

NEUROFIT: software for fitting Hodgkin–Huxley models to voltage-clamp data

Allan R. Willms

Department of Mathematics and Statistics, Biomathematics Research Centre, University of Canterbury, Private Bag 4800, Christchurch, New Zealand

Received 21 March 2002; received in revised form 24 July 2002; accepted 25 July 2002

Abstract

I introduce publicly available software for accurate fitting of Hodgkin–Huxley models to voltage-clamp data. I describe the model and non-linear fitting procedure employed by the software and compare its results with the usual method of fitting such models using potassium A-current data from a pyloric dilator cell of the lobster *Panulirus interruptus* and sodium current data from an electrocyte cell of the electric fish *Sternopygus macrurus*. The set of parameter values for the model determined by this software yield current traces that are substantially closer to the observed data than those determined from the usual fitting method. This improvement is due to the fact that the software fits all of the parameters simultaneously utilizing all of the data rather than fitting steady-state and time constant parameters disjointly using peak currents and portions of the rising and falling phases. I analyze the convergence properties of the software's fitting algorithm using simulated data showing that accurate parameter values are obtained for most of the parameters using any reasonable initial values. The software also incorporates a linear pre-estimation procedure to help in determining reasonable initial values for the full non-linear algorithm. I illustrate and discuss some of the inadequacies of voltage-clamp data.

© 2002 Elsevier Science B.V. All rights reserved.

Keywords: Hodgkin–Huxley model fitting; Computer software; Voltage-clamp data; Inactivating current

1. Introduction

The voltage-clamp method has been used extensively in the past 50 years to obtain data from which the behavioral properties of ion channels can be determined. These properties are usually described by the values of parameters for a particular model. The Hodgkin–Huxley (HH) model for voltage-dependent channels, which specifies a set of independent gates governed by first order kinetics, has been widely and successfully used (Hodgkin and Huxley, 1952; Rinzel and Lee, 1987; Buchholtz et al., 1992; Guckenheimer et al., 1993; De Schutter and Bower, 1994; Luo and Rudy, 1994), although other models, such as multi-state models (Goldman, 1975; Destexhe et al., 1994; Marom and Levitan, 1994) have also been employed.

As a technological development, the voltage-clamp allowed Hodgkin and Huxley to formulate and verify the idea that the excitability of the neuronal membrane was due to the independent flow of specific ions across the membrane through voltage-dependent channels (Hodgkin and Huxley, 1952). Mathematically speaking, the constant voltage environment allowed the equations of their model to be solved analytically, and permitted a relatively easy computational method of determining estimates of the parameter values from the recorded data. This method, which is still regularly employed today, uses peak currents from a series of voltage-clamp experiments to estimate the steady-state parameters for the gates, and portions of the rising and falling phases of the current to estimate the time constants of activation and inactivation. However, it must be emphasized that this disjoint method is primarily one of computational ease and not the most accurate method available for fitting parameters to data of this type. With the computing power available today, more sophisticated methods can be employed to obtain estimates of the

E-mail address: a.willms@math.canterbury.ac.nz (A.R. Willms).

parameters which are much more consistent with the data.

We have developed a method for fitting HH model parameters to voltage-clamp data which uses the full current trace rather than just peak currents or just rising or falling phases (Willms et al., 1999). This method finds least-squares estimates for both steady-state and dynamic parameters simultaneously by fitting the mathematical analytic solution of the differential equations to the entire current (or conductance) trace. The advantages and success of this method over the usual disjoint method are described fully in that paper. We also made the software (FULLTRCE) implementing this method available to others. However, FULLTRCE was not very user-friendly and lacked features such as the ability to graph the data and resulting fits, or read binary files produced by proprietary voltage-clamp software.

Here I wish to announce an improved user-friendly version of that software which is available for public use from <http://www.math.canterbury.ac.nz/~matharw/neurofit/>. This new software (called NEUROFIT) will read data files in either text format or Axon Instruments binary format (abf files). The data are plotted on the screen and the user is allowed to specify a region of the data which she/he wishes to fit. After giving the number of activation gates and other specifications for the model and the fitting algorithm, the software fits the model to the data, plots the resulting fitted curves on top of the displayed data, and displays the fitted values and standard errors for all of the parameters.

Two examples in Section 3.1 illustrate the superiority of the algorithm employed by NEUROFIT over the usual disjoint method. (A fuller comparison of this method with the disjoint method can be found in Willms et al. (1999).) In this paper, I extend the analysis of the algorithm employed by NEUROFIT by determining its convergence properties using simulated data. The results in Section 3.2 show that accurate parameter values are obtained for most of the parameters using any reasonable initial values. Some of the inadequacies of voltage-clamp data are discussed in Section 4.

2. Methods

This section gives a precise formulation of the HH model employed by NEUROFIT, and a short description of the fitting procedure.

2.1. Hodgkin–Huxley model

The electrical current through the ion channels is a product of the conductance, g , and the difference between the membrane potential, V , and the reversal potential, E_{rev} , for the current,

$$I = g(V - E_{\text{rev}}). \quad (1)$$

The conductance is given by,

$$g = g_{\text{max}} m^p h, \quad (2)$$

where g_{max} is the maximal conductance, p is the number of independent activation gates, and m and h are the fraction of open activation and inactivation gates, respectively.

