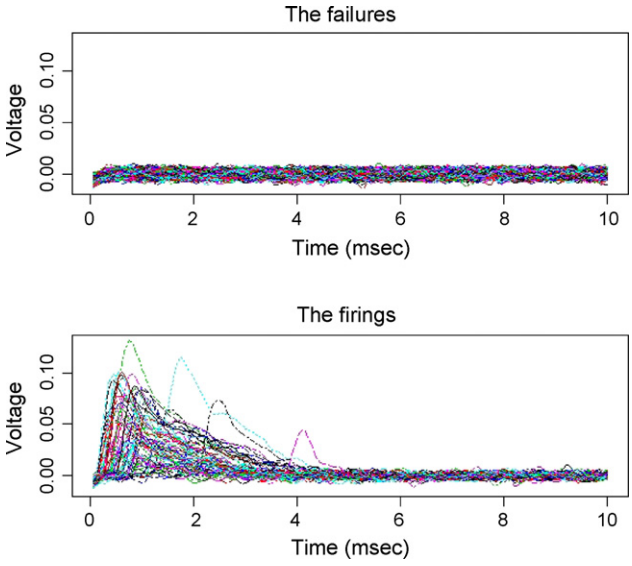

Fig. 11. The classification of the traces into failures and firings. This graph splits the data in Fig. 10 into failures (the top pane) and firings (the bottom pane). The classification is based on whether the maximal value of the trace exceeds the threshold value drawn in 10.

1. The firings in a matrix (the "all in one" functions handle this automatically).
2. The noise standard deviation produced by the *get.threshold* function.
3. The kHz rate at which the recording where computed. The function returns all times in milliseconds, and the acquisition rate is required to convert numbers of observations to time.
4. The amount of decay to be used in determining the decay time (the default is 0.37, as in Hille, 1992). See Section 3.4 for more details.
5. A parameter *auc.threshold* which determines the endpoints of the event for the purposes of computing the area under the curve (AUC). The endpoints are determined by finding the first points, moving away from the peak in each direction, where the trace first falls below auc.threshold times the noise standard deviation. Again, see Section 3.4 for more details.

The result of the *get.functions* command is a matrix containing the functionals. The software automatically produces histograms of all the functionals computed. These are shown for the sample data in Fig. 12.

5. Discussion

There are various commercially available computational methods to perform analysis of quantal events. However some of these packages have come and gone as they have not been maintained such as ones by Synaptosoft, Inc. that quit functioning in 2002 but still advertises on the web. The most commonly used analysis packaged is most likely pCLAMP 10 by Axon Instruments (Molecular Devices Corporation, California, USA) for analysis of evoked events. The shortfall of this program is the manual placement of cursors for events and baseline shifts that need to be controlled for in the data analysis. As with the method proposed in this paper, single and multiple evoked quantal re-

