Bayesian Modelling of Patch-clamp Records

J.A.Stark, W.J.Fitzgerald and S.B.Hladky
Departments of Pharmacology and Engineering,
University of Cambridge, CB2 1QJ, U.K.

Abstract

Efficient detection of opening and closing events in recordings of small numbers of ion channels is difficult. The identification of such events is an example of the more general statistical problem of the detection of changepoints in a time series. The simplest methods for tackling this involve filtering and thresholding, but these perform badly when current levels are closely spaced, when the intervals between events are short, or when there is substantial noise. More complex methods such as hidden Markov modelling have been successfully used in the analysis of patch-clamp recordings, but these impose quite detailed assumptions about the underlying physical process and have, so far, only been implemented using the assumption that successive data points are independent events, i.e. that the data has not been filtered.

We have adopted an alternative strategy. The signal is modelled as a sequence of filtered transitions between constant levels with no further assumptions about the process underlying the transitions. The investigator must specify the spectrum of the noise, the characteristics of the data acquisition filter, and initial guesses (strictly prior distributions) of the frequency of events and their magnitudes. This information is sufficient for a Bayesian statistical analysis of the data. Locations of changepoints and currents are chosen using Markov chain Monte Carlo sampling of their posterior distribution. By fitting this type of model to the data one achieves a better trade-off between sensitivity and reliability of detection compared to thresholding methods. Examples will be displayed of estimation of closely spaced current levels and the analysis of records with flicker type events.

Introduction

This research concerns the detection of events in patch-clamp records where the number of channels is small. This is a task of recovering an underlying process from noisy and filtered observations. It is also important to have estimates of the accuracy of the results. There are, broadly speaking, two approaches that have been taken to tackle this problem;

Filtering and thresholding: This method is simple and, perhaps surprisingly, often quite effective. The process is taken to remain in the same state between crossings of the signal and specified thresholds.

Hidden Markov models: This is a much more complex technique, and requires the experimenter to specify the structure of the process in detail. It is fairly slow, and cannot easily deal with filtered records.

We have taken an alternative approach in which a simple model is fitted to the data using Bayesian numerical methods. It makes better use of the information in the data than thresholding, and does not impose detailed structure on the underlying process.

Filtering and Thresholding

The technique of filtering and thresholding (F&T) is simple and fast, but demands that the morphology of the data be clear. In particular, the amount of noise must be small relative to the differences in current levels, and the time between events must be large enough with respect to the signal bandwidth and noise. The technique proceeds as follows:

Filter: A filter is applied to reduce noise.

Threshold: The points at which the filtered signal crosses predetermined thresholds are found, and these are taken as estimates of the locations of events.

Average: The current levels between the events are estimated, generally by averaging.

Example 1

The first example, shown below, is a synthesised signal of 500 points notionally sampled at 100kHz, Bessel filtered (before sampling, 4 poles) at 20kHz, and with a uniform noise spectrum up to 25kHz. There are four current levels, and the true underlying process is shown before and after filtering, and this is superimposed on the observed signal.

Example signal 1

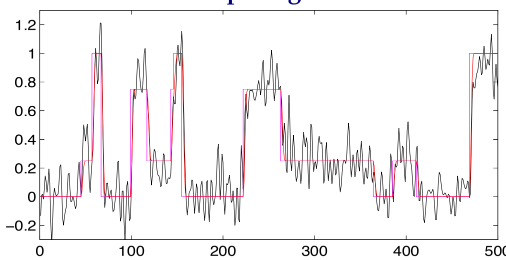

Results of thresholding

The example signal was filtered using 8-pole Bessel filters with different cutoff frequencies. The result was shifted in time to compensate for the overall delay of the filters and thresholded at levels halfway between those in the true underlying process. We assumed such knowledge to give this technique the best chance. The *original data points* were averaged between events, discarding the first three unless the duration was short.

The three figures show the results with estimated levels superimposed on the filtered signals when the lowpass filter cutoff was set to 10, 7.14 and 5.56kHz in turn. The middle value was chosen as a reasonable compromise. Notice that the first two have falsely detected events at around the point 190, whereas the last two miss the short event just before point 150.

