Improved Tracking of Time-Varying Encoding Properties of Visual Neurons by Extended Recursive Least-Squares

Nicholas A. Lesica, *Student Member, IEEE,* and Garrett B. Stanley, *Associate Member, IEEE*

Abstract—Traditional approaches to characterizing the transformation from stimulus to response in sensory systems assume both stationarity of the stimulus and time-invariance of the stimulus/response mapping. However, recent studies of sensory function under natural stimulus conditions have demonstrated important features of neural encoding that are in violation of these assumptions. Many sensory neurons respond to changes in the statistical distribution of the stimulus that are characteristic of the natural environment with corresponding changes in their encoding properties. In this paper, an extended recursive least-squares (ERLS) approach to adaptive estimation from stimulus/response observations is detailed. The ERLS approach improves the tracking ability of standard RLS approaches to adaptive estimation by removing a number of assumptions about the underlying system and the stimulus environment. The ERLS framework also incorporates an adaptive learning rate, so that prior knowledge of the relationship between the stimulus and the adaptive nature of the system under investigation can be used to improve tracking performance. Simulated and experimental neural responses are used to demonstrate the ability of the ERLS approach to track adaptation of encoding properties during a single stimulus/response trial. The ERLS framework lends tremendous flexibility to experimental design, facilitating the investigation of sensory function under naturalistic stimulus conditions.

Index Terms—Adaptive estimation, least-squares, receptive field, recursive.

I. INTRODUCTION

ONE MAJOR goal of sensory neurophysiology is to understand how information about the outside world is encoded in the electrical activity of neurons within various pathways of the brain. This is a system identification problem: encoding mechanisms (system) are characterized by observing neural responses (output) to sensory stimuli (input). The traditional approach to this problem is to record the firing rate modulations of neurons in response to simple stimuli, such as white-noise, and estimate a linear filter, or receptive field (RF), using least-squares estimation [1]. This approach has been successful in characterizing the basic function of many sensory areas. RFs have been used in functional models of the visual, auditory, and somatosensory pathways [2]–[4]. However, several assumptions

are generally involved in the traditional approach, namely that the stimulus is stationary (drawn from a fixed statistical distribution), and that the encoding properties of the neuron are time-invariant, both of which are only valid under the most artificial of experimental conditions.

In a natural setting, the statistical distribution of the stimulus is constantly changing. For example, the mean and contrast of the intensity of light incident upon the retina can vary over many orders of magnitude throughout the day due to changes in illumination and the constant motion of objects within the visual field. It is also known that sensory neurons adapt their encoding properties (gain, spatial, and temporal tuning, etc.) in response to such changes (for reviews, see [5] and [6]). These adaptive mechanisms are thought to serve important functional roles, such as maintaining contrast sensitivity over a wide range of mean values [5] and maximizing the flow of information to downstream neurons [7], [8]. Encoding models based on RF's estimated from responses to stationary stimuli neglect these adaptive mechanisms, and, thus, are insufficient to characterize function in the natural environment.

One approach that has been employed to characterize adaptive sensory encoding mechanisms from stimulus/response observations is recursive least-squares (RLS) estimation [9], [10]. In RLS, the estimated parameters of the encoding model are updated with each new observation, while the influence of past observations is downweighted, allowing the estimate to reflect the current state of the underlying system. However, in the standard RLS approach, a number of assumptions are made concerning the underlying system, and these assumptions limit the ability of RLS to track fast variations in model parameters. Additionally, the fixed learning rate in the standard RLS approach requires a trade-off between the ability to track parameter changes and the steady-state error level in the estimate.

In this paper, we develop an extended recursive least-squares (ERLS) approach to adaptive neural system identification to address these problems. The ERLS approach incorporates a stochastic model of parameter evolution (based on that which underlies the Kalman filter) and an adaptive learning rate that is modulated in response to changes in features of the visual stimulus. We use ERLS to analyze simulated and experimental responses of visual neurons to stimuli of varying complexity and demonstrate its ability to closely track adaptive changes in RF structure induced by nonstationary stimulus conditions during a single trial.
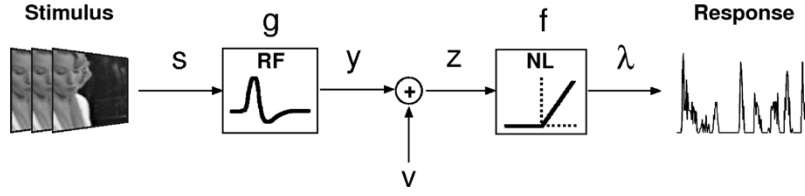
Fig. 1. Model of visual encoding. The spatiotemporal visual stimulus s is passed through a time-varying linear filter g (the spatiotemporal RF) to yield the intermediate signal y. This signal is then combined with additive, independent noise v to yield the generating function z, and passed through a rectifying static nonlinearity f to produce the nonnegative firing rate λ.