As a generalization of the usual HH scheme, the channels are classified into one of several groups depending on their kinetic inactivation properties. There are n_h different groups of inactivating channels. All of the channels in all of these groups have one inactivation gate and identical *steady-state* inactivation properties, but each group has different *kinetic* inactivation properties. There may also be one additional group with non-inactivating channels ($n_{\text{nonh}} = 0$ or 1). Thus the total number of groups is $n = n_h + n_{\text{nonh}}$, $n \geq 1$. The fraction of channels in group i is given by f_i ; these fractions add to one, $\sum_{i=1}^n f_i = 1$. The variable h_i represents the fraction of open channels in group i , and the overall fraction of channels with open (or non-existent) inactivation gates, h is given by,

$$h = \begin{cases} \sum_{i=1}^{n_h} f_i h_i, & \text{if } n_{\text{nonh}} = 0, \\ \sum_{i=1}^{n_h} f_i h_i + f_n, & \text{if } n_{\text{nonh}} = 1. \end{cases} \quad (3)$$

This generalization of the standard HH scheme was used to allow reasonable fitting of currents such as the crustacean transient potassium current, I_A , which shows at least a double exponential time course of inactivation (Harris-Warrick et al., 1995). The behavior of I_A could also be captured using a more complex generalization, such as having the f_i be voltage-dependent (Buchholz et al., 1992), allowing each group to have different steady-state inactivation properties, or by using a multi-state model, but I have found the above generalization to be sufficient for fitting I_A data.

The dynamics of the gates are governed by the first order differential equations,

$$\frac{dm}{dt} = \frac{m_{\infty}(V) - m}{\tau_m(V)}, \quad (4)$$

$$\frac{dh_i}{dt} = \frac{h_{\infty}(V) - h_i}{\tau_{h_i}(V)}, \quad i = 1, \dots, n_h, \quad (5)$$

where the steady-state curves are given by the Boltzmann functions,

$$m_{\infty}(V) = \frac{1}{1 + \exp((V - V_{2m})/s_m)}, \quad (6)$$

$$h_{\infty}(V) = \frac{1}{1 + \exp((V - V_{2h})/s_h)}, \quad (7)$$

with half potentials, V_{2x} , and slope factors, s_x , $x = m$ or h . The time constants of activation, τ_m , and inactivation

τ_{h_i} , are voltage-dependent and are specified independently for each step potential.

2.2. Fitting procedure

NEUROFIT fits voltage-clamp data. Each trace has a specified pre-step potential, V_p , and step potential, V_s . The data for each trace consist of a set of times and electrical currents, (t_j, I_j) , $j = 1, 2, \dots, N$. For each trace, Eqs. (4) and (5) are solved analytically with $V = V_s$, using initial conditions determined by Eqs. (6) and (7) evaluated at $V = V_p$. Thus it is implicitly assumed that the pre-step was held sufficiently long for the activation and inactivation to have achieved their steady-state values. For a given set, α , of parameter values, the current, $I(t, V_p, V_s, \alpha)$, is then calculated from Eqs. (1)–(3) and compared to the observed current. The software attempts to choose the parameter values, α , to minimize the quantity,

$$R = \sum_{j=1}^N \left(\frac{I(t_j, V_p, V_s, \alpha) - I_j}{\sigma_j} \right)^2,$$

where σ_j is the noise level for data point j . These noise levels are the same for each point in a trace and are estimated at the beginning of the optimization procedure by fitting the end of each trace (where the data are essentially constant) to a low-order polynomial. The optimization procedure employed by NEUROFIT is the well-known Levenberg–Marquardt algorithm (Moré, 1977) and has been implemented as a C translation of the Fortran code `lmdex.f`, the latter being freely available from Bell Lab’s Netlib facility.¹ The algorithm makes use of the analytic Jacobian of $I(t_j, V_p, V_s, \alpha)$ when determining how to adjust the parameter values in order to decrease R . The current is then re-calculated using the adjusted parameter values and the process is iterated until R converges to a minimum or another termination condition is encountered.

As with any non-linear least-squares algorithm, a set of initial values for the parameters must be supplied by the user. The program allows parameters to be held constant at their specified value or varied by the optimization algorithm. This feature allows the user to, for example, fix the reversal potential if it is known to be a specific value. If the data are particularly difficult to fit, this feature also aids in obtaining good results by allowing the user to fix a number of parameters at reasonable values and ask the algorithm to obtain optimal values for the remainder. The obtained optimal values can then be easily transferred to the initial values and the fitting algorithm re-run with more of the parameters allowed to vary.

Since the non-linear optimization algorithm is based on local quadratic approximations of the function R , there is always the possibility that the algorithm will converge to a local minimum for the parameters, α , rather than the global minimum. The existence of local minima depends on the complexity of the model employed and is a strong argument for keeping the model as simple as possible. To ensure that any obtained solution is not just a local minimum, it is necessary to choose initial values that are sufficiently close to the global minimum. I discuss this issue in Section 3.2 below.

As an aid to the user in determining reasonable initial values for the non-linear fitting algorithm, NEUROFIT also has a ‘Quick Fit’ option which generates fairly good parameter estimates and only requires an initial value for the reversal potential, E_{rev} . The ‘Quick Fit’ algorithm basically de-constructs the problem into a series of linear problems and is based on the method described by Tóth and Crunelli (1995). The values obtained by ‘Quick Fit’ may then be used as initial values for the full non-linear fitting algorithm, although in some instances these values may not be very accurate and may need to be adjusted by the user before doing the full non-linear fit. This is particularly the case for the time constants of inactivation if n_h is greater than one. Nonetheless, using this linear pre-estimation procedure should increase the likelihood that the non-linear algorithm will converge to the global minimum.