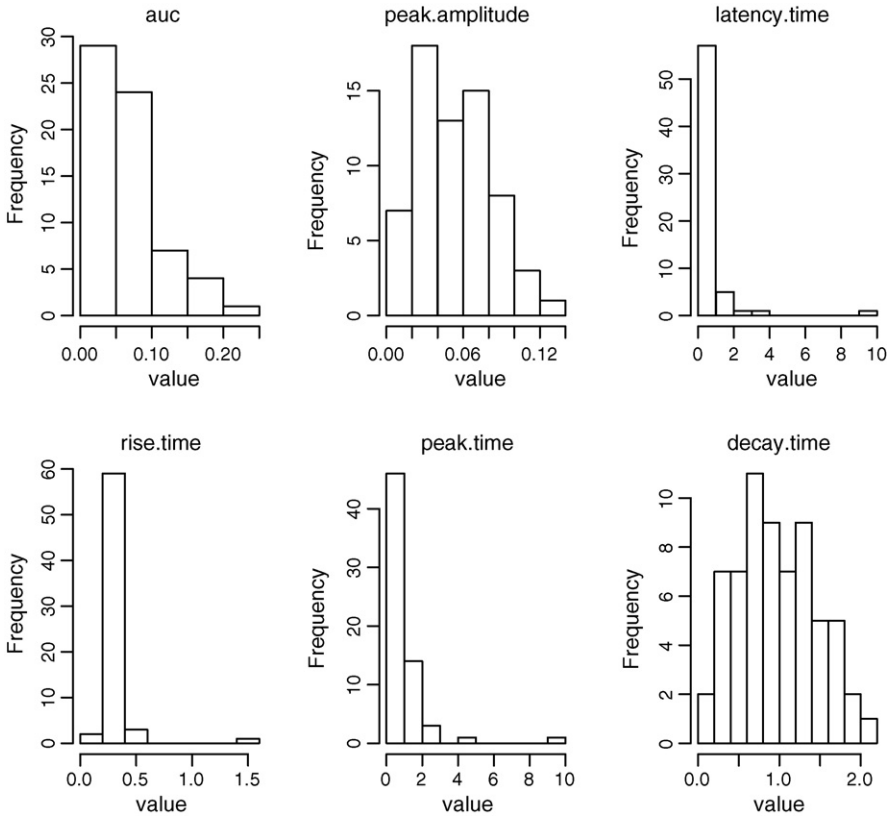

Fig. 12. Histograms of the six functionals computed on each of the firings. These observations are stored in a matrix and available for any future analyses.

sponses would also need to be visually inspected for determining single quantal responses for analysis. We feel that the automated procedure presented in this study will save time in not having to place cursors on events nor worry about adjusting baseline offset throughout an experimental run. In addition, the software "R" is freely obtained and maintained by a large number of statisticians dedicated to its improvement.

In this study we provide a procedure that automates the analysis of evoked single quantal events to provide characteristics of quanta. This provides a means to quantify any changes that affect the quantal responses experimentally. Regulatory processes that may alter presynaptic and/or postsynaptic properties can readily be resolved by the multiple measures in the quantal analysis presented. There are possible actions on postsynaptic receptor array such as the presence of antagonists or numbers of desensitized receptors that would decrease the peak amplitude and area of the quantal event without alteration of the decay tau, latency of occurrence or the number of occurrences (Nicoll and Malenka, 1999; Tang et al., 1994). Any difference in the postsynaptic density of receptors due to developmental differences (DiAntonio et al., 1999; Qin et al., 2005) or activity induced changes as observed during long-term potentiation in the CNS of vertebrates (Nicoll and Malenka, 1999) would also show such differences in quantal measures. Synaptic structural dimension in relation to the presynaptic vesicular fusion site can also impact the shape of a quantum (Bekkers and Stevens, 1990, 1991; Uteshev and Pennefather, 1997).

Factors that target presynaptic mechanisms and alter quantal size are vesicle packaging (Sulzer and Edwards, 2000; Wilson, 1998) or size differences of vesicles (Karunanithi et al., 2002; Kim et al., 2000). Many cellular processes such as degree of phosphorylation (Silverman-Gavrila et al., 2005) and handling of evoked calcium influx (Cooper et al., 1996a,b; Dawson-Scully et al., 2000; Millar et al., 2005) effect the number and latency of vesicular events. At the crayfish NMJ, the presence of serotonin is thought to lead to increased phosphorylation of synaptic proteins that increase the number of docked vesicles, thus enhancing evoked release and shortening latency of release (Cooper et al., 2003; Southard et al., 2000). It would be of interest to use this presented automated analysis to investigate the time dependent effect of action by neuromodulators such as serotonin. Previously we used the area measures of quantal analysis (Viele et al., 2003) to investigate clusters in occurrences in subsets of quantal events and estimate a probability for the various types of occurrences.

When the stimulation frequency was increased from 1 to 2 and 3 Hz the probability increased for particular subsets of events as well as new subsets of events appearing. This suggested that sites initially activated, which produce a given subset of quantal charges, increase in their occurrence, and that novel sites can also be recruited upon increased stimulation frequency. Automation of quantal analysis could readily enhance such measures in characterizing subsets of quantal responses.