II. MODEL OF SENSORY ENCODING

The development of the ERLS algorithm and the examples in this paper are based on the properties of neurons in the early visual pathway, but the results are easily generalized to other sensory systems. The mapping from stimulus to response in a visual neuron can be represented by a cascade of encoding elements consisting of a linear filter and a rectifying static nonlinearity. A schematic diagram of a cascade encoding model is shown in Fig. 1. Each component of the cascade model is described later in detail.

A. Spatiotemporal Receptive Field

The input is the spatiotemporal signal $s[p, n]$. For computer driven visual stimuli discretized in space-time, p represents the grid index of a pixel on the screen and n is the time sample. Note that pixel refers not to the atomic display units of the monitor, but, for instance, to squares in a white-noise checkerboard. No assumptions are made about the statistics of the stimulus, as natural signals are often nonstationary and correlated. To produce the intermediate signal y, which reflects the stimulus-related modulations in the membrane potential of the neuron, the stimulus is passed through the time-varying linear filter $g_n[p, m]$ (convolution in time, integration in space) representing P (total pixels in stimulus) separate temporal filters each with M parameters. This filter is known as the spatiotemporal RF and captures the spatial and temporal integration of the stimulus within the visual pathway. If s and g_n are organized appropriately, then this discrete time operation can be written as a dot product $y[n] = s_n^T g_n$, where s_n and g_n are the column vectors defined as

$$s_n \triangleq [s[P, n - M + 1], s[P - 1, n - M + 1], \ldots,$$
$$s[1, n - M + 1], s[P, n - M + 2], \ldots, s[1, n]]^T \quad (1)$$
$$g_n \triangleq [g_n[P, M], g_n[P - 1, M], \ldots,$$
$$g_n[1, M], g_n[P, M - 1], \ldots, g_n[1, 1]]^T \quad (2)$$

and T denotes matrix transpose.

B. Additive Noise

Before being passed through the static nonlinearity, y is combined with additive noise v to yield z, which is known as the generating function. The noise v represents spontaneous fluctuations in the membrane potential of the neuron, which are reflected in the variability of firing rate responses to repeated presentations of the same stimulus. These fluctuations have been shown to be uncorrelated over time and independent of the baseline membrane potential of the neuron, with a distribution that

is approximately Gaussian [11]. Thus, the noise v is assumed to be independent of the stimulus and Gaussian with zero mean $\mathcal{N}(0, \sigma_v^2[n])$.

C. Static Nonlinearity

The generating function z is passed through a static nonlinearity $f(\cdot)$ to yield the nonnegative firing rate λ. This static nonlinearity captures the transformation from the membrane potential of the neuron to its observed firing rate. A common model for the nonlinearity present in many experimentally observed visual neurons is linear half-wave rectification [12]

$$f(z) = \begin{cases} z, & z \geq 0 \\ 0, & z < 0 \end{cases}. \quad (3)$$

The linear rectifying static nonlinearity implies that the neuron is silent when the membrane potential is below a threshold and that modulations in the membrane potential above that threshold are reflected as proportional modulations in firing rate. Note, however, that the ERLS algorithm is developed for a general static nonlinearity $f(\cdot)$ and the specific form of the function should be chosen to match the properties of the neuron under investigation [13].

III. RECURSIVE ESTIMATION

The basic premise of recursive system identification is that the estimate of the underlying parameters at time $n + 1$ is computed by combining the previous estimate from time n with an update based on the new observation and prior knowledge of the evolutionary dynamics of the parameters. RLS estimation is one particular form of this approach that has been used successfully to characterize the encoding properties of visual neurons [9], [10].

A. RLS Estimation

We have previously detailed an RLS technique for identifying the encoding properties of visual neurons based on a model structure such as the cascade shown in Fig. 1 [9]. The RLS algorithm is based on the minimization of the prediction error, which is defined as the difference between the observed firing rate in an interval, and the expected firing rate in the interval given the current stimulus and the estimated model parameters. In the model shown in Fig. 1, the only parameters to be estimated are those of the RF g_n. Note, however, that the estimation techniques described in the following are easily adapted to include simultaneous estimation of the RF and other parameters [14].

