Controlled Hidden Markov Models for Dynamically Adapting Patch Clamp Experiment to Estimate Nernst Potential of Single-Ion Channels

Vikram Krishnamurthy*, *Fellow, IEEE*, and G. George Yin, *Fellow, IEEE*

***Abstract*—This paper presents novel kernel-based stochastic learning algorithms for controlling the kinetics of single-ion channels in a patch clamp experiment. The algorithms yield efficient estimates of the equilibrium (Nernst) potential of an ion channel. The equilibrium potential of an ion channel is the applied external potential difference required to maintain electrochemical equilibrium across the ion channel. The algorithm adaptively controls the exploration of the learning algorithm to achieve an optimal balance between exploration and exploitation. An important feature of the resulting algorithm is that it is guaranteed to minimize the experimental effort. We illustrate the efficiency of the algorithms for the experimentally determined current voltage curve of a bi-ionic single potassium ion channel.**

***Index Terms*—Adaptive exploration, discrete stochastic optimization, equilibrium potential, hidden Markov model (HMM), ion channel, patch clamp experiment.**

I. Introduction

CELL membranes of all animal, plant, and bacterial cells are composed of a lipid bilayer (two layers of lipid molecules) making them impermeable to ions. Ion transportation into and out of cells is mediated by large membrane protein molecules called ion channels. Ion channels are biological nanotubes whose opening and closing may be intrinsic or gated. Their primary function is to facilitate the diffusion of ions across the cell membrane.

The measurement of ionic currents flowing through single-ion channels in cell membranes has been made possible by the giga-seal *patch clamp* technique [1], [2] for which the authors of [1] won the 1991 Nobel Prize in medicine. Defects in ion channel functions is the cause of several physiological and degenerative diseases, see [3]. Understanding the functioning of ion channels is a fundamental problem in biology. There are two broad classes of problems that study the dynamics of ion channels—we refer the reader to [4] and the special issue [5] for a tutorial exposition. At the femtosecond time scale and angstrom unit spatial scale, the permeation problem deals with modeling the dynamics of individual ions propagating through the protein nanotube [6]. At the millisecond time scale, the kinetic model of an ion channel deals with probabilistic models for the gating of the ion channel. In this paper we focus on controlling the kinetic behavior of ion channels. Typically a gated ion channel has two states—a "closed" state that does not allow ions to flow through, and an "open" state that does allow ions to flow through. In the open state, the single-ion channel current denoted by $I(v)$, where v denotes the applied external voltage to the patch clamp experiment, is typically of the order of picoamps (i.e., 10^{-12} amps). The measured single-ion channel current (obtained by sampling at 10 kHz) is obfuscated by large amounts of thermal noise. In several recent papers [7], [8], hidden Markov models (HMMs) have been used to model these ion channel currents and obtain an estimate $\hat{I}(v)$ of the open-state current level $I(v)$.

In characterizing different types of ion channels, current–voltage (*I–V*) curves are widely used by neurobiologists, since they yield a unique signature of a particular ion channel, revealing its operating characteristics. The *I–V* curve represents the variation of the open-state current level $I(v)$ of the single-ion channel as a function of the applied voltage value v. The zero point of the *I–V* curve, i.e., the voltage v^* at which the open-state current level $I(v^*)$ is zero, is known as the *equilibrium potential* (Nernst potential). The equilibrium potential gives information about the relative concentrations at the two faces of the ionic channel. The value of the open-state current level $I(v)$ is described by the Nernst–Planck equation that combines Ohm's and Fick's laws. Once the equilibrium potential v^* is determined, by computing estimates of the current $I(v)$ at several values of v around v^*, the experimenter can straightforwardly determine if the *I–V* response of the single-ion channel is ohmic (linear) or not.

In this paper, we present novel learning algorithms to adaptively control the applied voltage v of a patch clamp experiment in order to dynamically learn the equilibrium potential v^* of a single-ion channel with minimum experimental effort. We formulate the problem as the solution v^* of a *discrete stochastic optimization problem* of the form

$$\text{Compute } v^* = \arg\min_{v \in V} \left[\mathbf{E} \left\{ \hat{I}(v) \right\} \right]^2$$

where V is a discrete set of possible applied voltages to the patch clamp experiment, and for any fixed voltage $v \in V$, $\hat{I}(v)$ denotes the maximum likelihood estimate obtained from experimental observations in a patch clamp experiment.

Manuscript received April 25, 2005; revised October 19, 2005. This work was supported in part by the Natural Sciences and Engineering Research Council of Canada (NSERC) and in part by the U.S. National Science Foundation. *Asterisk indicates corresponding author.*

*V. Krishnamurthy is with the Department of Electrical and Computer Engineering, University of British Columbia, Vancouver V6T 1Z4, Canada (e-mail: vikramk@ece.ubc.ca).

G. G. Yin is with the Department of Mathematics, Wayne State University, Detroit, MI 48202 USA (e-mail: gyin@math.wayne.edu).

Digital Object Identifier 10.1109/TNB.2006.875038

1536-1241/$20.00 © 2006 IEEE

We also use the learning algorithm to estimate the applied voltage at which the operating signal-to-noise ratio of the single-ion channel current is maximized. With $\hat{\sigma}^2(v)$ denoting the estimate of the measured ion channel noise variance $\sigma^2(v)$, we use the learning algorithm to estimate the potential v^* at which the $\sigma^2(v)$ is minimized.

There are several methods that can be used to solve the above discrete stochastic optimization problem; see [9] for a recent survey. In our recent paper [10], we presented a *discrete stochastic approximation algorithm* for estimating the equilibrium potential. The idea of discrete stochastic approximation [9], [11] is to dynamically control (schedule) the choice of voltages at which the patch clamp experiment operates to provide more observations near v^* and less in other areas efficiently obtain the equilibrium potential.

However, a key issue with the discrete stochastic approximation algorithms in [9] and [11] is that they are designed for *simulation-based optimization* and not control of *real-life experiments*. That is, they make the assumption that there is no effort involved in evaluating $\hat{I}(v)$ for any chosen v. In estimating the equilibrium potential of an ion channel, each evaluation of $\hat{I}(v)$ is expensive, since it involves running a patch clamp experiment. Thus, what is needed is a stochastic approximation algorithm that optimally trades off *exploration* with *exploitation*, as we now explain.

Exploration refers to the fact that the learning algorithm needs to examine other candidate solutions compared to its current solution in order to determine if these candidates are better than the current solution. However, to estimate the equilibrium potential, exploration is expensive, since evaluation of an alternative candidate involves running a patch clamp experiment on this alternative candidate. While exploration is expensive, the advantage of exploration is that it allows the algorithm to rapidly converge to the vicinity of the optimum v^*. This is intuitively clear, since, the more one explores, the quicker one gains information about which candidates are good. Hence, there is a tradeoff between convergence rate (advantage of exploration) versus exploration cost (disadvantage of exploration). *Exploitation* refers to the fact that if one is confident that the current estimate is close to the optimum, then it is better to exploit this estimate—i.e., keep running experiments at this estimate to keep the system performance close to optimal.

All learning algorithms have an inherent tradeoff between exploration and exploitation. Typically during initial iterations of a learning algorithm, it is desirable to aggressively explore more candidates, since one is uncertain how good the current estimate is. After more confidence has been obtained about the candidates, it is desirable to reduce exploration and exploit the best candidates. The key idea of this paper is to introduce a theoretical framework and practical algorithms that determine an optimal tradeoff between exploration and exploration in controlling patch clamp experiments to determine the equilibrium potential v^* of single-ion channels. Fig. 1 schematically illustrates our kernel-based learning algorithm—details are presented in Section III-D. The key idea is to adaptively update the exploration probabilities based on the ion channel behavior via a kernel-based function. As the learning algorithm gains confidence of the ion channel behavior, the exploration probability

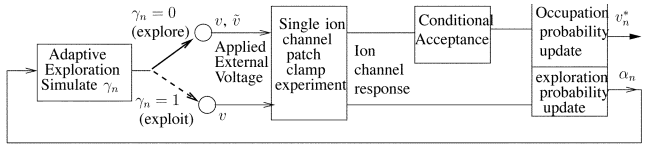

Fig. 1. Schematic setup of kernel-based learning algorithm for controlling single-ion channel patch clamp. The variable $\gamma_n \in \{0, 1\}$ is randomly generated according to exploration probabilities $(\alpha_n, 1 - 2\alpha_n)$, respectively. Algorithm 2 in Section III-D presents details.