3. Results

3.1. Examples

Here I compare the results obtained by NEUROFIT with those obtained by the usual disjoint method for voltage-clamp data from two voltage-dependent inactivating currents: the potassium current I_A from *Panulirus* and the sodium current I_{Na} from *Sternopygus*. Since the overall goal of fitting the data is to minimize the variance between the observed data and the model predictions, I used an F -test to compare the calculated variances resulting from two different fits. The F -test statistic is the ratio of the variances of the two fits, which, assuming these variances have χ^2 distributions, has Snedecor’s F -distribution (Beyer, 1987). This assumption is valid if the fits are reasonably good, but becomes less so if the fits are bad and the residues are not Gaussian. Nonetheless it is a reasonable comparative measure of the goodness of two fits. In what follows, ‘better’ means a statistically significant difference in the variance using an F -test at the 5% level.

¹ <http://www.netlib.org>.

3.1.1. Results for I_A

The I_A data were from the pyloric dilator (PD) cell of the stomatogastric ganglion in *Panulirus interruptus*. The cell was treated² to block as much of the other currents and synaptic transmission as possible. The voltage-clamp protocols, and the recorded currents are shown in Fig. 1. In order to remove any of the remaining potassium delayed rectifier current, I_K , from the records, additional protocols identical to those in the inset of the top panel of Fig. 1 but lacking a pre-step were also performed. Such protocols do not remove the inactivation of I_A and therefore represent only the remaining I_K current. The currents from these protocols (which were very small) were subtracted from the currents obtained from the other protocols. The subtracted resultant currents are displayed in Fig. 1.

The parameter values obtained from these data using the usual disjoint method and NEUROFIT are displayed in Table 1. Since the disjoint method does not determine an estimate for E_{rev} , its value was fixed at -86 mV (a typical value for pyloric cells in the STG, Baro et al., 1997) when using NEUROFIT to ease comparison of both methods. Some remarks about the model specifications, that is, values of p , n_h , and n_{nonh} , given in Table 1 are in order. I tested values of p from 1 to 10.




Fig. 1. Voltage-clamp experiments on a PD cell of *P. interruptus*. Recorded I_A from varying step potential experiments (top panel), and from varying pre-step potential experiments (bottom panel). Insets: voltage protocols.

² 10^{-7} M tetrodotoxin (TTX), 5×10^{-6} M picrotoxin, 2×10^{-2} M

For the disjoint method, although increasing p by one always reduced the variance in the fit to the steady-state activation curve and the rising phases of the currents, this improvement was not statistically significant. This is not surprising since fitting steady-state curves to peak currents is not a good way to estimate p (Willms et al., 1999). For NEUROFIT, increasing p by one gave a better overall fit up to $p = 4$. I decided to use $p = 3$ since it was a reasonable value for both methods and is historically a commonly used value for I_A . For the disjoint method, setting n_h to two was better than one, but three was not better than two. For NEUROFIT, incrementing n_h by one made the fit better, up to $n_h = 3$. However, the values for $n_h = 2$ are displayed in the table, for ease of comparison. The data do not have any non-inactivating component, so n_{nonh} was set to zero. (Setting n_h to two and n_{nonh} to one in NEUROFIT resulted in a fitted value for f_3 , the fraction of non-inactivating channels, of zero.) Table 1 thus compares essentially the best fit from the disjoint method with a good, but not the best, fit from NEUROFIT.

Using the two sets of parameter values in Table 1, the electrical currents of the cell under the voltage-clamp protocols can be re-constructed. Some of these reconstructions are shown in Fig. 2 (only a few are shown to prevent the figure from being too cluttered). It is clear from these figures that NEUROFIT out-performs the disjoint method and does a better job at fitting the data. The difference in the variances of the two fits (excluding the first 5 ms of the traces which are contaminated by noise and were not used in the fitting procedures) was so large, it was easily statistically significant for the F -test at even the 0.1% level.

3.1.2. Results for I_{Na}

The I_{Na} data were from an electrocyte cell of *Sternopygus macrurus*. The cell was treated with TEA to block potassium currents and was bathed in saline in which NaCl was replaced with sodium methyl sulfate to reduce the chloride level of the saline, thereby reducing the passive leakage. The remaining leak current was estimated from the traces after the sodium current had inactivated and this was digitally subtracted from the data. The voltage-clamp protocols, and the recorded currents (after leak subtraction) are shown in Fig. 3. The capacitive current dominates the first half millisecond of the trace and this portion was excluded in the fitting procedures. Voltage steps were in 5 mV increments from -40 to $+90$ mV. The reversal potential for this cell is in the range 50–60 mV (McAnelly and Zakon, 1996) so voltage steps above this level should elicit a positive current. However, in this experiment, the cell was not well clamped at high depolarizations, that is, the actual membrane voltage did not achieve the command voltage and/or was not sufficiently constant near the start of the trace. The data in Fig. 3 show that only the upper three

Table 1
Estimated parameter values for A-current of the PD cell

	Disjoint	NEUROFIT				
<i>Model specifications</i>						
p	3	3				
n_h	2	2				
n_{nonh}	0	0				
<i>Other parameters^a</i>						
E_{rev}	−86	−86				
g_{max}	0.46	0.48				
f_1	0.57 ^b	0.49				
V_{2m}	−50.9	−45.3				
s_m	−20.9	−11.6				
V_{2h}	−69.6	−69.4				
s_h	5.2	4.8				
Voltage (mV)	Disjoint	NEUROFIT				
	τ_m	τ_{h_1}	τ_{h_2}	τ_m	τ_{h_1}	τ_{h_2}
<i>Time constant parameters (ms)</i>						
40	3.0	113	703	3.0	113	728
30	3.3	118	685	3.4	101	667
20	3.7	128	658	3.6	92	575
10	4.1	133	649	4.6	72	506
0	4.5	115	510	5.3	61	416
−10	5.3	119	504	6.5	58	372
−20	6.3	112	485	7.9	66	332
−30	8.7	107	407	10.0	95	294
−40	16.0	142	348	12.3	220 ^c	220 ^c

^a Units for these parameters are mV except for g_{max} which has units μS , and f_1 which is dimensionless.