Using the variables defined in the encoding model shown in Fig. 1, the RLS algorithm for the estimation of the RF parameters can be written as follows:

$$e[n] = \lambda[n] - f\left(s_n^T \hat{g}_{n|n-1}\right) \quad \text{Prediction Error}$$

$$G_n = \frac{\gamma^{-1} K_{n|n-1} s_n}{\gamma^{-1} s_n^T K_{n|n-1} s_n + 1} \quad \text{Update Gain}$$

$$\hat{g}_{n+1|n} = \hat{g}_{n|n-1} + G_n e[n] \quad \text{Update Parameter Estimates}$$

$$K_{n+1|n} = \gamma^{-1} K_{n|n-1} - \quad \text{Update Inverse of}$$
$$\gamma^{-1} G_n s_n^T K_{n|n-1} \quad \text{Stimulus Autocovariance}$$

where $0 \ll \gamma \leq 1$ serves to downweight past information, and is, therefore, often called the forgetting factor. At each time step, the gain G is calculated based on the estimate of the inverse of the stimulus autocovariance matrix K and combined with the prediction error e to update the RF parameter estimate $\hat{g}$. Note, however, that the recursive framework avoids the explicit inversion of the stimulus autocovariance matrix. The subscript $n|n-1$ denotes an estimate at time n given all observations up to and including time $n-1$.

The prediction error e is the difference between the observed and predicted firing rates. Given the stimulus and the current estimate of the RF, the expected firing rate is

$$E\left\{\lambda[n]|s_n, \hat{g}_{n|n-1}\right\} = \int_\lambda \lambda[n] p\left(\lambda[n]|s_n, \hat{g}_{n|n-1}\right) d\lambda[n]$$
$$= \int_v f\left(s_n^T \hat{g}_{n|n-1} + v[n]\right) p\left(v[n]\right) dv[n]$$

where $p(\lambda[n]|s_n, \hat{g}_{n|n-1})$ is the probability density function of the predicted response conditioned on the current stimulus and estimated model parameters. For small v relative to $s_n^T \hat{g}_{n|n-1}$, the expectation can be approximated as $E\{\lambda[n]|s_n, \hat{g}_{n|n-1}\} \approx f(s_n^T \hat{g}_{n|n-1})$ through a series expansion about $s_n^T \hat{g}_{n|n-1}$. This approximation is valid when the signal-to-noise ratio (SNR) is large (>2), as is typically the case in visual neurons under dynamic stimulation. In the event that this approximation is not valid, the integral expression for the expected firing rate can be evaluated at each time step.

The dynamics of the system shown in Fig. 1 can be represented by the following state-space model:

$$g_{n+1} = F_n g_n + q_n \tag{4}$$
$$\lambda[n] = f\left(s_n^T g_n + v[n]\right) \tag{5}$$

where F_n (known as the state evolution matrix) and q_n (known as the state evolution noise) specify the deterministic and stochastic components of the evolutionary dynamics of the RF parameters, respectively. Investigation of the model underlying RLS reveals that the technique is designed to estimate time-invariant RF parameters ($F_n = \mathbf{I}$ and $q_n = 0$) based on the assumption that the variance of the noise in the observed response decreases exponentially over time ($\sigma_v^2[n] \propto \gamma^n$) [15]. This assumption causes the current observations to be weighted more heavily in the computation of the parameter estimates than those in the past, and provides RLS with the ability to track a slowly varying RF. However, as the parameters change more quickly, the ability of RLS to track them decreases. The tracking behavior of the RLS algorithm can be greatly improved by assuming a model in which the RF parameters are time-varying and their evolution is treated as a more general stochastic process.

B. ERLS Estimation

The optimal algorithm for tracking the time-varying parameters of a state-space model, in terms of minimizing the mean-squared error (MSE), is the Kalman filter [16]. However, implementation of the Kalman filter requires exact knowledge of the quantities F_n, $\Sigma_q[n]$ (the covariance of the state evolution noise), and $\sigma_v^2[n]$, which are generally unknown during the experimental investigation of neural systems. Fortunately, some of the tracking ability of the Kalman filter can be transferred to the RLS framework by replacing the deterministic model that underlies RLS estimation with a stochastic one that approximates that which underlies the Kalman filter. The result, known as extended recursive least-squares (ERLS), was developed in [15] based on the correspondence between RLS and the Kalman filter presented in [17]. Here, we develop a particular form of ERLS that is designed to track adaptation of visual encoding properties in response to changes in the statistical properties of the stimulus.

The model underlying the Kalman filter assumes that the RF parameters evolve according to the general model $g_{n+1} = F_n g_n + q_n$, where q_n is a vector of nonstationary Gaussian white-noise $\mathcal{N}(0, \Sigma_q[n])$. To simplify the incorporation of this model into the RLS framework, assume that the parameter evolution is completely stochastic ($F_n = \mathbf{I}$) and that the parameters evolve independently and at equal rates ($\Sigma_q[n] = \sigma_q^2[n]\mathbf{I}$). In this stochastic model, the evolution of the parameter estimates is constrained only by the variance $\sigma_q^2[n]$, and this parameter can be used to control the tracking behavior of the algorithm based on knowledge of the underlying system. When the prediction error is likely to be the result of changing encoding properties (during adaptation to changes in the statistics of the stimulus), a large value of $\sigma_q^2[n]$ is used to allow the estimate to track these changes. Conversely, if the parameters are not likely to be changing, a small value of $\sigma_q^2[n]$ is used to avoid tracking the noise in the observed response. Thus, $\sigma_q^2[n]$ functions as an adaptive learning rate.