α_n decreases to zero, so that exploration is minimized and the best estimate is increasingly exploited to optimize the behavior of the patch clamp experiment.

The rest of this paper is organized as follows. Section II describes the patch clamp experiment setup and formally presents a stochastic signal model (HMM) for a channel ion current. In Section III the learning algorithm is described. Finally, Section IV illustrates the performance of these algorithms in computer simulations. These simulations show that using the learning algorithms result in a remarkable improvement in overall efficiency. To ease the presentation, all proofs of results are placed in an appendix at the end of the paper.

II. Ion Channel Model and Current Voltage (I–V) Curve of Single Ion Channel

In this section we give a precise formulation of the ion channel current signal model and the experimental setup. This allows us to formulate the equilibrium potential learning problem in Section III mathematically.

A. I–V Curve, Equilibrium Potential

The single-ion channel current measurement from a patch clamp experiment (after anti-aliasing filtering and sampling) reveals a piecewise constant discrete time ion channel current that randomly jumps between two values—zero amperes, which denotes the *closed state* of the channel, and $I(v)$ amperes, which denotes the *open* state. $I(v)$ is called the *open-state* current level. The open-state current level $I(v)$ depends on the voltage v that is applied by the experimenter to the ion channel.

Fig. 2 shows the experimentally determined *I–V* curve of a potassium ion channel. In patch clamp experiments, the applied voltage v is usually chosen from a finite set, i.e.,

$$v \in V = \{\theta(1), \ldots, \theta(M)\}$$

denotes the finite set of possible voltage values that the experimenter can choose. For example, for the experimentally determined *I–V* curve for the potassium ion channel of Fig. 2, $M = 31$ and $\theta(i)$ are uniformly spaced in 10-mV steps from $\theta(1) = -150$ mV and $\theta(M) = 150$ mV. The shape of the *I–V* curve illustrated in the figure incorporates several features observed in many experimentally observed *I–V* curves: linear and nonlinear segments, a saturation of current with increasing driving force (voltage), and an asymmetry between outward and inward currents (i.e., the *I–V* curve for $I(v) > 0$ and $I(v) < 0$, respectively). The equilibrium potential (voltage at which vertical broken line intersects the v axis) is $v^* = -30$ mV.

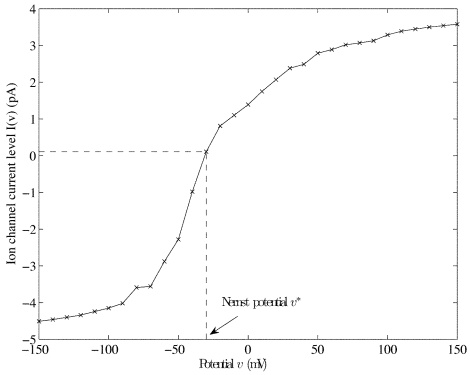

Fig. 2. Experimentally determined *I–V* curve of potassium ion channel in bi-ionic condition with 100 mM K external and 100 mM Tl (Thallium) internal. The × represents *I–V* experimentally determined values of $I(v)$ versus v at 10-mV intervals ranging from -150 mV to 150 mV. The equilibrium potential is the voltage v at which $I(v)$ is closest to zero—approximately 30 mV (dashed lines).

Markov Model for Ion Channel Current: Suppose the patch clamp experiment is conducted with a voltage v applied across the ion channel. Then, as described in [7], [8], the ion channel current $\{i_n(v)\}$, can be modeled as a three-state homogeneous first-order Markov chain. with state space $\{0_g, 0_b, I(v)\}$ corresponding to the physical states of *gap mode*, *burst-mode-closed* and *burst-mode-open*. For convenience, we will refer to the burst mode closed and burst-mode-open states as the *open* and *closed* states, respectively. In the gap mode and the closed state the ion channel current is zero. In the open state, the ion channel current has a value of $I(v)$. The transition probability matrix $A(v)$ of the Markov chain $\{i_n(v)\}$ has elements $a_{ij}(v) = P(i_{n+1}(v) = j | i_n(v) = i)$. Note that in general, the applied voltage v affects both the transition probabilities and state levels of the ion channel current $\{i_n(v)\}$.

Before describing our adaptive learning and control methodology, we briefly outline two existing methods for estimating the equilibrium potential and *I–V* curves.

1) *Filtered trace brute force approach*: This method is widely used by neurobiologists to estimate the *I–V* curve and determine the equilibrium potential. First, several possible candidate voltages v are chosen. For each of these voltages v, a patch clamp experiment is run for a long time—typically several minutes. The *measured* ion channel current sequence $\{y_n(v)\}$ is then heavily low-pass filtered—and an estimate of the current level $I(v)$ is determined visually. The equilibrium potential is typically deduced by measuring the current levels near its vicinity and then linearly extrapolating the data points. This approach is highly unreliable when the signal to noise ratio is low—which is almost always the case in typical patch clamp experiments. It is also very expensive, since the experiments are not controlled to extract maximum information about the equilibrium potential.
2) *HMM Brute Force Approach*: Let $\{y_n(v)\}$ denote the measured noisy ion channel current at the electrode when conducting a patch clamp experiment:

$$y_n(v) = i_n(v) + w_n(v), \quad n = 1, 2, \ldots \tag{1}$$

Here $\{w_n(v)\}$ is thermal noise and is modeled as zero mean white Gaussian noise with variance $\sigma^2(v)$. The ionic channels can be modeled extremely well by an HMM [7], [8] parameterized by

$$\lambda(v) = \left\{A(v), I(v), \sigma^2(v)\right\}. \tag{2}$$

Hence, an obvious improvement to the above approach is to replace the visual estimation step by an HMM parameter estimation algorithm. For each possible voltage $v \in V$, run an independent patch clamp experiment to gather the sample path $\{y_1(v), y_2(v), \ldots, y_\Delta(v)\}$ for a very large batch size Δ. Compute the MLE $\hat{I}(v)$ via an HMM parameter estimator. Finally pick $\hat{v}^* = \arg\min_{v \in V} |\hat{I}(v)|^2$. Since for any fixed $v \in V$, the MLE $\hat{I}(v)$ is strongly consistent [12], $\hat{I}(v) \to I(v)$ w.p. 1, as the batch size $\Delta \to \infty$. This and the finiteness of V imply that as $\Delta \to \infty$,

$$\arg\min_{v \in V} \left(\hat{I}(v)\right)^2 \to \arg\min_{v \in V} \left(I(v)\right)^2 \text{ w.p.1.} \tag{3}$$

Thus in principle, the above brute force simulation method can solve the discrete stochastic optimization problem (6) for large batch size Δ and the estimate is *consistent*, i.e., (3) holds. However, the method is highly inefficient, since $\hat{I}(v)$ needs to be evaluated for each $v \in V$. The evaluations of $\hat{I}(v)$ for $v \neq v^*$ are wasted because they contribute nothing to the estimation of the ion channel current $i(v^*)$ at the equilibrium potential v^*. Also the brute force approach does not exploit the fact that the *I–V* curve is monotonically increasing.

For $M = 80$, the brute force approach to compute the equilibrium potential requires conducting a total of 80 experiments, one at each value of $v \in V$. For a typical sampling rate of 100 kHz and 10 min of data per experiment, 6×10^7 observations are obtained per experiment which need to be processed by an HMM MLE estimator.

Examples: To illustrate the *ad hoc* nature of approach 1) above, we used the potassium channel *I–V* curve in Fig. 2 to generate computer-simulated ion channel currents (denoted as $i_n(v)$), and the noisy measured ion channel current (denoted as $y_n(v)$) from a patch clamp experiment for 500 time points—Section IV contains precise details. Figs. 3, 4, and 5 show the plots of the sequences $\{i_n(v)\}$ and $\{y_n(v)\}$ for three different applied voltages $v = -30$ mV, $v = 0$ mV, and $v = 30$ mV, respectively. For each of these voltages v, the ion channel current $i_n(v)$ was simulated as Markov chain) with state levels 0 (off-state) and $I(v)$ (on-state). The on-state level $I(v)$ is chosen using the *I–* curve of Fig. 2. From the traces of the measured currents $\{y_n(v)\}$ in Figs. 3, 4, and 5, it is virtually impossible to determine the on-state level $I(v)$ by merely eye-balling the plots. However, as demonstrated in [7] and [8], accurate maximum likelihood estimates of $I(v)$ for various values of applied voltage v can be obtained using HMM signal processing methods. Then using the HMM brute force approach 2) above, the equilibrium potential i.e., the voltage v^* at which $I(v^*) = 0$ can be determined by exhaustive experimentation and computation over all possible voltages v.