^b Different values of f_1 and f_2 are estimated for each voltage level in the disjoint method. The value for f_1 reported here is an average of the values obtained from all the traces from the varying step protocol except the -40 mV step which produced almost no signal. The range of values for f_1 was 0.46–0.71.

^c The current trace for the step to -40 mV did not have a sufficiently large signal for the NEUROFIT algorithm to discern two distinct inactivation time constants.

voltage traces, corresponding to steps to 80 mV and greater, gave a positive current. For this reason, only traces with voltage steps below 60 mV were used in the fitting procedures.

The parameter values obtained from these data using the usual disjoint method and NEUROFIT are displayed in Table 2. The reversal potential for the disjoint method was set to 59.6 mV being the optimal value obtained by NEUROFIT, and the number of activation gates was set to $p = 5$ since NEUROFIT found this to be best. A single inactivation time constant was used.

The most striking difference between the two sets of parameter values is the estimate for g_{max} ; the disjoint method gave 53 μS but NEUROFIT gave 211 μS . This large difference is due to the fact that the disjoint method does not take into account the simultaneous nature of the activation and inactivation processes, a deficiency which is relatively minor when the time constants for the two processes are widely separated as in I_A above, but which is drastically damaging when the two processes occur at nearly the same rates as for I_{Na} .

Fig. 4 shows a few of the electrical currents reconstructed from the parameter values of Table 2. The substantial under-estimation of g_{max} by the disjoint method is immediately apparent in this figure where the amplitude of the disjoint re-constructions are about a factor of three too small. The time constants of activation found by NEUROFIT were mostly larger than those obtained by the disjoint method, while the time constants of inactivation were mostly smaller. Another notable difference in the two fits is the values of V_{2m} and s_m where the disjoint method has a very shallow steady-state curve shifted to the left compared to NEUROFIT. This is likely due to the fact that the steady-state curve for this data is non-symmetrical and is not well fit by a Boltzmann function of the form Eq. (6); a more general model for this data may be appropriate. The overall fit obtained by NEUROFIT for this current is not quite as good as that obtained for I_A but the improvement over the disjoint method is even greater.

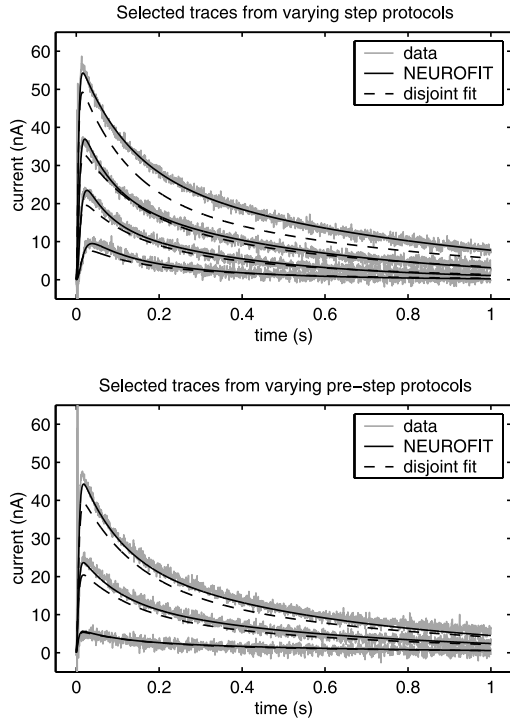


Fig. 2. Re-constructed current traces for the data from Fig. 1 (gray lines) using parameter values obtained by the disjoint method (dashed lines) and by NEUROFIT (solid lines). Top panel: varying step protocols, steps to -30 , -10 , 10 , and 40 mV. Bottom panel: varying pre-step protocols, pre-steps to -60 , -70 , and -100 mV.

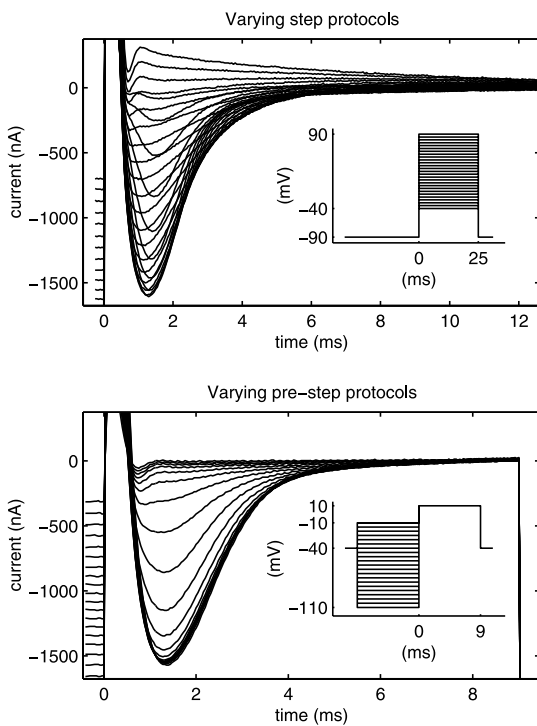


Fig. 3. Voltage-clamp experiments on an electrocyte cell of *Sternopygus*. Recorded I_{Na} from varying step potential experiments (top panel), and from varying pre-step potential experiments (bottom panel). Insets: voltage protocols.