F&T analysis of example 1: 10kHz cutoff

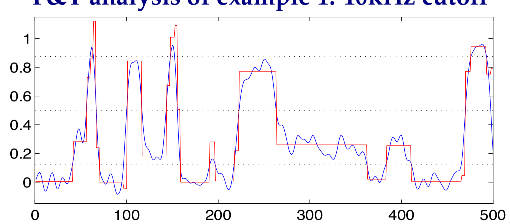

7.14kHz cutoff

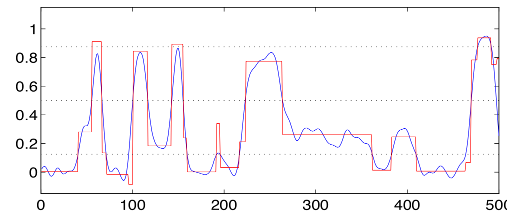

5.56kHz cutoff

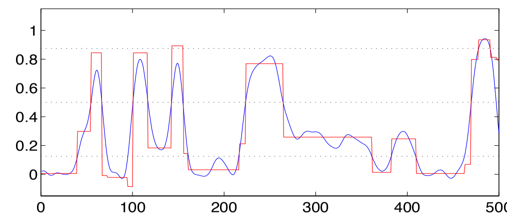

Assessment

Filtering and thresholding has become the standard method for analysing patch-clamp records. It is fast, and it is effective when the underlying process is clearly discernible in the observations. There are significant drawbacks, however, as illustrated in this example.

- The threshold levels have to be supplied; this is quite specific information, and the results can be sensitive to small variations in it.
- The filter bandwidth is problematic, since there is no obviously optimal value. If it is too high then there are too many 'false alarm' events, whereas if it is too low then events will be missed. The example shows how both of these problems can occur together with a compromise value.
- The technique makes very poor use of information in the signal. When it works well, it does so because the information is clear. As soon as ambiguities arise it begins to fail.

In this example the technique was implemented in quite a basic form. There are many enhancements that can be made [1]. In particular, it is possible to eliminate short events and to estimate the levels more intelligently. It is easy to criticise many of these as being 'ad-hockeries'. For instance, there may well be very short events in the data.

Hidden Markov Models

A radically different approach to the analysis of ion channel records [2] employs Markov models (HMMs).

- A highly structured model is proposed for the data. This assumes that the underlying process moves between a number of states, and that the probability of it being in any state at one time step is solely dependent on the state at the previous step. The possible transitions form a *transition probability matrix*.
- There are an infinite number of model configurations, comprising the pattern of events, the transition matrix, and the set of current levels for each state. A function is defined that describes the likelihood of any configuration. An algorithm is applied that attempts to find the one that maximises the probability.
- If the model accurately reflects the nature of the process, HMMs make very effective use of the information in the signal. They can allow for short events, and can account for the possibility of aggregate states.
- The algorithm for fitting the model is quite slow. Currently there are no efficient methods for dealing with filtered transitions, and hence it is commonly assumed that the data is unfiltered. This can seriously affect the interpretation of short events. This form of analysis also imposes a very specific type of model. If this does not model the process well, or if the nature of the process is unknown, then it would be inappropriate. To some extent the results can look unduly convincing: such models are, in a sense, 'self-fulfilling prophecies'.
- There are currently no rigorous statistical methods for selecting the number of states in the model.

Objectives

In attempting to develop an alternative technique, a number of features were sought:

- Efficient use of information in the signal.
- Compensation for filtering.
- Indications of the accuracy of reconstruction.
- Scalability in bandwidth. With F&T the signals are lowpass filtered which means that higher frequency information is lost. It is better for a technique to be robust to such noise and to leave it in.
- Meaningful tuning parameters. Any form of analysis will require the experimenter to specify some parameters. It is better to ask for an approximate rate of events than a filter bandwidth.
- Robustness. The results should not change discontinuously when tuning parameters are varied.

Bayesian Modelling

The problem of analysing ion-channel records is an example of the statistical task of changepoint analysis of time series. We have explored the use of a simple model, and analysed it using Bayesian inference[3]. Aspects of this are:

- The formulation of a formal statistical model, as with HMMs, but simpler. The parameters of the model are the number and locations of the change-points, and the current levels between them. The *likelihood* function relates the observations to the parameters.
- The likelihood is calculated from the difference between the recorded data and the idealized trace after it has been filtered in the same manner.
- Additional parameters are the rate of transitions and variance of the changes in current. Estimates of these are supplied and formulated as priors in the model.
- The posterior distribution is formulated from the likelihood and priors.

Sampling from the posterior

The posterior distribution is a function with many parameters. Finding the maximum value of such multivariate distributions is difficult and unreliable. Furthermore, knowledge of the best configuration of the parameters does not provide any indication of the reliability of the estimate, nor does it give any comparison with alternatives. An alternative to maximization, which is the subject of much current statistical research, involves drawing samples from the distribution using 'Markov chain Monte Carlo' techniques. The probability of a model configuration being sampled is proportional to its quality of fit as given by the posterior.