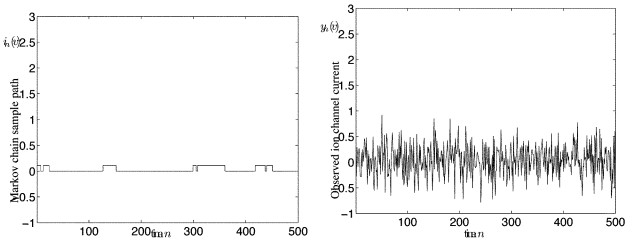

Fig. 3. Synthetic generated patch clamp currents for bi-ionic potassium channel with 100 mM K external and 100 mM Tl internal. Applied voltage $v = -30$ mV, $I(v) = 0.11$ pA.

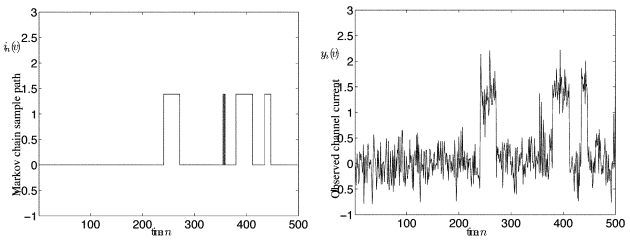

Fig. 4. Synthetic generated patch clamp currents for bi-ionic potassium channel with 100 mM K external and 100 mM Tl internal. Applied voltage $v = 0$ mV, $I(v) = 1.39$ pA.

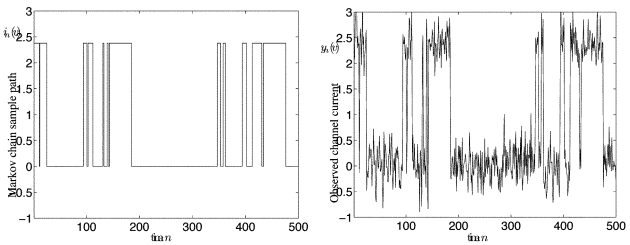

Fig. 5. Synthetic generated patch clamp currents for bi-ionic potassium channel with 100 mM K external and 100 mM Tl internal. Applied voltage $v = 30$ mV, $I(v) = 2.38$ pA.

Our Approach: In Section III we present a novel algorithm that dynamically learns the equilibrium potential with minimum effort. The key feature of the algorithm is its tradeoff between exploration and exploitation. The adaptive learning algorithms we propose are based on a novel extension of discrete stochastic approximation algorithms that have recently been developed in the operations research literature [9], [11]. The basic idea is to generate a reversible homogeneous Markov chain taking values which spends more time at the zero point v^* than at any other voltage v and to control how much exploration this Markov chain does to minimize the experimental effort.

B. Measurements of Ionic Channel Current

The measurement of ionic currents flowing through single-ion channels in cell membranes has been made possible by the giga-seal patch clamp technique [1], [2]. A tight seal between the rim of the electrode tip and the cell membrane drastically reduces the leakage current and extraneous background noise, enabling the resolution of the discrete changes in conductance, which occur when single-ion channels open or close. To record currents from single-ion channels in the *cell attached patch* configuration, the tip, an electrode with the diameter of about 1 μm, is pushed against the surface of a cell, and then a tight seal is formed between the rim of the electrode tip and the cell membrane. A patch of the membrane surrounded by the electrode tip usually contains one or more single-ion channels. The current flowing from the inside of the cell to the tip of the electrode through a single-ion channel is monitored. This is known as "cell-attached" configuration of patch clamp techniques for measuring ion channel currents through a single-ion channel.

The ionic strength in the electrode is made the same as that in the outside of the cell. Let E_i and E_o, respectively, denote the resting membrane potential and the potential applied to the electrode. If E_o is identical to the membrane potential, there will be no potential gradient across the membrane patch confined by the tip of the electrode. Let c_i denote the intracellular ionic concentration and c_o the ionic concentration in the electrode. Here the intracellular concentration c_i inside the cell is unknown as is the resting membrane potential E_i. c_o and E_o are set by the experimenter and are known. Let $v = E_o - E_i$ denote the potential gradient. Both the potential gradient v and concentration gradient $c_o - c_i$ drive ions across an ion channel resulting in an ion channel current $\{i_n(v)\}$.

The potential E_o (and hence potential difference v) is adjusted experimentally until the current $I(v)$ goes to zero. This voltage v^* at which the current $I(v^*)$ vanishes is called the *equilibrium potential* and satisfies the Nernst equation

$$v^* = -\frac{kT}{e} \ln \frac{c_o}{c_i} = -59 \log_{10} \frac{c_o}{c_i} (\text{mV}) \tag{4}$$

where $e = 1.6 \times 10^{-19}$ C denotes the charge of an electron, k denotes Boltzmann's constant, and T denotes the absolute temperature. The Nernst equation (4) gives the potential difference v required to maintain electrochemical equilibrium when the concentrations are different on the two faces of the membrane.

Another widely used configuration is the *excised patch* experiment. In this excised configuration of patch clamp technique, both c_i and c_o, as well as E_i and E_o, are known. Again it is important to determine the reversal potential E_{rev} accurately. For example, the ionic solutions in the electrode and the bath may contain a mixture of ionic species, such as Na^+, Cl^-, and K^+. The ion channel contained in the membrane patch may be permeable to, for example, Na^+ and K^+, as the case with the acetylcholine receptor or Na^+ and Cl^- as the case with certain mutant glycine receptors. It is important to deduce the permeability ratio of the ion channel, namely, the ratio between the number of Na^+ ions and K^+ ions that move across the ion channel per unit time. This ratio can be deduced by accurately determining the equilibrium potential using the Goldman–Hodgkin–Katz voltage equation [13] of the form

$$E_{\text{rev}} = \frac{kT}{e} \ln \frac{P_K[K]_o + P_{\text{Na}}[\text{Na}]_o + P_{\text{Cl}}[\text{Cl}]_o}{P_{\text{K}}[\text{K}]_i + P_{\text{Na}}[\text{Na}]_i + P_{\text{Cl}}[\text{Cl}]_i}. \tag{5}$$

Here P_{K}, P_{Na}, and P_{Cl} refer to the permeability of K^+, Na^+, and Cl^-, respectively.

III. Patch Clamp Control Algorithms for Estimating Equilibrium Potential

Here we formulate the ion channel control problem as a discrete stochastic optimization problem. A novel discrete stochastic approximation algorithm is presented for efficiently solving this problem.

A. Formulation as Discrete Stochastic Optimization Problem

As explained in the HMM formulation above, due to the presence of large amounts of thermal noise, $I(v)$ cannot be exactly evaluated and only unbiased estimates $\hat{I}(v)$ are available. Thus, computing the equilibrium potential is equivalent to the following discrete stochastic optimization problem:

$$\text{Compute } v^* = \arg\min_{v \in V} \left[\mathbf{E} \left\{ \hat{I}(v) \right\} \right]^2 \tag{6}$$

where $\hat{I}(v)$ is the MLE of the parameter $I(v)$ of the HMM. Note that the equilibrium potential v^* (zero crossing point) does not necessarily belong to the discrete set V—instead we will find the point in V that is closest to v^* (with resolution $\theta(2) - \theta(1)$). With slight abuse of notation we will denote the element in V closest to the equilibrium potential as v^*. Our choice of using a quadratic objective function (instead of, for example, $\arg\min_{v \in V} |I(v)|$) is because it allows us to conveniently reformulate the optimization problem to be linear in the expected value—see (11) below.