Table 2
Estimated parameter values for I_{Na} of the electrocyte cell

	Disjoint	NEUROFIT		
<i>Model specifications</i>				
p	5	5		
n_h	1	1		
n_{nonh}	0	0		
<i>Other parameters^a</i>				
E_{rev}	59.6	59.6		
g_{max}	53.3	211.0		
V_{2m}	-48.1	-33.3		
s_m	-25.7	-7.15		
V_{2h}	-48.4	-49.3		
s_h	7.21	5.56		
Voltage (mV)	Disjoint		NEUROFIT	
	τ_m	τ_h	τ_m	τ_h
<i>Time constant parameters (ms)</i>				
55	0.22	2.56	0.25	2.27
50	0.22	2.36	0.40	1.61
45	0.24	2.16	0.46	1.39
40	0.25	2.03	0.49	1.27
35	0.26	1.85	0.51	1.19
30	0.27	1.68	0.52	1.13
25	0.28	1.54	0.53	1.08
20	0.29	1.39	0.54	1.03
15	0.31	1.28	0.54	0.99
10	0.31	1.15	0.59	1.00
5	0.33	1.05	0.56	0.93
0	0.35	0.99	0.58	0.90
-5	0.37	0.95	0.60	0.88
-10	0.40	0.94	0.61	0.87
-15	0.42	0.93	0.62	0.87
-20	0.45	0.96	0.61	0.90
-25	0.46	1.00	0.56	1.00
-30	0.45	1.08	0.46	1.24
-35	0.42	1.23	0.29	1.91
-40	0.38	1.46	0.29	1.69

^a Units for these parameters are mV except for g_{max} which has units μS .

3.2. Convergence properties

Here I analyze the size of the convergence region, that is, the maximum distances the initial values for the parameters can be from the 'true' values and still have the non-linear algorithm converge to a reasonably accurate solution.

The convergence of the algorithm to a solution and the accuracy of that solution will depend on several factors. The degree to which the model's form accurately reflects the behavior of the ion channels and the amount of noise in the data will of course have significant effects. Further, since the model is non-linear and the algorithm essentially approximates this non-linearity, both the true values and the initial values of *all* of the parameters will affect the convergence of *each* parameter. Thus, even in the absence of noise, the convergence regions will in general be different for data

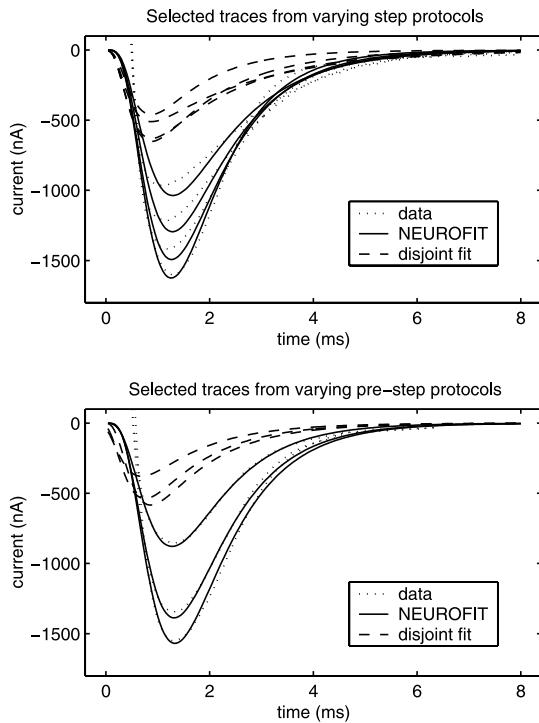


Fig. 4. Re-constructed current traces for the data from Fig. 3 (dotted lines) using parameter values obtained by the disjoint method (dashed lines) and by NEUROFIT (solid lines). Top panel: varying step protocols, steps to 0, 20, 30, and 40 mV. Bottom panel: varying pre-step protocols, pre-steps to -50 , -60 , and -90 mV.

representing models with different ‘true’ values of the parameters. It is therefore impossible to make clear statements such as ‘Provided the initial value for g_{\max} is within 40% of its true value the algorithm will converge to a solution where g_{\max} is within 5% of its true value’. However, some idea of the size of the convergence region can be obtained.

I generated simulated data for two types of inactivating channels: the crustacean potassium A-current, I_A ; and a classical sodium current, I_{Na} . The true values of the parameters used to generate these data are shown in Table 3. For each current, two sets of voltage-clamp experiments were simulated, the first with varying step potentials and the second with varying pre-step potentials. Five simulated data sets for each current were generated by adding a different realization of Gaussian noise to the model currents (standard deviation: 2.0 nA for I_A , 1.0 nA for I_{Na}). A portion of the data for one of these realizations for each current type, and the generating voltage-clamp protocols, are shown in Figs. 5 and 6. The first 350 ms of the voltage step were fit for I_A , while for I_{Na} the first 5 ms were fit. The noise in each trace was calculated by fitting a line to the regions 350–450 and 6–10 ms for I_A and I_{Na} , respectively.