The value of $\sigma_q^2[n]$ is adjusted based on knowledge of how features of the stimulus affect the parameters. For example, adaptation in the visual system generally occurs in the interval directly following a change in a feature of the stimulus (mean, variance, etc.). With this knowledge, $\sigma_q^2[n]$ is increased following a stimulus transition, allowing the parameter estimate to change quickly. Similarly, if the statistics of the stimulus have been stationary for some time and the underlying parameters are not likely to be adapting, $\sigma_q^2[n]$ is decreased. The dynamics of the adaptive learning rate should be based on knowledge of the adaptive properties of the system under investigation, reflecting the expected rate of change (per estimation time step) of the parameters under the given stimulus conditions. For a situation where the relevant stimulus feature or the adaptive properties are not known *a priori*, $\sigma_q^2[n]$ should be set to a

relatively small constant value throughout the trial. This gives the estimate some degree of adaptability (although fast changes in parameter values will likely be missed), while keeping the steady-state noise level in a reasonable range. The initial estimate provides some information about the adaptive behavior of the system, and the estimation can be performed again with a more appropriate choice of $\sigma_q^2[n]$. This process can be iterated and, as more is learned about the adaptive dynamics of the underlying system, the value of $\sigma_q^2[n]$ can be updated. For further discussion of the choice of $\sigma_q^2[n]$, see Section V.

The ERLS algorithm for the model in Fig. 1 is given by

$$e[n] = \lambda[n] - f\left(s_n^T \hat{g}_{n|n-1}\right) \quad \text{Prediction Error}$$

$$G_n = \frac{K_{n|n-1}s_n}{s_n^T K_{n|n-1}s_n + 1} \quad \text{Update Gain}$$

$$\hat{g}_{n+1|n} = \hat{g}_{n|n-1} + G_n e[n] \quad \text{Update Parameter}$$
$$\text{Estimates}$$

$$K_{n+1|n} = K_{n|n-1} - \quad \text{Update Inverse of}$$
$$G_n s_n^T K_{n|n-1} + \sigma_q^2[n]\mathbf{I} \quad \text{Stimulus Autocovariance.}$$

Again, the estimate is generated by solving the above equations sequentially at each time step. If possible, the algorithm should be initialized when the underlying system is at steady state and the parameter values are known. In the absence of any prior knowledge of the parameter values, the initial conditions $\hat{g}_{0|-1} = 0$ and $K_{0|-1} = \delta \times \mathbf{I}$ should be used. The regularization parameter δ effects the convergence properties and steady-state error of the ERLS estimate by placing a smoothness constraint on the parameter estimates, removing some of the error introduced by highly correlated natural stimuli [18]. For estimation from responses to uncorrelated white-noise stimuli in the examples below, δ was set to 10^{-4}. For estimation from responses to correlated naturalistic stimuli, δ was set to 10^{-2}. Note that the ERLS algorithm can easily be extended to simultaneously estimate the encoding properties of an ensemble of neurons simply by augmenting the stimulus, parameter, and response vectors accordingly. For example, an adaptive estimation technique similar to the ERLS approach presented here has been used to estimate the RF parameters of an ensemble of motor neurons from responses to changes in hand velocity [19].

IV. SIMULATIONS

In the following simulations, examples of adaptive encoding in retinal ganglion cells are used to demonstrate the ability of ERLS to track changes in RF parameters during nonstationary stimulation. Ganglion cells are the output neurons of the retina and provide the only pathway for the transmission of visual information from the retina to the brain. Responses to a variety of nonstationary stimuli were simulated using the cascade encoding model shown in Fig. 1, which has been shown to provide accurate predictions of ganglion cell responses [20].