The most popular way of computing the maximum likelihood estimate (MLE) $I(v)$ is via the expectation maximization (EM) algorithm (Baum–Welch equations). Let $\hat{I}_\Delta(v)$ denote MLE of $I(v)$ based on the Δ-point measured channel current sequence $(y_1(v), \ldots, y_\Delta(v))$. For sufficiently large batch size Δ of observations, due to the asymptotic normality of the MLE for an HMM [14]

$$\sqrt{\Delta} \left(\hat{I}_\Delta(v) - I(v) \right) \sim N(0, \Sigma(v)) \tag{7}$$

where $\Sigma^{-1}(v)$ is the Fisher information matrix. Thus, asymptotically $\hat{I}_\Delta(v)$ is an unbiased estimator of $I(v)$, i.e., $\mathbf{E}\{\hat{I}_\Delta(v)\} = I(v)$ where $\mathbf{E}\{\cdot\}$ denotes the mathematical expectation operator. Since for an HMM, no closed form expression is available for $\Sigma^{-1}(v)$ in (7), the above expectation cannot be evaluated analytically. This motivates the need to develop a simulation based (stochastic approximation) algorithm [15].

B. Learning Algorithm for Controlling Patch Clamp

In this subsection we present a generic learning algorithm for estimating the equilibrium potential v^*. We use the term "generic," since the algorithm has a parameter α where $0 < \alpha \leq 1/2$ that determines how much the algorithm explores the candidate solutions. Later, we will denote α as the exploration probability. Algorithm 1 below is described for a fixed α. In subsequent sections we will show how the exploration probability α can be chosen and dynamically adjusted as the response of the ion channel is better learned. It is important to note that for the special case $\alpha = 0.5$, the algorithm specializes to that proposed in [10] and [11]. However, as described below, $\alpha = 0.5$ is expensive from an experimental point of view in the sense that patch clamp experiments need to be conducted at each iteration on alternative candidate solutions.

Let $k = 1, 2, \ldots$ denote discrete time. The generic learning algorithm described below is recursive and requires conducting patch clamp experiments on batches of data. Since the patch clamp experiments will be conducted over batches of data, it is convenient to introduce the following notation. Group the discrete time into batches of length Δ—typically $\Delta = 10\,000$ in experiments. We use the index $n = 1, 2, \ldots$ to denote batch number. Thus, batch n comprises the Δ discrete time instants $k \in \{n\Delta, n\Delta + 1, \ldots, (n+1)\Delta - 1\}$.

In the algorithm below, the process $\{X_n, n = 1, 2 \ldots\}$ denotes the state of the algorithm. Indeed, as shown later, X_n is a finite-state Markov chain on the state space V. The most important feature of the learning algorithm below is that its state $\{X_n\}$ spends more time at the global optimum v^* than any other value in V—indeed this feature will be used to estimate v^*. To quantify the amount of time the algorithm state $\{X_n\}$ spends at each candidate value in V, define the vector of empirical occupation probabilities of $\{X_n\}$ as $\hat{\pi}_n = (\hat{\pi}_n(1), \ldots, \hat{\pi}_n(M))'$, $n = 1, 2, \ldots$. The elements $m = 1, 2, \ldots, M$ of this occupation probability vector are

$$\hat{\pi}_n(m) = \frac{\#\text{ of times } X \text{ visits state } m \text{ in batches 1 to } n}{n}$$

For the state X_n, define the neighborhood set

$$\mathcal{N}_{X_n} = \begin{cases} \{2\} & \text{if } X_n = 1 \\ \{X_{M-1}\} & \text{if } X_n = M \\ \{X_n - 1, X_n + 1\} & \text{otherwise} \end{cases} \tag{8}$$

Finally, denote the M-dimensional standard unit vectors by e_m, $m = 1, \ldots, M$, where

$$e_m = [0 \quad \cdots \quad 0 \quad 1 \quad 0 \quad \cdots \quad 0]' \tag{9}$$

with 1 in the mth position and zeros elsewhere.

It is convenient to formulate the cost function to be minimized as $\min_{v \in V} \mathbf{E}\{C_n(v)\}$. Equation (6) can be converted to this form as follows. Let $\hat{I}_1(v)$, $\hat{I}_2(v)$ be two statistically independent HMM-ML estimates of $I(v)$, each obtained from a Δ length patch clamp experiment batch of HMM observations. Then assuming that the ML estimates are unbiased, defining

$$\hat{C}_n(v) = \hat{I}_{n,1}(v) \hat{I}_{n,2}(v), \tag{10}$$

it straightforwardly follows that for any fixed $v \in M$

$$\mathbf{E} \left\{ \hat{C}_n(v) \right\} = \left[\mathbf{E} \left\{ \hat{I}_n(v) \right\} \right]^2 = |I(v)|^2 \tag{11}$$

The discrete stochastic approximation algorithm we propose is as follows.

Algorithm 1: [Constrained Exploration Cost Stochastic Approximation Algorithm for Estimating Equilibrium Potential]

- `Step 0`: (Initialization.) At batch time $N = 0$, initialize state of the algorithm $X_0 \in \{1, \ldots, M\}$ randomly.

Initialize state occupation probabilities $\hat{\pi}_0 = e_{X_0}$.
Initialize equilibrium potential estimate as $\widehat{v}_0^* = \theta(X_0)$.

- `Step 1`: (Adaptive Sampling and Exploration.) At time n, given current algorithm state X_n, apply voltage $v = v(X_n)$ to patch clamp experiment and evaluate $C_n(X_n)$. Then perform the following two-level sampling procedure: simulate a Bernoulli random variable $\gamma_n \in \{0,1\}$ with probabilities $P(\gamma_n = 0) = 2\alpha$ and $P(\gamma_n = 1) = 1 - 2\alpha$, where $0 < \alpha \leq 1/2$.
  — If $\gamma_n = 0$, perform exploration as follows: generate an alternative candidate state $\tilde{X}_n$ by sampling uniformly from the neighborhood $\mathcal{N}_{X_n}$ of current state X_n. Apply voltage $\tilde{v} = v(\tilde{X}_n)$ to patch clamp experiment and evaluate $C_n(\tilde{X}_n)$ and go to Step 2.
  — If $\gamma_n = 1$ (perform no exploration), go to Step 3.
- `Step 2`: (Conditional Acceptance Test.) If $\hat{C}_n(\tilde{v}) < \hat{C}_n(v)$, set $X_{n+1} = \tilde{X}_n$, else, set $X_{n+1} = X_n$.
- `Step 3`: Update empirical state occupation probabilities π_N as

$$\hat{\pi}_{n+1} = \hat{\pi}_n + \mu_n \left(e_{X_{n+1}} - \hat{\pi}_n\right), \quad \hat{\pi}_0 = e_{X_0}. \tag{12}$$

- `Step 4`: (Update estimate of equilibrium potential.) $\widehat{v}_n^* = \theta(m^*)$ where $m^* = \arg\max_{m\in\{1,\ldots,M\}} \hat{\pi}_{n+1}(m)$, set $n \to n+1$, go to Step 1.

Processing a batch size $\Delta = 10\,000$ takes negligible time compared to the data acquisition time. Actually in the theorem below, we only require that $\hat{I}_N^{(1)}(v)$, $\hat{I}_N^{(2)}(v)$ have symmetric probability density functions with mean $I(v)$. We assume the following.

(A) The batch size Δ is sufficiently large (e.g $\Delta = 10\,000$) so that due to (7), the MLE estimates $\hat{I}_N^{(1)}(v)$, $\hat{I}_N^{(2)}(v)$ are $N(I(v), \Sigma(v))$ Gaussian random variables.
(B) For any $m \in \{1, \ldots, M-1\}$

$$\begin{aligned} &I^2(\theta(m+1)) > I^2(\theta(m)) \\ &\qquad \Longrightarrow P\left(\hat{C}(\theta(m+1)) > \hat{C}(\theta(m))\right) > 0.5 \\ &I^2(\theta(m+1)) < I^2(\theta(m)) \\ &\qquad \Longrightarrow P\left(\hat{C}(\theta(m+1)) > \hat{C}(\theta(m))\right) < 0.5 \end{aligned}.$$

The following theorem proved in the appendix shows that the above algorithm is convergent to the equilibrium potential v^* of the ion channel irrespective of the choice of the exploration parameter α.