The following tests were performed for both I_A and I_{Na} . Starting at the true parameter values, a direction in parameter space was chosen at random and the initial

Table 3
True parameter values for simulated currents

	I_A	I_{Na}
<i>Model specifications</i>		
p	3	3
n_h	2	1
n_{nonh}	0	0
<i>Other parameters^a</i>		
E_{rev}	-86	50
g_{\max}	3.9	5.3
f_1	0.36	
V_{2m}	-42	-8
s_m	-15	-10
V_{2h}	-67	-46
s_h	6	4
<i>Voltage (mV)</i>		
	I_A	I_{Na}
	τ_m	τ_m τ_h
<i>Time constant parameters (ms)</i>		
40	–	– 0.35 1.0
30	–	– 0.33 1.0
20	2.0	25 106 0.30 1.0
10	2.2	27 117 0.26 1.0
0	2.6	30 128 0.22 1.0
-10	3.0	34 139 0.18 1.0
-20	3.6	39 152 0.14 1.0
-30	4.4	45 171 0.10 1.2
-40	5.4	52 192 0.06 1.5
-50	6.7	60 213 – –

^a Units for these parameters are mV except for g_{\max} which has units μS , and f_1 which is dimensionless.

values for the parameters were systematically moved outward along this direction to determine the method’s convergence properties in that direction for each of the five data sets. The total number of directions chosen was 16 times the number of parameters for the current (496 for I_A and 384 for I_{Na}).

A test was labeled successful for a particular parameter if the final value obtained by the fitting algorithm for that parameter was within 5% of its true value, that is, if

$$\left| \frac{\text{true value} - \text{obtained value}}{\text{true value}} \right| < 0.05.$$

I pooled all of the data from the five data sets and then, for each parameter, attempted to find the largest ellipsoidal region in parameter space (centered at the true parameter values) such that all of the non-successful tests lay outside this region. Thus the region approximates a lower bound on the algorithm’s domain of convergence for that parameter. Convergence will almost always occur for initial value sets inside this region, but, due to the highly non-linear nature of the problem, will likely also occur for a substantial number of initial value sets outside this region. The ellipsoidal




Fig. 5. Simulated I_A voltage-clamp data. Left panel: varying step potential experiments; right panel: varying pre-step potential experiments; insets: voltage-clamp protocols.

region in parameter space is a piece-wise ellipsoid with semi-axes aligned along the coordinate axes, and is defined by $2n$ semi-axis values, where n is the number of parameters (the dimension of the space). An example of such an ellipsoidal region in two dimensions is shown in Fig. 7. In this figure, P and Q are parameters and the ellipsoidal region is defined by the semi-axis values P_+ , P_- , Q_+ , and Q_- . Distance is measured in units of relative distance from the true parameter values. Fig. 7A shows the region itself, while Fig. 7B shows a plot of the semi-axis values for each parameter. It is this second representation that I will use for the high-dimensional parameter spaces.

In the determination of these domains of convergence, I restricted the size of the semi-axis values to a maximum of 2.0, which represents a 200% change in the parameter value. In addition, I restricted some semi-axis values to a maximum of 1.0 if values greater than this were invalid (for example, the semi-axis values measuring decrease of the time constants; time constants must be positive and hence cannot decrease more than 100% from their true values). In practice, the initial values (either obtained from 'Quick Fit' or entered manually) should be well within this distance, for most of the parameters.

Although in general there is one domain of convergence (ellipsoidal region) for *each* parameter, my determination of these domains, which was not completely thorough, resulted in only a small number of




Fig. 7. (A) Example ellipsoidal region in two dimensions with semi-axis values P_+ , P_- , Q_+ and Q_- . (B) Semi-axis representation of the region.

distinct regions. That is, many of the parameters have approximately the same domain of convergence. In order to obtain more differentiation, a finer gradation of distances along each direction would need to be used, perhaps a larger number of directions would need to be chosen, and more effort at finding the largest ellipsoidal region, which is itself a difficult optimization problem, would need to be exerted. Nonetheless, the results are quite informative regarding the sensitivity of the data to the various parameters, and the convergence properties of the algorithm.

3.2.1. Convergence for simulated I_A

For I_A there were three distinct domains of convergence as displayed in Fig. 8. The largest of these had most semi-axis values around 0.5 indicating that as long as the relative distances of the initial parameter values

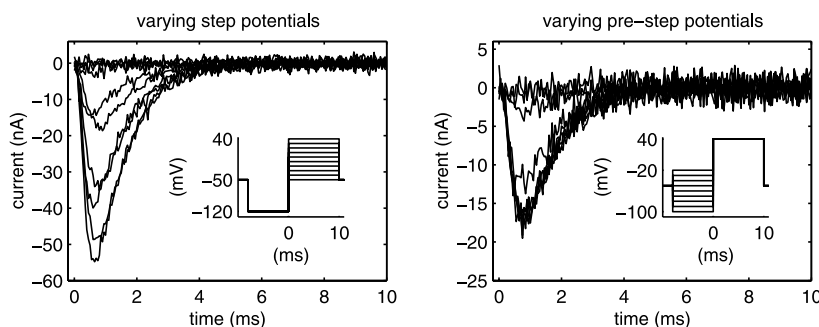


Fig. 6. Simulated I_{Na} voltage-clamp data. Left panel: varying step potential experiments; right panel: varying pre-step potential experiments; insets: voltage-clamp protocols.




Fig. 8. Ellipsoidal domains of convergence for I_A . The domains of convergence of the parameters for I_A fell into one of three (top, middle, bottom) domains shown here; shaded parameter names indicate the parameters for which the plotted domain applies. In each plot, the parameters are listed horizontally and for each, a relative distance above and below the true value of the parameter is plotted. (The plots are scaled in order to see the details near zero, hence some of the values are off the plots; these values were typically large or at the maximum distance.) If the initial values of the parameters are within these ranges, the algorithm converges to within 5% (dotted lines) of the true value for the shaded parameter names.

were within about 50% of their true values, the algorithm would converge successfully for the following parameters: f_1 , V_{2m} , V_{2h} , s_h , $\tau_m(20)$, \dots , $\tau_m(-30)$, $\tau_{h_1}(20)$, \dots , $\tau_{h_1}(-20)$, $\tau_{h_2}(20)$, \dots , $\tau_{h_2}(-30)$. The second largest region had semi-axis values around 0.2, indicating that initial guesses needed to be within about 20% of their true values for successful convergence. The parameters with this domain of convergence were: E_{rev} , $\tau_m(-40)$, $\tau_m(-50)$, $\tau_{h_1}(-30)$, \dots , $\tau_{h_1}(-50)$, $\tau_{h_2}(-50)$.