Bayesian analysis of example 1: Family of samples

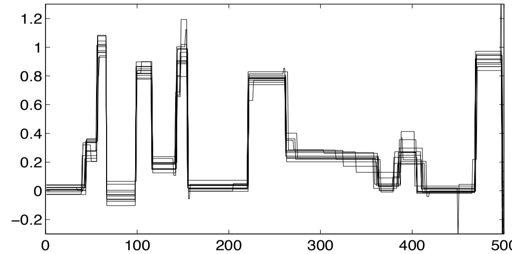

Grand average over 50 000 samples

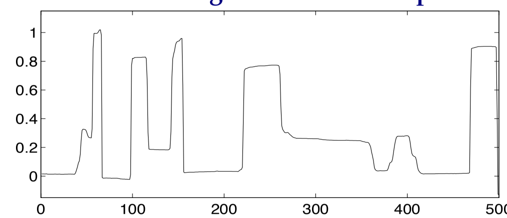

Results of sampling from the model

The two graphs illustrate some of the key features of this form of analysis. The results should be viewed in conjunction with the original data. For example, around the 50th point, the average of every 100 samples out of $50\,000$ shows a small drop. Examination of the samples (12 of which are shown as a family) reveals that there is a certain probability of an extra event here, in which case the separate levels are displaced. The probability associated with this event can be estimated from the frequency with which it occurs over the set of samples. Throughout the plot the variability of estimates of the levels can be seen. Some information, such as the interdependence between the estimates of neighbouring levels, is unclear from the family plot. Also, the number of samples that can be usefully displayed together is limited. This problem, of displaying the results of the sampler, is the subject of current research.

Closely-spaced Levels

The second example is one where some current levels are close with respect to the noise variance.

Example 2: Closely-spaced levels

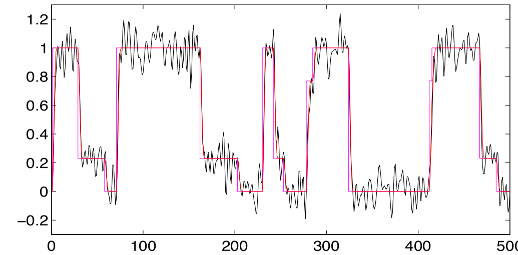

F&T analysis: 5kHz cutoff

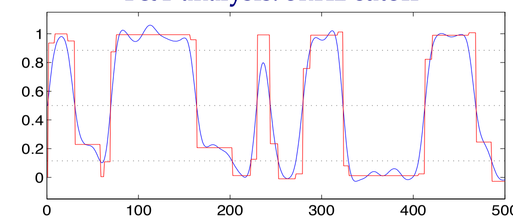

Bayesian analysis: Family of samples

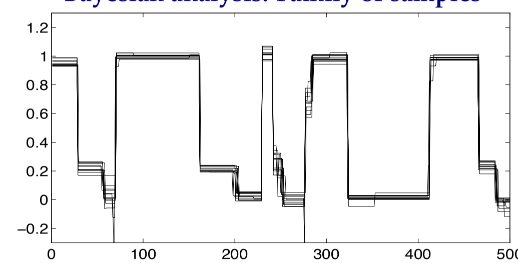

Bayesian analysis: Grand average

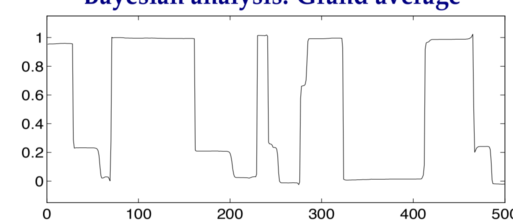

Note again that the levels in the F&T analysis are estimated from the *original data*, and not from the filtered series. Setting the cutoff frequency at 5kHz yields quite reasonable results. However, there are a number of false events. The nature of the Bayesian method of analysis are demonstrated further in this example. There are two short-lived events at around points 275 and 415. The question is: how much evidence does the data provide for these? The family of samples suggest that the evidence for the first is strong, but that for the second is less so. A weakness of F&T is that it gives no indication of this. The details of the change at point 250 are less clear than that at 200.

One further feature should be noted from the family of plots: the very short events such as those at point 275 and at point 450 in example 1. The effect of short spikes on the data as recorded is quite small, and the statistical model compares the recorded version of the idealized trace with the data. Hence the main penalty against the presence of such events in the samples is the event-rate distribution. This is realistic in the sense that there might be events such as these in the data. If they were to be deemed unreasonable, then the prior distribution could be modified to eliminate them.