Theorem 1: Under conditions (N) and (O), the sequence $\{\theta(X_n)\}$ generated by Algorithm 1 is a homogeneous, aperiodic, irreducible, and reversible Markov chain with state space V. The stationary distribution π of the Markov chain X_n is independent of α, and the occupation probabilities $\hat{\pi}_n$ converge to π geometrically fast. Furthermore, Algorithm 1 is attracted to the equilibrium potential v^*, i.e., for sufficiently large n, the sequence $\{\theta(X_n)\}$ spends more time at v^* than any other states. (Equivalently, if $\theta(m^*) = v^*$, then $\hat{\pi}_n(m^*) > \hat{\pi}_n(j)$ for $j \in \{1, \ldots, M\} - \{m^*\}$.)

C. Constrained Exploration Equilibrium Potential Estimation

Our aim is to choose the exploration probability α in Algorithm 1 above so as to maximize the exploration reward while satisfying a constraint on the exploration cost. The exploration reward and exploration cost are defined as follows.

Exploration Cost: In each iteration of the above algorithm if $\gamma_n = 0$ in Step 1, exploration of an alternative candidate is performed. Define the average exploration cost of the algorithm over N iterations as (with γ shown as an explicit function of α)

$$J(\alpha) = \lim_{N\to\infty} \frac{1}{N} \sum_{n=1}^{N} \mathbf{1}\left(\gamma_n^{(\alpha)} = 0\right). \tag{13}$$

The summation denotes the number of times $\gamma_n = 0$, $n = 1, \ldots, N$, i.e., the number of times in N iterations that exploration is performed. Hence, the above cost reflects the average exploration cost per iteration of running a patch clamp experiment on an alternative candidate voltage. Since the sequence $\{\gamma_N\}$ consists of independent identically distributed (i.i.d.) Bernoulli random variables, by the strong law or large numbers, the right-hand side of (13) converges to $\alpha = P(\gamma_1 = 0)$. Thus, the average exploration cost per iteration of the algorithm is

$$J(\alpha) = \alpha. \tag{14}$$

Exploration Reward: As shown in Theorem 1, $\hat{\pi}_n$ converges as $n \to \infty$ to the invariant distribution π of the Markov chain $\{X_n\}$. As is standard in statistical inference, the *asymptotic rate of convergence* of $\hat{\pi}_n$ to π is measured by the asymptotic covariance

$$\Gamma(\alpha) = \lim_{N\to\infty} N\mathbf{cov}\left\{\hat{\pi}_N^{\alpha}\right\} \tag{15}$$

The smaller this asymptotic covariance, the faster the estimate $\widehat{v}_n^*$ generated by Algorithm 1 converges to v^*. Thus, $\Gamma(\alpha)$ in (15) defines the exploration reward of Algorithm 1 when the exploration parameter is set to α with $0 < \alpha \leq 1/2$.

We now consider two schemes for choosing α. In the first scheme, the exploration cost is constrained as

$$\lim_{N\to\infty} \frac{1}{N} \sum_{n=1}^{N} \mathbf{1}\left(\gamma_n^{(\alpha)} = 0\right) = \alpha \leq \bar{\alpha} \tag{16}$$

Here $0 < \bar{\alpha} < 1/2$ denotes a user defined exploration cost. Thus, we have the following constrained exploration discrete stochastic optimization problem: Consider Algorithm 1 with its attraction property $\hat{\pi}_N(v^*) > \hat{\pi}_N(v)$ for $v \in M - \{v^*\}$ to

$$\text{Compute } v^* = \arg\min_{v\in V} \left[\mathbf{E}\left\{\hat{I}(v)\right\}\right]^2 \tag{17}$$

Optimize the following exploration/exploitation optimization tradeoff of Algorithm 1

$$\min \Gamma(\alpha) \quad \text{subject to } \alpha \leq \bar{\alpha}. \tag{18}$$

Theorem 2: The asymptotic covariance $\Gamma(\alpha)$ of $\hat{\pi}_n$ generated by Algorithm 1 is a decreasing function of α. Thus, for $\alpha \in [0, \bar{\alpha}]$, the exploration reward is maximized by choosing $\alpha = \bar{\alpha}$.

The above theorem states that the exploration reward is an increasing function of α. On the other hand the exploration cost $J(\alpha)$ in (14) is an increasing function of α. Thus, there is an inherent tradeoff between the exploration cost and exploration reward in choosing α.

D. Kernel-Based Adaptive Exploration Learning Algorithm

Section III-C presented an exploration constrained algorithm for estimating the equilibrium potential v^* where the exploration probability α was held constant. A natural extension of idea is to vary α dynamically over time. This is consistent with the idea that for early iterations, α should be chosen close to 0.5 to maximize exploration, while after several iterations once confidence has been gained about the behavior of the patch clamp experiment, the exploration probability α can be decreased to reduce the exploration cost. The aim in this section is to dynamically adapt the exploration probability α with time. This adaptation of α is done via a kernel-based learning algorithm. The key idea is Step 3 below, where a kernel-based update on α_n is used.

Algorithm 2: [Kernel-Based Adaptive Exploration/Exploitation Stochastic Approximation Algorithm for Estimating Equilibrium Potential]

- `Step 0`: Identical to Algorithm 1.
- `Step 1`: (Adaptive Sampling and Exploration.) At time n, given current algorithm state X_n, apply voltage $v = v(X_n)$ to patch clamp experiment and evaluate $C_n(X_n)$. Then perform the following two-level sampling procedure: simulate a Bernoulli random variable $\gamma_n \in \{0, 1\}$ with probabilities $P(\gamma_n = 0) = 2\alpha_n$ and $P(\gamma_n = 1) = 1 - 2\alpha_n$, where $0 \leq \alpha_n \leq 1/2$.
  - If $\gamma_n = 0$, perform exploration as follows: generate an alternative candidate state $\tilde{X}_n$ by sampling uniformly from the neighborhood $\mathcal{N}_{X_n}$ of current state X_n. Apply voltage $\tilde{v} = v(\tilde{X}_n)$ to patch clamp experiment and evaluate $C_n(\tilde{X}_n)$ and go to Step 2.
  - If $\gamma_n = 1$ (perform no exploration), go to Step 3.
- `Step 2`: (Conditional Acceptance Test.) If $\hat{C}_n(\tilde{v}) < \hat{C}_n(v)$, set $X_{n+1} = \tilde{X}_n$, else, set $X_{n+1} = X_n$.
- `Step 3`: (Kernel-based exploration probability α_n update and empirical state occupation probabilities π_n update.)

$$\hat{\pi}_{n+1} = \hat{\pi}_n + \mu_n \left(e_{X_{n+1}} - \hat{\pi}_n \right), \quad \hat{\pi}_0 = e_{X_0}, \tag{19}$$

$$T_{n+1} = T_n + \mu_n \left(e_{X_{n+1}} - \hat{\pi}_n - T_n \right), \tag{20}$$

$$\alpha_{n+1} = 0.5 \left[1 - K \left(\frac{T_{n+1}}{\delta_{n+1}} \right) \right] \tag{21}$$

where $\delta_n \geq 0$, $\delta_n \to 0$, and $\mu_n/\delta_n \to 0$ as $n \to \infty$. [In the actual implementation, we often choose δ to be a small constant to simplify the computation.]

- `Step 4`: (Update estimate of equilibrium potential.) $\widehat{v}_n^* = \theta(m^*)$ where $m^* = \arg\max_{m \in \{1,\ldots,M\}} \hat{\pi}_{n+1}(m)$, set $n \to n+1$, go to Step 1.

In the above algorithm, for any $x \in \mathbb{R}^M$ the kernel

$$K(x) = \begin{cases} (1 - x'x) & \text{if } \|x\|_2 = \sqrt{x'x} \leq 1 \\ 0 & \text{otherwise} \end{cases}.$$

The intuition in the above kernel-based learning in Step 3 is as follows: in early iterations as the algorithm knows little about the behavior of the ion channel. Hence, the average error T_n between e_{X_n} and the empirical occupation $\hat{\pi}_n$ is large. When this average error T_n is large, then $K(T_{n+1}/\delta_n)$ is close to 0 and α_n given in (21) is close to 0.5. Thus, in early iterations the algorithm is forced to explore the space of candidate solutions. As the iterations progress and the algorithm learns the behavior of the ion channel, the average error T_n is getting smaller, then $K(T_{n+1}/\delta_n)$ is close to 1 and $\alpha_n \to 0$. As a result as the algorithm becomes more confident about the ion channel behavior, it reduces the exploration probability to reduce the exploration cost. To proceed, we present a theorem that reveals the convergence of the algorithm presented in Step 3 above.