It is clear that the algorithm successfully obtained good estimates for most of the I_A parameters even if the initial guesses were well away from their true values. The time constants at voltages where the current was barely activated (for example, the traces in Fig. 5 corresponding to voltage steps to -40 and particularly -50 mV)

were difficult to fit, however, this was not due to the algorithm, but rather the lack of sufficient signal in the data themselves.

The first inactivation time constant at -30 mV, $\tau_{h_1}(-30)$, also had a small domain of convergence. This is a manifestation of the fact that the signal at this voltage level was also not sufficiently large to allow a clear separation of the two time constants of inactivation. Indeed, fitting sums of exponentials with varying weights is in general a highly ill-posed problem (Acton, 1990). For this model, the problem is not as ill-posed because the weights (values of f_i) are required to be the same for each current trace. Nonetheless, as the number of channel types, n_h , increases, the problem becomes more ill-posed and one should expect large standard

errors for the parameter estimates, and small domains of convergence. (It should be noted that the disjoint method obtains different weightings, f_i , for each voltage level and does not even attempt to obtain a consistent weighting valid for all voltage levels.)

The domain of convergence for g_{\max} was of moderate size, and that of E_{rev} was quite small. This is due to the fact that there is some degree of degeneracy or trade-off between these two parameters as seen in Eqs. (1) and (2). If E_{rev} is relatively distant from the voltage step levels, small changes in E_{rev} , will affect all traces to almost the same degree, and therefore be equivalent to small changes in g_{\max} . This is precisely the situation for the I_A data, where E_{rev} is -86 mV but the current traces with substantial signal are at voltage levels above -20 mV. It is therefore sometimes necessary with such data to obtain estimates of E_{rev} , from other means and hold that parameter constant in the fitting procedure to allow an accurate estimate of g_{\max} .

3.2.2. Convergence for simulated I_{Na}

For I_{Na} there were just two distinct domains of convergence as displayed in Fig. 9. The largest of these regions corresponded to the parameters E_{rev} , s_m , V_{2h} , s_h , $\tau_m(40), \dots, \tau_m(0)$, and $\tau_h(40), \dots, \tau_h(-10)$. For this region, most of the parameters could

be initially about 200% larger than their true value and the algorithm would still converge. The only relatively small semi-axis values were those corresponding to decreases in g_{\max} and V_{2m} , and those corresponding to some of the time constants. The smaller region had most of its semi-axis values within the 5% region indicating a need for very accurate initial values. The parameters with this domain of convergence were g_{\max} , V_{2m} , $\tau_m(-10), \dots, \tau_m(-40)$, $\tau_h(-20), \dots, \tau_h(-40)$.

For I_{Na} , the results were comparable with I_A particularly in that the time constants at voltages where the current was barely activated and the signal to noise ratio was very small, were the most difficult parameters to fit. But even so, as seen in Fig. 9, the time constants would generally tolerate a large initial over-estimate; it was the under-estimates, particularly with the time constants of activation, which caused failure to converge.

In contrast to I_A , for I_{Na} , the parameter E_{rev} , had a large domain of convergence, that is, was easily estimated. This is due mostly to the fact that the effect of E_{rev} on these data is clearly distinct from g_{\max} , in that the largest signal currents were at voltages relatively near the reversal potential.

The fact that g_{\max} and V_{2m} had a small domain of convergence for I_{Na} was due to a lack of sensitivity between these parameters in these data. Only six of the




Fig. 9. Ellipsoidal domains of convergence for I_{Na} . The domains of convergence of the parameters for I_{Na} fell into one of two (top, bottom) domains shown here; shaded parameter names indicate the parameters for which the plotted domain applies. See the legend to Fig. 8 for details.

nine voltage steps generated a substantial signal since the current was nearly completely deactivated at about -20 mV. The complicated interaction of a small value for s_m , and the various values for the time constants of activation, coupled with essentially just six relevant data traces resulted in a small domain of convergence for these parameters.

4. Discussion

The two examples of real data in Section 3.1 illustrate that NEUROFIT is substantially better than the disjoint method at fitting a HH model to voltage-clamp data. This success is due to the fact that NEUROFIT uses the full current trace, that is, all of the data, rather than just the peak currents and parts of the rising and falling phases. Further, by fitting all of the data simultaneously, it obtains a consistent set of parameters, unlike the disjoint method which obtains different weights, f_i , for the multi-exponential inactivation at each voltage level. The end result is a set of parameter values which better reflects the information contained in the data.

The analysis in Section 3.2 shows that NEUROFIT will converge to accurate estimates of most of the parameters starting from any reasonable initial values. It should again be emphasized that the ellipsoidal regions determined in that section approximate a *lower bound* on the domain of convergence, and initial guesses outside this region will also, in many cases, result in successful convergence. The linear ‘Quick Fit’ pre-estimation feature of the software provides reasonable initial values for the full non-linear fitting procedure.

It is still true that NEUROFIT is constrained by the data set itself, and difficulties associated with its quality or the lack of information in it will affect the fit. For example, the time constants at voltage levels where the current was barely activated and hence the signal to noise ratio was small, had small domains of convergence due to the fact that the data did not contain sufficiently discriminatory information for these parameters. The analysis in Section 3.2 also reveals some other inadequacies in standard voltage-clamp data.