All-points histograms

Example 3: All-points histogram, 5000 points

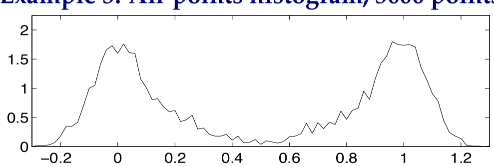

After filtering: 5kHz cutoff

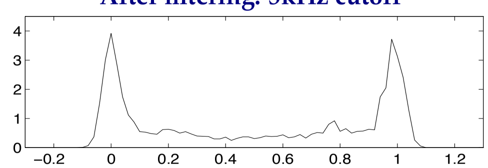

Bayesian analysis: 100 000 samples

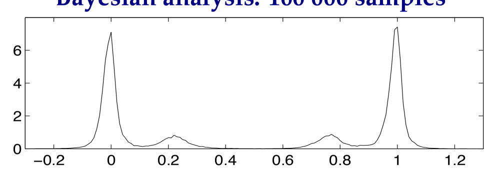

The all-points current-level histograms show the performance of the Bayesian method very clearly. A longer version (5000 points) of example 2 was taken and 100 000 samples drawn from the posterior distribution. Whilst filtering narrows the distributions, the time taken to sweep between levels obscures the results. The sampler produces neat and distinct component distributions.

Closely-spaced Events

The detection of temporally close events is difficult. This is especially true for F&T when the events span multiple thresholds. To test this we generated a slightly artificial example (4) with a large number of these, and once again F&T introduces many false events, even though the optimal thresholds have been supplied. Whilst entertaining the possibility of short spikes, the sampler shows that they have low probability in the posterior distribution. Even more difficult is the estimation of levels, particularly for the first two spikes and at point 470. The Bayesian model assigns a distribution of heights and widths to these: it is, so-to-speak, convinced of their existence and the locations of their centres. An average, such as that provided by F&T does not tell the whole story.

Example 4: Closely-spaced events

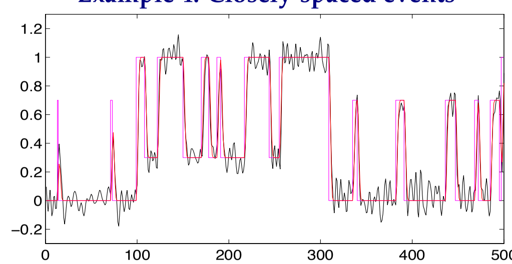

F&T analysis: 16.7kHz cutoff

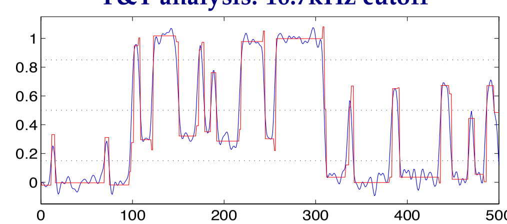

Bayesian analysis: Family of samples

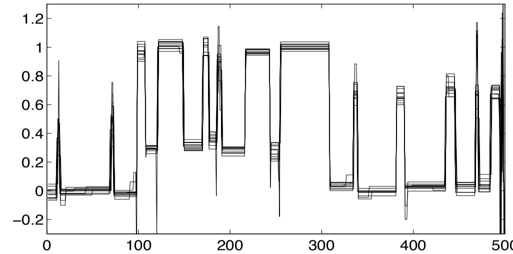

Bayesian analysis: Grand average

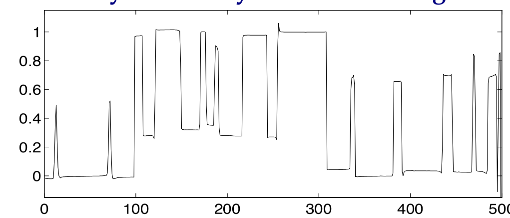

Discussion

This research has shown Bayesian modelling to be a powerful means of analysing patch-clamp records. It can deal with closely-spaced current levels and flicker-type events more effectively than filtering and thresholding, and it does so without imposing a particularly structured model such as HMMs. It can compensate for the effect of filtering. There are issues to be explored, especially concerning the display and interpretation of the results.

References

[1] D. Colquhoun and F. J. Sigworth. Fitting and statistical analysis of single channel records. In B. Sackmann and E. Neher, editors, *Single-Channel Recording*. Plenum, New York, second edition, 1995.

[2] S. H. Chung, V. Krishnamurthy, and J. B. Moore. Adaptive processing techniques based on Hidden Markov Models for characterizing very small channel currents buried in noise and deterministic interferences. *Phil. Trans R. Soc Lond. B*, 334:357–384, 1991.

[3] J. J. K. Ó Ruanaidh and W. J. Fitzgerald. *Numerical Bayesian Methods Applied to Signal Processing*. Statistics and Computing. Springer Verlag, 1996.