Theorem 3: Assume that the conditions of Theorem 1 are satisfied. Then

- the sequence $(\hat{\pi}_n, T_n, \alpha_n)'$ given in Step 3 above converges to $(\pi, 0, 0)'$ w.p. 1. Also
- Algorithm 2 is attracted to the equilibrium potential v^*, i.e., for sufficiently large n, the sequence $\{\theta(X_n)\}$ spends more time at v^* than any other states. (Equivalently, if $\theta(m^*) = v^*$, then $\hat{\pi}_n(m^*) > \hat{\pi}_n(j)$ for $j \in \{1, \ldots, M\} - \{m^*\}$.)

Remark: Let us comment on the rate of convergence of the algorithm under consideration. For Algorithm 2, the rate of convergence can be studied by use of diffusion approximation methods. The basic ideas and technical details are in [15, Ch. 10] and we only outline the main ideas below.

Define

$$\begin{aligned} U_n &= \sqrt{n}(\hat{\pi}_n - \pi), \widehat{T}_n = \sqrt{n} T_n, \\ U^0(t) &= U_n, \widehat{T}^0(t) = T_n, t \in [t_n, t_{n+1}), \\ U^n(t) &= U^0(t + t_n), \widehat{T}^n(t) = \widehat{T}^0(t + t_n). \end{aligned} \tag{22}$$

Then using the martingale averaging method, we can show that the following limit result holds.

Theorem 4: Under the conditions of Theorem 3, the scaled sequence of errors $(U^n(\cdot), \widehat{T}^n(\cdot))'$ converges weakly to a limit process $(U(\cdot), \widehat{T}(\cdot))'$ such that the limit process is a solution of the following system of stochastic differential equations:

$$\begin{aligned} dU &= -U dt + \Sigma^{1/2} dw, \\ d\widehat{T} &= -(U + \widehat{T}) dt + \Sigma^{1/2} dw \end{aligned} \tag{23}$$

where Σ is given by the limit

$$\sum_{j=n_0}^{n+n_0} \sum_{k=n_0}^{n+n_0} \left(e_{X_{j+1}} - E e_{X_{j+1}} \right) \left(e_{X_{j+1}} - E e_{X_{j+1}} \right)' \to \Sigma \text{ w.p.1.}$$

Remark: The convergence rates of $(U_n, \widehat{T}_n, \alpha_n)'$ may be obtained by the error bounds in the L_2 sense. In fact, it is easily seen that

$$nE|U_n|^2 = O(1), nE|\widehat{T}_n|^2 = O(1), \text{and } n^{2-4\gamma_1}E|\alpha_n|^2 = O(1)$$

provided we choose $\delta_n = n^{-\gamma_1}$ for some $\gamma_1 < 1$. The above result gives further insight on the asymptotic distribution $(U_n, \widehat{T}_n)$. In fact, we have $\sqrt{n}(U_n, \widehat{T}_n)$ converges in distribution to a normal random variable $N(0, S)$ where S is given by

$$S = \int_0^\infty \exp(At)\text{diag}(\Sigma, \Sigma)\exp(A't)dt$$

and $A = \begin{pmatrix} -I & 0 \\ -I & -I \end{pmatrix}$ with I being the $M \times M$ identity matix. Note that S is in fact, the stationary covariance of the diffusion process.

IV. Numerical Results—Bi-Ionic Potassium Ion Channel Equilibrium Potential Estimation

Using computer generated synthetic data for the potassium ion channel *I–V* curve, we illustrate the efficiency of the adaptive exploration algorithm compared to the nonadaptive exploration algorithm. the performance of the performance of Algorithms 1 and 2.

Sample paths of the potassium ion channel current $\{i_n(v)\}$ as a Markov chain with transition probability matrix A and open-state current level $I(v)$. Here

$$A = \begin{bmatrix} 0.97 & 0.03 & 0 \\ 0.3 & 0.6 & 0.1 \\ 0 & 0.1 & 0.9 \end{bmatrix} \tag{24}$$

and $I(v)$ was generated using the experimentally determined *I–V* curve of Fig. 2. The choice of A in (24) implies that the steady state probability vector of the Markov chain is $\pi = [0.834, 0.083, 0.083]'$. This is consistent with patch clamp experimental data which shows that typically the ion channel current spends about 80% of its time in the gap mode, and about 10% of its time in each of the closed and open states. Note that the *I–V* curve of the potassium ion channel in Fig. 2 it is a difficult case to handle, since it is nonohmic (nonlinear) and has asymmetric behavior for positive and negative currents.

The observed channel current at the electrode was simulated by adding white Gaussian noise with standard deviation $\sigma(v) = 0.3$ to the simulated ion channel current sequence $\{i_n(v)\}$, resulting in the HMM sequence $\{y_n(v)\}$ (see (1)). Given the simulated ion channel current for the potassium ion channel with *I–V* curve in Fig. 2, we used Algorithm 1 and Algorithm 2 to estimate the equilibrium potential v^*. Experiments were run over batch sizes $\Delta = 10\,000$.

Both algorithms were initialized with $X_0 = 1$, i.e., initial applied voltage $v = -150$ mV. In Step 2, the EM algorithm was run for 200 iterations on each Δ-length batch of HMM observations. As described at the end of Section III-B, this takes only about 0.001 secs on a 2 GFlop Pentium 4. The resulting

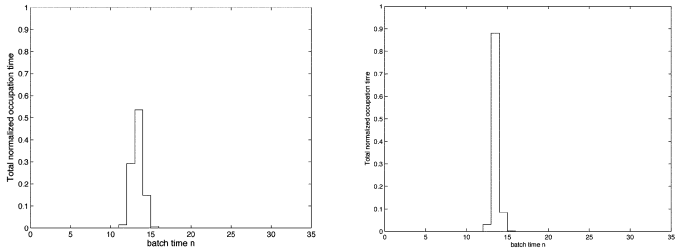

(a) Algorithm 1 (b) Algorithm 2

Fig. 6. Total normalized time Algorithm 1 and Algorithm 2 spend at various candidate potentials. The efficiency of the algorithms are the normalized amount of time spent at the equilibrium potential v^*. Efficiency of Algorithm 2 is 89% while the efficiency of Algorithm 1 is 54%. (a) Algorithm 1. (b) Algorithm 2.

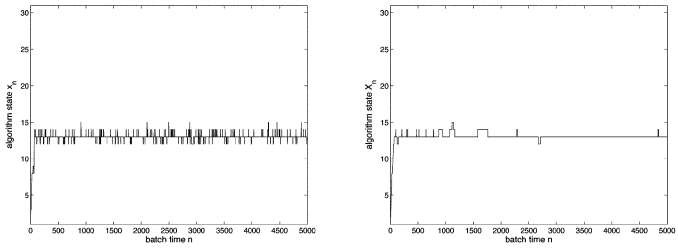

(a) Algorithm 1 (b) Algorithm 2

Fig. 7. State evolution of adaptive exploration Algorithm 2 and fixed exploration Algorithm 1. For the adaptive exploration algorithm the exploration probability α_n evolves as shown in Fig. 8. For the fixed exploration algorithm, $\alpha = 0.5$. (a) Algorithm 1. (b) Algorithm 2.

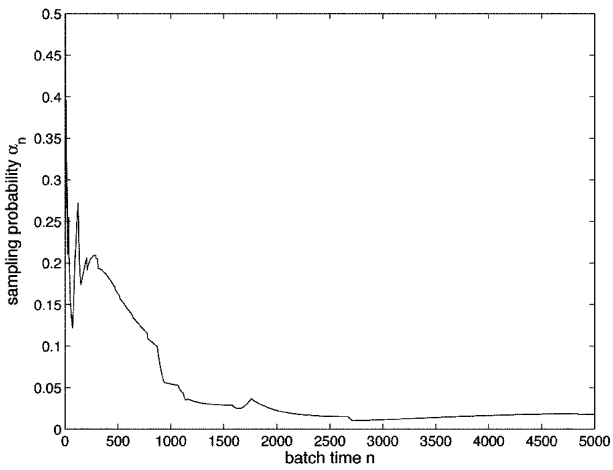

Fig. 8. Evolution of exploration probability α_n versus batch time n of Algorithm 2. In comparison with Algorithm 1, α is kept fixed at 0.5.