A. Comparison of RLS and ERLS

A biphasic temporal RF typical of retinal ganglion cells, with a time course of 300 ms, was used in the cascade encoding model shown in Fig. 1 to simulate responses to a contrast switching, spatially uniform Gaussian white-noise stimulus. A
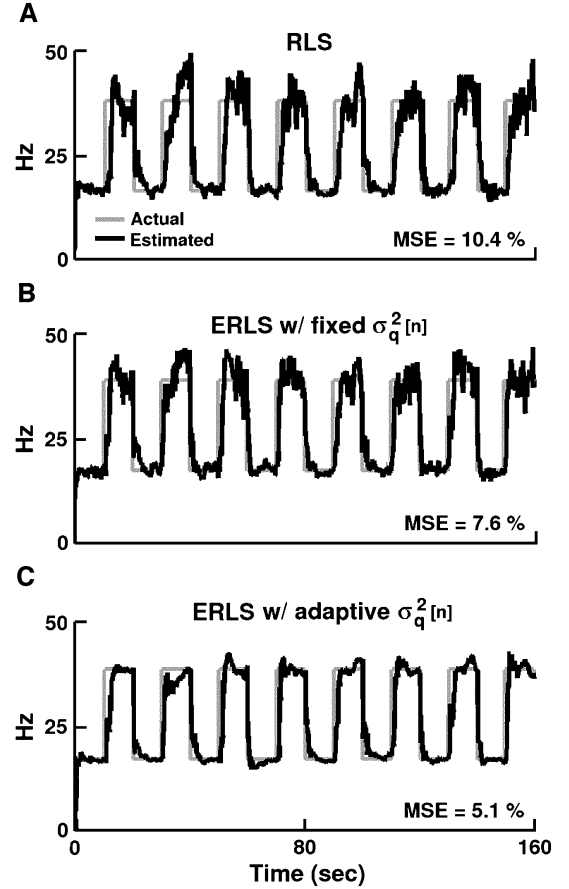

Fig. 2. Comparison of RLS and ERLS. (A) The gain of the actual RF (gray) along with the gain of the RLS estimate (black) for a 160-s segment of the contrast-switching white-noise stimulus. The value of the forgetting factor γ used to compute the RLS estimate was 0.96. The MSE in the RF estimate over the entire trial was 10.4% of the variance of the actual RF. (B) The gain of the actual RF along with that of the ERLS estimate, computed with $\sigma_q^2[n] = 10^{-5}$ for the entire trial. The MSE in the RF estimate over the entire trial was 7.6%. (C) The gain of the actual RF along with that of the ERLS estimate, computed with $\sigma_q^2[n] = 10^{-4}$ for the 1 s following each contrast transition, and 10^{-6} at all other times. The MSE in the RF estimate over the entire trial was 5.1%.

new luminance value for the stimulus was chosen every 30 ms and the root mean square (RMS, ratio of standard deviation to mean) contrast was switched from 0.05 to 0.30 every 30 s. Contrast gain control has been shown to modulate the RF gain of visual neurons in response to changes in stimulus contrast, with a time course that is approximately equal to the integration time of the neuron [21], [22]. To simulate the adaptive changes that have been observed experimentally, the gain of the RF (magnitude of peak value) was increased by a factor of 2 following a decrease in contrast, and decreased by a factor of 2 following an increase in contrast. The variance of the noise v was adjusted to produce responses with an SNR of 5. This value is consistent with those measured in the experimental responses of retinal ganglion cells [14].

ERLS and standard RLS were used to track changes in the RF parameters from the simulated responses. The results are shown in Fig. 2. Fig. 2(A) shows the gain of the actual RF (gray), along with the gain of the RLS RF estimate (black). The RLS estimate was generated with forgetting factor $\gamma = 0.96$ (which corresponds to a memory time constant of approximately 7 s).

This value was optimal in the sense that it yielded the lowest MSE in the RF estimate (10.4% of the variance of the actual RF) over the entire trial for all $0 \ll \gamma \leq 1$.

Fig. 2(B) shows the gain of the actual RF, along with the gain of the ERLS RF estimate, computed with a fixed learning rate $\sigma_q^2[n] = 10^{-5}$. This value of $\sigma_q^2[n]$ was also chosen to minimize the MSE in the RF estimate over the entire trial. The MSE in the ERLS estimate with fixed learning rate (7.6%) is lower than that of the optimal RLS estimate, illustrating the enhanced tracking ability that results from incorporating the stochastic model of parameter evolution.

The tracking performance of the ERLS estimate can be further improved by using an adaptive learning rate to exploit the relationship between the stimulus and the adaptive nature of the system. Since contrast gain control only modulates the gain of the RF in the short interval following a change in contrast, the learning rate $\sigma_q^2[n]$ is set to a large value in those intervals to allow the estimate to adapt quickly, and to a small value otherwise, so that noise in the observed response is not attributed to changes in the encoding properties of the neuron. Accordingly, $\sigma_q^2[n]$ was set to 10^{-4} during the transient intervals (1 s following each contrast transition), and 10^{-6} during steady-state intervals (all other times). The gain of the resulting RF estimate is shown in Fig. 2(C). The adaptive learning rate allows the estimate to closely track the fast changes in gain, while maintaining a low steady-state error between transitions. The MSE in the ERLS estimate with adaptive learning rate (5.1%) is half of that in the standard RLS estimate. The values of $\sigma_q^2[n]$ used to generate the ERLS estimate with an adaptive learning rate were chosen based on the adaptive dynamics of the simulated neuron, but were not optimized. Similar results were obtained with a range of values for $\sigma_q^2[n]$ during the transient and steady-state intervals (not shown), indicating the robust improvement in tracking provided by the adaptive learning rate.