At the crayfish NMJ the glutamatergic ligand gated receptor are a qusiqualate type with rapid sodium conductance (Shinozake and Shibuya, 1974) which are similar as those at the neuromuscular junction in the genetically favorable model *Drosophila melanogaster* (Bhatt and Cooper, 2005). We are now

examining various mutants in *Drosophila* lines that have been identified to produce alterations in synaptic strength to determine the mechanistic reasons by this automated high throughput quantal analysis. In addition, this analysis will provide new approaches to address quantal subsets to better described the number of functional release sites (n) and the probability of release (p) in not only crayfish NMJs, but all synapses that allow quantal measures to be obtained.

Acknowledgements

Funding was provided by NSF-IBN-0131459 (RLC, ML and KV) and a G. Ribble Fellowship for graduate studies in the School of Biological Sciences at the University of Kentucky (AFMJ).

References

Bekkers JM, Stevens CF. Presynaptic mechanism for long-term potentiation in the hippocampus. Nature 1990;346:724–29.

Bekkers JM, Stevens CF. Application of quantal analysis to the study of long term potentiation: errors, assumptions, and precautions. In: Baudry M, Davis JL, editors. Long term potentiation. A debate of current issues. Cambridge, MA: MIT Press; 1991. p. 63–76.

Benjamini Y, Hochberg Y. Controlling the false discovery rate: a practical and powerful approach to multiple testing. J Roy Stat Soc Ser B 1995;57:289–300.

Bhatt D, Cooper RL. The pharmacological and physiological profile of glutamate receptors at the Drosophila larval neuromuscular junction. Physiol Entomol 2005;30:1–6.

Brockwell P, Davis R. Time series: theory and methods. 2nd ed. New York: Springer; 1991.

Bykhovskaia M, Worden MK, Hackett JT. An algorithm for high-resolution detection of postsynaptic quantal events in extracellular records. J Neurosci Meth 1996;65(2):173–82.

Casella G, Berger R. Statistical inference. 2nd ed. New York: Duxbury Press; 2001.

Cooper RL, Donmezer A, Shearer J. Intrinsic differences in sensitivity to 5-HT between high- and low-output terminals innervating the same target. Neurosci Res 2003;45:163–72.

Cooper RL, Harrington C, Marin L, Atwood HL. Quantal release at visualized terminals of crayfish motor axon: intraterminal and regional differences. J Comp Neurol 1996a;375:583–600.

Cooper RL, Marin L, Atwood HL. Synaptic differentiation of a single motor neuron: conjoint definition of transmitter release, presynaptic calcium signals, and ultrastructure. J Neurosci 1995;15:4209–22.

Cooper RL, Winslow J, Govind CK, Atwood HL. Synaptic structural complexity as a factor enhancing probability of calcium-mediated transmitter release. J Neurophysiol 1996b;75:2451–66.

Dawson-Scully K, Bronk P, Atwood HL, Zinsmaier KE. Cystein-string protein increases the calcium sensitivity of neurotransmitter exocytosis in Drosophila. J Neurosci 2000;20:6039–47.

DiAntonio A, Petersen SA, Heckmann M, Goodman CS. Glutamate receptor expression regulates quantal size and quantal content at the Drosophila neuromuscular junction. J Neurosci 1999;19:3023–32.

Del Castillo J, Katz B. Quantal components of the end-plate potential. J Physiol 1954a;124:560–73.

Del Castillo J, Katz B. Statistical factors involved in neuromuscular facilitation and depression. J Physiol 1954b;124:574–85.

Dudel J. The effect of reduced calcium on quantal unit current and release at the crayfish neuromuscular junction. Pfugers Arch 1981;391:35–40.

Faber DS, Korn H, Redman SJ, Thompson SM, Altman JS. Central synapses: quantal mechanisms and plasticity. Strasbourg: Human Frontier Science Program; 1998.

Ginsborg BL. Electrical changes in the membrane in junctional transmission. Biochim Biophys Acta 1973;300(3):289–317.