- if the reversal potential, E_{rev} , is distant from the voltage step levels at which an appreciable current is recorded, there is difficulty in distinguishing changes in E_{rev} from those of the maximal conductance, g_{max} ;
- there must be a clear multi-exponential time course in a sufficient number of the traces in order for a consistent estimate of the time constants of inactivation, τ_{h_i} , and the fraction of channels of each inactivation type, f_i , to be obtained;
- if the voltage-clamp apparatus is not successful in depolarizing the cell precisely to the command voltage or is unable to do so sufficiently instantaneously,

due to, for example, the spatial extension of the cell, electrode resistance, or other physical or technological constraints, the recorded current will not correspond to the analytic voltage-constant solutions of the differential equations. The parameters most affected by this problem are E_{rev} , g_{max} , V_{2m} , s_m and the time constants at these depolarized levels;

- if the membrane time constant is the same order of magnitude as the time constants of activation, the rising phase of the current will be substantially concealed by the capacitive surge leaving little data from which to determine τ_m and p ;
- if there are a restricted number of voltage step levels at which recordings are made or if they are not sufficiently spread over the region of interest, the ability to determine some of the parameters (notably the half activation potential, V_{2m} , and the activation slope factor, s_m) will be limited.

No fitting algorithm will fully overcome these difficulties as they are associated with a lack of (or erroneous) information in the data set itself. Different experimental protocols which yield more informative data are necessary.

The fitting algorithm employed by NEUROFIT is not dependent on the HH model. Indeed, I intend in the future to generalize the software to allow fitting of other models such as multi-state models or models with non-symmetric steady-state curves.

NEUROFIT is easy to use and can obtain accurate estimates for HH model parameters. Its use should benefit the neurobiological community in fitting HH models in a more rigorous and consistent manner to standard voltage-clamp data.

Acknowledgements

The development of NEUROFIT was partially completed by Greg Ewing and Howie Kuo. The voltage-clamp data for I_A and I_{Na} were supplied by Ronald Harris-Warrick, and Lynne McAnelly, respectively. The comments of two anonymous reviewers were greatly appreciated. This research has been supported financially by the University of Canterbury, internal grant U6322, and the Royal Society of New Zealand, International Science and Technology Linkages Fund, NZ/USA Co-operative Science Programme, grants 99-CSP-49-WILL and 01-CSP-44-WILL.

References

Acton FS. Numerical Methods That Work. Washington DC: Mathematical Association of America, 1990.

- Baro DJ, Levini RM, Marshall TK, Willms AR, Cole Lanning C, Rodriguez HE, Harris-Warrick RM. Quantitative single-cell-reverse transcription-PCR demonstrates that A-current magnitude varies as a linear function of *shal* gene expression in identified stomatogastric neurons. *J Neurosci* 1997;17:6597–610.
- Beyer WH, editor. CRC Standard Mathematical Tables, 28th ed. Boca Raton, FL: CRC Press, 1987:536–78.
- Buchholtz F, Golowasch J, Epstein IR, Marder E. Mathematical model of an identified stomatogastric ganglion neuron. *J Neurophysiol* 1992;67:332–40.
- De Schutter E, Bower JM. An active membrane model of the cerebellar Purkinje cell. I. Simulation of current clamps in slice. *J Neurophysiol* 1994;71:375–400.
- Destexhe A, Mainen ZF, Sejnowski TJ. Synthesis of models for excitable membranes, synaptic transmission and neuromodulation using a common kinetic formalism. *J Comput Neurosci* 1994;1:195–230.
- Goldman L. Quantitative description of the sodium conductance of the giant axon *Myxicola* in terms of a generalized second-order variable. *Biophys J* 1975;15:119–36.
- Guckenheimer J, Gueron S, Harris-Warrick RM. Mapping the dynamics of a bursting neuron. *Phil Trans R Soc Lond B* 1993;341:345–59.
- Harris-Warrick RM, Coniglio LM, Barazangi N, Guckenheimer J, Gueron S. Dopamine modulation of transient potassium current evokes phase shifts in a central pattern generator network. *J Neurosci* 1995;15:342–58.
- Hodgkin AL, Huxley AF. A quantitative description of membrane current and its application to conduction and excitation in nerve. *J Physiol* 1952;117:500–44.
- Luo CH, Rudy Y. A dynamic model of the cardiac ventricular action potential I. Simulations of ionic currents and concentration changes. *Circ Res* 1994;74:1071–96.
- Marom S, Levitan IB. State-dependent inactivation of the Kv-3 potassium channel. *Biophys J* 1994;67:579–89.
- McAnelly L, Zakon HH. Protein kinase A activation increases sodium current magnitude in the electric organ of *Sternopygus*. *J Neurosci* 1996;16:4383–8.
- Moré JJ. The Levenberg–Marquardt algorithm: implementation and theory. In: Watson GA, editor. *Numerical Analysis Lecture Notes in Mathematics*, vol. 630. Berlin: Springer-Verlag, 1977:105–16.
- Rinzel J, Lee YS. Dissection of a model for neuronal parabolic bursting. *J Math Biol* 1987;25:653–75.
- Tóth TI, Crunelli V. A numerical procedure to estimate kinetic and steady-state characteristics of inactivating ionic currents. *J Neurosci Methods* 1995;63:1–12.
- Willms AR, Baro DJ, Harris-Warrick RM, Guckenheimer J. An improved parameter estimation method for Hodgkin–Huxley models. *J Comput Neurosci* 1999;6:145–68.