B. Tracking RF Modulation During Natural Stimulation

In a natural setting, tracking RF changes is complicated by the lack of clear transitions in the relevant stimulus features (as opposed to the contrast-switching example above). The following contrast gain control simulation demonstrates the ability of ERLS to track RF modulation from responses to a stimulus with continuously varying statistics. The stimulus was the temporal intensity of one pixel of a grayscale natural scene movie recorded in the forest with a home video camera, updated every 30 ms, as shown in Fig. 3(A). For more details regarding the natural scene movie, see [23]. For this example, a stimulus was chosen in which the mean intensity was relatively constant over time, while the contrast was constantly changing. The response of a retinal ganglion cell was simulated as above. The gain of the temporal RF was varied inversely with the contrast of the stimulus, and, thus, varied continuously throughout the trial. At each time step, the contrast was defined as the RMS contrast of the previous 300-ms segment of the stimulus, in accordance with the time course of contrast gain control [Fig. 3(B)], and the gain of the RF was set to twice the inverse of the contrast. Since the contrast of the stimulus was constantly varying and transitions were not well defined, the value of the adaptive learning
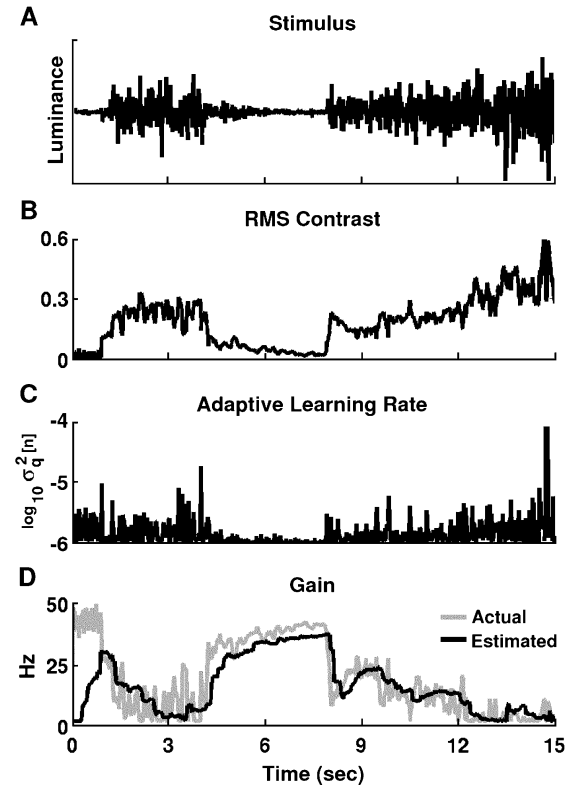

Fig. 3. Tracking RF modulation during natural stimulation. (A) The stimulus was spatially uniform and the luminance was modulated according to the intensity of a typical pixel in a natural scene movie, updated every 30 ms. (B) The RMS contrast of the natural stimulus throughout the trial. (C) The value of the learning rate $\sigma_q^2[n]$ throughout the trial. (D) The gain of actual RF during the simulation is shown (gray), along with the gain of the ERLS RF estimate (black).

rate $\sigma_q^2[n]$ was proportional to the derivative of the stimulus contrast, as shown in Fig. 3(C). At each time step, $\sigma_q^2[n]$ was defined as 10^{-4} times the absolute value of the first-order difference in the contrast of the stimulus. Fig. 3(D) shows the results of the estimation. The gain of the ERLS RF estimate (black) closely tracks that of the actual RF (gray). Aside from the error associated with the initial conditions and some of the very fast transients, the ERLS RF estimate captures most of the gain changes in the actual RF.