MLEs for the 4 batches, namely $\hat{I}_N^{(1)}(\tilde{v})$, $\hat{I}_N^{(2)}(\tilde{v})$, $\hat{I}_N^{(1)}(v)$ and $\hat{I}_N^{(2)}(v)$ were used to determine $\hat{C}_N(\tilde{v})$ and $\hat{C}_N(v)$.

The most important comparison between Algorithms 1 and the adaptive exploration Algorithm 2 is the total amount of experimental effort used to estimate the equilibrium potential. This total effort (normalized by the number of batch iterations) is depicted in Fig. 6. It is seen from Fig. 6 that the adaptive exploration algorithm results in an overall efficiency of 89% compared to the fixed exploration algorithm with efficiency of 54%. That is, Algorithm 2 conducted 35% fewer experiments. Over 5000 batch iterations the savings amount to 1750 patch

clamp experiment batches which is an enormous saving in experimental effort. Fig. 7 shows how the state X_n of the fixed exploration algorithm jumps around with batch iterations whereas the state of the adaptive exploration algorithm remains relatively constant after the first 1500 iterations. Again this illustrates the considerable experimental savings in the adaptive exploration algorithm. Finally, Fig. 8 shows how the exploration probability α_n in Algorithm 2 evolves with time. It can be seen that initially α_n is large so that the algorithm explores the set of candidate voltages aggressively. As the algorithm learns more about the equilibrium potential, it reduces the exploration probability α_n and α_n gets close to zero for large n.

V. Conclusions and Extensions

In this paper we presented two learning-based discrete stochastic approximation algorithms for learning the behavior of nerve cell membrane ion channels and controlling the patch clamp experiment. The main idea was to take into account the cost of exploration of the learning algorithm. We gave two performance criteria—the exploration reward (asymptotic covariance) and exploration cost—and showed how the exploration probability α can be chosen as a tradeoff between these criteria. In simulation studies we showed that the adaptive exploration algorithm achieves a significant efficiency improvement over earlier simulation based learning algorithms. Indeed for a 5000 batch length patch clamp experiment, the savings were shown to be 1750 patch clamp experiment batches. In future work we will explore the use of similar learning-based control algorithms for controlling Gramicidin biosensors [16].

Appendix

A. Proof of Theorem 1

Proof: By construction of Algorithm 1, the transition probability matrix $P(\alpha)$ for $\alpha \in (0, 0.5]$ results in X_n being an irreducible aperiodic (ergodic) Markov chain with a unique invariant distribution π. The invariant distribution π of the Markov chain X_n with transition probability matrix $P(\alpha)$ is straightforwardly computed as

$$\pi(m) = \pi(1) \prod_{l=2}^{m} \frac{\alpha p_{l-1}}{\alpha q_l} = \pi(1) \prod_{l=2}^{m} \frac{p_{l-1}}{q_l} \tag{25}$$

which is independent of the exploration probability α. In other words, the stationary distribution of the Markov chain X_n with transition probability matrix α is independent of the choice of α. Since X_n is an ergodic finite-state Markov chain, the empirical occupation probabilities $\hat{\pi}_n$ converge geometrically fast to π. From (25) we also have

$$\pi(m+1) = \pi(m) \frac{p_m}{q_{m+1}}.$$

From Assumption (O) it then follows that if $I^2(\theta(m+1)) > I^2(\theta(m))$, then $\pi(m+1) < \pi(m)$, if $I^2(\theta(m+1)) < I^2(\theta(m))$, then $\pi(m+1) > \pi(m)$. Hence, the algorithm converges to the local set of minimizers.

Finally, since $I^2(v)$ is a strictly convex function (as the *I–V* curve is strictly increasing), it has a unique minimum v^* which is the global minimum. Therefore, the algorithm converges to the global minimum which by definition is the equilibrium potential v^*.

To check that X_n is reversible it is sufficient to verify that $P(\alpha)$ satisfies $\pi(i)P_{ij}(\alpha) = \pi(j)P_{ji}(\alpha)$. This condition is straightforwardly verified. Hence, X_n is a reversible Markov chain. ■

B. Proof of Theorem 2

As shown in Theorem 1, $\hat{\pi}_n^* \to \pi$ irrespective of the exploration probability α. Thus, we can apply Theorem 7.2 [17, p. 299] which states that if $P(\alpha_1)$, $P(\alpha_2)$ are two transition matrices for reversible ergodic Markov chains on a finite-state space with the same stationary distribution π. If all off-diagonal terms in $P(\alpha_1)$ are greater than the corresponding terms in $P(\alpha_2)$, then $\Gamma(\alpha_1) \leq \Gamma(\alpha_2)$. The nonzero diagonal terms are αq_i and αp_i. So if $\alpha_1 > \alpha_2$, the off-diagonal terms in $P(\alpha_1)$ are greater than the corresponding terms in $P(\alpha_2)$. Thus, $\gamma(\alpha)$ is a decreasing function of α. Hence, $\Gamma(\alpha)$ is minimized for $\alpha = \bar{\alpha}$.

C. Proof of Theorem 3

The second assertion can be established as in the proof of Theorem 1. Thus, we shall concentrate on the proof of the first assertion. The proof is divided into two parts.

Part 1: We first consider the vector $(\hat{\pi}_n, T_n)$. To prove the convergence, we use a different approach than that of Theorem 1, namely, the ODE approach [15]. For future use, define

$$L_n = (\hat{\pi}_n, T_n)' \in \mathbb{R}^{2m}, \quad B = \begin{bmatrix} -1 & 0 \\ -1 & -1 \end{bmatrix}, \quad 1_2 = \begin{bmatrix} 1 \\ 1 \end{bmatrix}. \tag{26}$$

Step 1: We shall show that $\{L_n\}$ is bounded w.p. 1. From (21), the recursion for L_n can be written as

$$L_{n+1} = L_n + \mu_n B L_n + \mu_n 1_2 e_{X_{n+1}}. \tag{27}$$

That is, the algorithm that we are dealing with is a linear stochastic approximation algorithm. It is easily seen that

$$\begin{aligned} |L_{n+1}|^2 = {} & |L_n|^2 + 2\mu_n L_n' B L_n + \mu_n L_n' 1_2 e_{X_{n+1}} \\ & + \mu_n \left(1_2 e_{X_{n+1}}\right)' L_n + \mu_n^2 L_n' B' B L_n \\ & + \mu_n^2 L_n' B' 1_2 e_{X_{n+1}} + \mu_n^2 \left(1_2 e_{X_{n+1}}\right)' B L_n \\ & + \mu_n \left|1_2 e_{X_{n+1}}\right|^2. \end{aligned}$$

Note that although B is not symmetric, obviously $B = \widetilde{B} + \widehat{B}$, where $\widetilde{B}' = \widetilde{B}$ and $\widehat{B}' = -\widehat{B}$ and that for any quadradic form $y'By = y'\widetilde{B}y$. That is, the antisymmetric part $\widehat{B}$ does not contribute anything to the value of $y'By$. For our case,

$$\widetilde{B} = (B + B')/2 = \begin{bmatrix} -1 & -\frac{1}{2} \\ -\frac{1}{2} & -1 \end{bmatrix},$$

which has eigenvalues $-1/2$ and $-3/2$. Since $O(\mu_n^2)(1+|L_n|^2) = O(\mu_n)(1+|L_n|^2)$ for sufficiently small μ_n, we have

$$E|L_{n+1}|^2 \le (1-\mu_n)E|L_n|^2 + O(\mu_n)\left(1+E|L_n|^2\right) \le C_{n|0}E|\hat{\pi}_0|^2 + K\sum_{j=0}^{n}\mu_j C_{n|j}\left(1+E|L_j|^2\right).$$

Since $\mu_n = 1/(n+1)$, $\sum_{j=0}^{n}\mu_j C_{n|j} < \infty$, where

$$C_{n|j} = \begin{cases} \prod_{l=j}^{n}\left(1-2/(l+1)\right), & n > j, \\ 1, & n = j. \end{cases}$$

An application of the Grownwall's inequality implies that

$$E|L_{n+1}|^2 \le K\left(C_{n|0}E|L_0|^2 + K\sum_{j=0}^{n}\mu_j C_{n|j}\right)\exp\left(K\sum_{j=0}^{n}\mu_j C_{n|j}\right) < \infty.$$

Moreover, the bound holds uniformly. That is, $\sup_n E|L_{n+1}|^2 < \infty$. Consequently, by virtue of Theorem 6.7.2 [15, p. 192], $\{L_n\}$ is bounded w.p. 1.