Hille B. Ionic channels of excitable membranes. 2nd ed. Sunderland, MA, USA: Sinauer Associates, Inc.; 1992.

Karunanithi S, Marin L, Wong K, Atwood HL. Quantal size and variation determined by vesicle size in normal and mutant Drosophila glutamatergic synapses. J Neurosci 2002;22:10267–76.

Kass R, Raftery A. Bayes factors. J Am Stat Assoc 1995;90:773.

Kim S, Atwood HL, Cooper RL. What are the real sizes of synaptic vesicles in nerve terminals. Brain Res 2000;877:209–17.

Magrassi L, Purves D, Lichtman JW. Fluorescent probes that stain living nerve terminals. J Neurosci 1987;7:1207–14.

Martin AR. A further study of the statistical composition on the end-plate potential. J Physiol 1955;130(1):114–22.

McLachlan EM. The statistics of transmitter release at chemical synapses. In: Porter R, editor. International review of physiology, neurophysiology III, vol. 17. Baltimore: University Park Press; 1978. p. 49–117.

Millar AG, Zucker RS, Ellis-Davies GCR, Charlton MP, Atwood HL. Calcium sensitivity of neurotransmitter release differs at phasic and tonic synapses. J Neurosci 2005;25:3113–25.

Nicoll RA, Malenka RC. Expression mechanisms underlying NMDA receptor-dependent long-term potentiation. Ann N Y Acad Sci 1999;868:515–25.

Olkin I, Petkau AJ, Zidek JV. A comparison of n estimators for the binomial distribution. J Am Stat Assoc 1981;76:637–42.

Qin G, Schwarz T, Kittel RJ, Schmid A, Rasse TM, Kappei D, et al. Four different subunits are essential for expressing the synaptic glutamate receptor at neuromuscular junctions of Drosophila. J Neurosci 2005;25:3209–18.

R Development Core Team. R: a language and environment for statistical computing. Vienna, Austria: R Foundation for Statistical Computing; 2004. ISBN 3-900051-07-0, URL: http://www.R-project.org.

Sacchi O, Rossi ML, Canella R, Fesce R. Synaptic current at the rat ganglionic synapse and its interactions with the neuronal voltage-dependent currents. J Neurophysiol 1998;79:727–42.

Shinozake H, Shibuya I. New potent excitant, quisqualate acid: effects on the crayfish neuromuscular junction. Neuropharmacol 1974;13:665–72.

Silver RA. Estimation of nonuniform quantal parameters with multiple-probability fluctuation analysis: theory, application and limitations. J Neurosci Meth 2003;130(2):127–41.

Silverman-Gavrila LB, Orth PMR, Charlton MP. Phosphorylation-dependent low-frequency depression at phasic synapses of a crayfish motoneuron. J Neurosci 2005;25:3168–80.

Southard RC, Haggard J, Crider ME, Whiteheart SW, Cooper RL. Influence of serotonin on the kinetics of vesicular release. Brain Res 2000;871:16–28.

Sulzer D, Edwards R. Vesicles: equal in neurotransmitter concentration but not in volume. Neuron 2000;28:5–6.

Tang C-M, Margulis M, Shi Q-Y, Fielding A. Saturation of postsynaptic glutamate receptors after quantal release of transmitter. Neuron 1994;13:1385–93.

Uteshev VV, Pennefather PS. Analytical description of the activation of multi-state receptors by continuous neurotransmitter signals at brain synapses. Biophys J 1997;72:1127–34.

Viele K, Stromberg A, Cooper RL. Determining the number of release sites within the nerve terminal by statistical analysis of synaptic current characteristics. Synapse 2003;47:15–25.

Wilson M. The possible origin of variability in miniature PSC amplitude in cultured amacrine neurons. In: Faber DS, Korn H, Redman SJ, Thompson SM, Altman JS, editors. Central synapses: quantal mechanisms and plasticity. Strasbourg: Human Frontier Science Program; 1998. 99–108.