C. Tracking Contrast Gain Control in the Retina

The tracking ability of the ERLS technique can be demonstrated under experimental conditions using the response of a retinal ganglion cell to a contrast switching stimulus. Spatially uniform Gaussian white-noise was projected onto an isolated salamander retina and ganglion cell action potentials were recorded extracellularly. These experiments were performed in the laboratory of Markus Meister at Harvard University and details of the preparation are given in [20]. A new luminance value for the stimulus was chosen every 30 ms and the RMS contrast was switched from 0.05 to 0.3 every 30 s [see Fig. 4(A) and (B)]. As described in example A, contrast gain control is known to modulate gain in the short interval following a change in contrast. Thus, the learning rate $\sigma_q^2[n]$ was set to 10^{-4} during the transient intervals (1 s following each contrast transition), and 10^{-6} during steady-state intervals (all other times), as
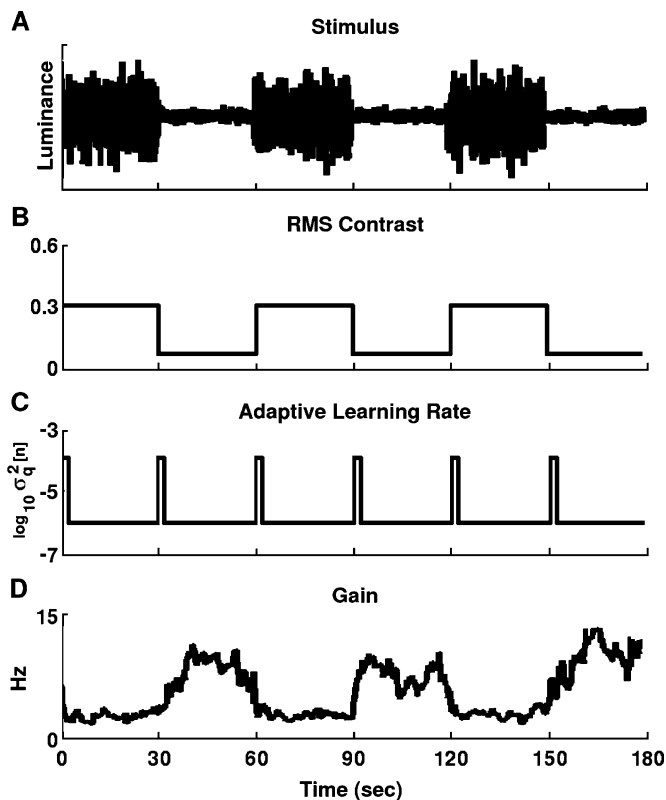
Fig. 4. Tracking gain changes in the retina. Experimentally observed responses of a salamander retinal ganglion cell to a contrast switching stimulus were used for adaptive estimation of encoding properties. (A) Three minutes of the spatially uniform Gaussian white-noise stimulus. A new luminance value for the stimulus was chosen every 30 ms. (B) The contrast of the stimulus was switched between 0.05 and 0.3 every 30 s. (C) The adaptive learning rate of the ERLS algorithm, $\sigma_q^2[n]$, was set to 10^{-4} during the 1 s following each contrast transition, and 10^{-6} during all other times. (D) The gain of the ERLS RF estimate.

shown in Fig. 4(C). The ERLS estimate tracks the changes in gain following stimulus transitions during a single trial, as shown in Fig. 4(D). Note that in this example, an additional offset parameter was estimated simultaneously with the RF to account for contrast induced changes in baseline membrane potential (not shown, for details see [14]).

V. CONCLUSION

In this paper, an ERLS approach to the identification of adaptive neural encoding properties was developed for use under nonstationary stimulus conditions. ERLS removes the simplifying assumptions of stimulus stationarity and time-invariance of the stimulus/response mapping that underly traditional RF estimation approaches. ERLS is an extension of standard RLS estimation with improved tracking abilities due to its incorporation of a stochastic model of parameter evolution and adaptive learning rate. Simulated and experimental examples demonstrated the superiority of ERLS to standard RLS in tracking changes in encoding properties during contrast-switching stimulation and the ability of ERLS to track adaptation under highly nonstationary stimulus conditions. Other studies have shown that ERLS outperforms both standard RLS and least-mean-square (LMS) estimation in a variety of nonstationary system identification problems, such as tracking a chirped sinusoid in noise [15].

The ERLS algorithm that we have presented requires specification of the learning rate parameter $\sigma_q^2[n]$. In the form of the algorithm that we presented, the value of this parameter is varied externally to match the adaptation dynamics of the system under investigation. The ERLS technique could be made more powerful by specifying a relationship between stimulus features and adaptation in the underlying system, such that the value of $\sigma_q^2[n]$ could be controlled by the algorithm internally based on changes in the relevant stimulus features. In the text, we described an iterative procedure for optimizing the learning rate when little is known about the adaptive dynamics of the underlying system. The internal control of the learning rate could be taken one step further by automating this iterative process such that the relationship between changes in the relevant stimulus features and the learning rate is also optimized during this process. In this form, the ERLS algorithm would be extremely powerful, as it would require no *a priori* specification of adaptation dynamics and would produce not only time-varying estimates of the model parameters, but also a description of the relationship between changes in stimulus features and adaptation of these parameters.