Step 2: We work with a piecewise constant interpolation of L_n. Define

$$t_n = \sum_{j=0}^{n-1}\mu_j = \sum_{j=0}^{n-1}\frac{1}{j+1},$$
$$m(t) = \max\left\{m : \sum_{j=0}^{m-1}\frac{1}{j+1} \le t\right\},$$
$$L^0(t) = L_n, \text{ for } t \in [t_n, t_{n+1}),$$
$$L^n(t) = L^0(t+t_n).$$

Then $\{L^n(\cdot)\}$ is uniformly bounded.

Note that $\{e_{X_n}\}$ is a sequence of uniformly bounded random variables. This boundedness, the linearity in L_n, and the interpolation imply that $\{L^n(\cdot)\}$ is equicontinuous in the extended sense [15, p. 102]. [Note that since we do not use a truncation algorithm, we work with the interval $[0,\infty)$ than $(-\infty,\infty)$.] By virtue of [15, Theorem 4.2.2], we can extract a convergent subsequence and still denote it by $L^n(\cdot)$ for notational simplicity. Denote the limit by $L(\cdot)$. Then $L^n(\cdot) \to L(\cdot)$ w.p. 1 and the convergence is uniform on any bounded interval.

Step 3: Using the ODE method for stochastic approximation, we can show that the limit L satisfies the following system of ordinary differential equations

$$\dot{L}(t) = BL(t) + 1_2\pi. \tag{28}$$

Note that (28) can be written equivalently as

$$\dot{\pi}(t) = \pi - \pi(t)$$
$$\dot{T}(t) = \pi - \pi(t) - T(t). \tag{29}$$

Step 4: It is readily seen that $L_* = (\pi, 0)'$ is a unique asymptotic stable point of (28). As $t \to \infty$, $L(t) = (\pi(t), T(t))' \to (\pi, 0)'$. Consequently, the ODE method in [15] enables us to conclude that $L_n \to L_* = (\pi, 0)'$ w.p. 1.

It now remains to deal with the sequence $\{\alpha_n\}$. By virtue of the structure of the kernel $K(\cdot)$, $K(T_n/\delta_n) \to 1$ as $n \to \infty$. As a result, $\alpha_n \to 0$ w.p. 1 as $n \to \infty$.

REFERENCES

[1] E. Neher and B. Sakmann, "Single-channel currents recorded from membrane of denervated frog muscle fibers," *Nature*, vol. 260, pp. 799–802, 1976.

[2] O. Hamill, A. Marty, E. Neher, B. Sakmann, and F. Sigworth, "Improved patch-clamp techniques for high-resolution current recording from cells and cell-free membrane patches," *Pflüders Arch.-Eur. J. Physiol.*, vol. 391, pp. 85–100, 1981.

[3] F. Ashcroft, *Ion Channels and Disease*. New York: Academic, 2000.

[4] S. Chung, O. Andersen, and V. Krishnamurthy, *Handbook of Biological Membrane Ion Channels: Dynamics, Structure and Applications*. New York: Springer-Verlag, 2006.

[5] V. Krishnamurthy, S. Chung, and G. Dumont, Eds., *IEEE Trans. NanoBiosci. (Special Issue on Ion Channels—Bionanotubes)*, vol. 4, no. 1, pp. 1–132, Mar. 2005.

[6] V. Krishnamurthy and S. Chung, "Brownian dynamics simulation for modeling ion permeation across bio-nanotubes," *IEEE Trans. NanoBiosci.*, vol. 4, no. 1, pp. 102–111, Mar. 2005.

[7] S. Chung, V. Krishnamurthy, and J. Moore, "Adaptive processing techniques based on hidden Markov models for characterising very small channel currents buried in noise and deterministic interferences," *Proc. Phil. Trans. Roy. Soc. Lond. B*, vol. 334, pp. 357–384, 1991.

[8] L. Venkataramanan, R. Kuc, and F. Sigworth, "Identification of hidden Markov models for ion channel currents-part ii: bandlimited, sampled data," *IEEE Trans. Signal Process.*, vol. 48, no. 2, pp. 376–385, Feb. 2000.

[9] S. Andradottir, "Accelerating the convergence of random search methods for discrete stochastic optimization," *ACM Trans. Model. Comput. Simul.*, vol. 9, no. 4, pp. 349–380, Oct. 1999.

[10] V. Krishnamurthy and S. Chung, "Adaptive learning algorithms for nernst potential and current-voltage curves in nerve cell membrane ion channels," *IEEE Trans. NanoBiosci.*, vol. 2, no. 4, pp. 266–278, Dec. 2003.

[11] S. Andradottir, "A method for discrete stochastic optimization," *Manage. Sci.*, vol. 41, no. 12, pp. 1946–1961, 1995.

[12] B. Leroux, "Maximum-likelihood estimation for hidden Markov models," *Stoch. Process. Their Appl.*, vol. 40, pp. 127–143, 1992.

[13] B. Hille, *Ionic Channels of Excitable Membranes*, 3rd ed. Sunderland, MA: Sinauer Assoc., 2001.

[14] P. Bickel, Y. Ritov, and T. Rydén, "Asymptotic normality of the maximum-likelihood estimator for general hidden Markov models," *Ann. Stat.*, vol. 26, pp. 1614–1635, 1998.

[15] H. Kushner and G. Yin, *Stochastic Approximation Algorithms and Recursive Algorithms and Applications*, 2nd ed. New York: Springer-Verlag, 2003.

[16] B. Cornell, V. Braach-Maksvytis, L. King, P. Osman, B. Raguse, L. Wieczorek, and R. Pace, "A biosensor that uses ion-channel switches," *Nature*, vol. 387, pp. 580–583, Jun. 1997.

[17] P. Bremaud, *Markov Chains*. New York: Springer-Verlag, 1999.

Vikram Krishnamurthy (S'90-M'91-SM'99-F'05) was born in 1966. He received the B.S. degree in electrical engineering from the University of Auckland, New Zealand, in 1988 and the Ph.D. degree from the Australian National University, Canberra, in 1992.

He was a Chaired Professor at the Department of Electrical and Electronic Engineering, University of Melbourne, Australia, where he also served as Deputy Head of Department. Since 2002, he has been a Professor and Canada Research Chair at the Department of Electrical Engineering, University of British Columbia, Vancouver, BC, Canada. He has served as Associate Editor for *Systems and Control Letters* and *European Journal of Applied Signal Processing*. His research interests span several areas, including ion channels and nanobiology, stochastic scheduling and control, statistical signal processing, and wireless telecommunications.

Dr. Krishnamurthy has served as Associate Editor for the IEEE TRANSACTIONS ON SIGNAL PROCESSING, the IEEE TRANSACTIONS ON AEROSPACE AND ELECTRONIC SYSTEMS, and the IEEE TRANSACTIONS ON CIRCUITS AND SYSTEMS—II. He was guest editor of a special issue of IEEE TRANSACTIONS ON NANOBIOSCIENCE in March 2005 on bionanotubes.

G. George Yin (S'87-M87-SM'96-F'02) received the B.S. degree in mathematics from the University of Delaware, Newark, in 1983, and the M.S. degree in electrical engineering and the Ph.D. degree in applied mathematics from Brown University, Providence, RI, in 1987.

He joined the Department of Mathematics, Wayne State University, and became a professor in 1996. He is an Associate Editor of *SIAM Journal on Control and Optimization*, *Automatica*, and is (or was) on the editorial board of five other journals.

Dr. Yin served on the Mathematical Review Date Base Committee, IFAC Technical Committee on Modeling, Identification and Signal Processing, and various conference program committees; he was the editor of SIAM Activity Group on Control and Systems Theory Newsletters, the SIAM Representative to the 34th CDC, Co-Chair of 1996 AMS-SIAM Summer Seminar in Applied Mathematics, Co-Chair of 2003 AMS-IMS-SIAM Summer Research Conference: Mathematics of Finance, and Co-Organizer of the 2005 IMA Workshop on Wireless Communications. He was an Associate Editor of IEEE TRANSACTIONS ON AUTOMATIC CONTROL from 1994 to 1998