While single trial changes in the ERLS parameter estimates reflect changes in the underlying system, they are also influenced by a number of other factors, including the SNR in the response measurements, the correlation structure of the stimulus, and the learning rate $\sigma_q^2[n]$. As a result, it is desirable to develop some measure of confidence in the parameter estimates to determine what fraction of the changes should be attributed to adaptive function. If the SNR of the measurements and the degree of the nonstationarity of the system can be specified a priori, then confidence bounds can be determined analytically, as in the Kalman filter [16]. A more practical approach is to record the neural responses to repeated presentations of the same stimulus and estimate the parameters with ERLS independently for each trial. The mean of these estimates will reflect actual changes in the underlying system and the variance of these estimates can be used as a measure of confidence in other single trial estimates under similar experimental conditions.

REFERENCES

[1] P. Z. Marmarelis and V. Z. Marmarelis, *Analysis of Physiological Systems*. New York: Plenum, 1978.
[2] J. Keat, P. Reinagel, R. C. Reid, and M. Meister, "Predicting every spike: A model for the responses of visual neurons," *Neuron*, vol. 30, pp. 817–830, 2001.
[3] F. E. Theunissen, K. Sen, and A. J. Doupe, "Spectral-temporal receptive fields of nonlinear auditory neurons obtained using natural sounds," *J. Neurosci.*, vol. 20, pp. 2315–2331, 2000.
[4] J. J. DiCarlo, K. O. Johnson, and S. S. Hsiao, "Structure of receptive fields in area 3b of primary somatosensory cortex in the alert monkey," *J. Neurosci.*, vol. 18, pp. 2626–2645, 1998.
[5] R. Shapley and C. Enroth-Cugell, "Visual adaptation and retinal gain controls," *Prog. Ret. Res.*, vol. 3, pp. 263–346, 1984.
[6] M. Meister and M. J. Berry, "The neural code of the retina," *Neuron*, vol. 22, pp. 435–450, 1999.

[7] N. Brenner, W. Bialek, and R. de Ruyter van Steveninck, "Adaptive rescaling maximizes information transmission," *Neuron*, vol. 26, pp. 695–702, 2000.

[8] A. L. Fairhall, G. D. Lewen, W. Bialek, and R. R. de Ruyter van Stevenick, "Efficiency and ambiguity in an adaptive neural code," *Nature*, vol. 412, pp. 787–790, 2001.

[9] G. B. Stanley, "Adaptive spatiotemporal receptive field estimation in the visual pathway," *Neural Computation*, vol. 14, pp. 2925–2946, 2002.

[10] D. L. Ringach, M. J. Hawken, and R. Shapley, "Receptive field structure of neurons in monkey visual cortex revealed by stimulation with natural image sequences," *J. Vision*, vol. 2, pp. 12–24, 2002.

[11] I. Lampl, I. Reichova, and D. Ferster, "Synchronous membrane potential fluctuations in neurons of the cat visual cortex," *Neuron*, vol. 22, pp. 361–374, 1999.

[12] P. Dayan and L. F. Abbott, *Theoretical Neuroscience*. Cambridge, MA: MIT Press, 2001.

[13] E. J. Chichilnisky, "A simple white noise analysis of neuronal light responses," *Network*, vol. 12, pp. 199–213, 2001.

[14] N. A. Lesica and G. B. Stanley, "Adaptive interactions in the Visual Pathway," in *Soc. Neurosci. Abstr.*, 2003, Prog. 485.10.

[15] S. Haykin, A. H. Seyed, J. R. Zeidler, P. Yee, and P. C. Wei, "Adaptive tracking of linear time-variant systems by extended RLS algorithms," *IEEE Trans. Signal Process.*, vol. 45, no. 5, pp. 1118–1128, May 1997.

[16] R. E. Kalman, "A new approach to linear filtering and prediction problems," *Trans. ASME—J. Basic Eng.*, vol. 82, pp. 35–45, 1960.

[17] A. H. Sayed and T. Kailath, "A state-space approach to adaptive RLS filtering," *IEEE Signal Process. Mag.*, vol. 11, no. 3, pp. 18–60, Jul. 1994.

[18] B. Willmore and D. Smyth, "Methods for first-order kernel estimation: Simple-cell receptive fields from responses to natural scenes," *Network: Comput. Neural Syst.*, vol. 14, pp. 553–577, 2003.

[19] U. T. Eden, L. M. Frank, R. Barbieri, V. Solo, and E. N. Brown, "Dynamic analysis of neural encoding by point process adaptive filtering," *Neural Comp.*, vol. 16, pp. 971–998, 2004.

[20] S. A. Baccus and M. Meister, "Fast and slow contrast adaptation in retinal circuitry," *Neuron*, vol. 36, pp. 909–919, 2002.

[21] R. M. Shapley and J. D. Victor, "The effect of contrast on the transfer properties of cat retinal ganglion cells," *J. Physiol.*, vol. 285, pp. 275–298, 1978.

[22] J. D. Victor, "The dynamics of the cat retinal x cell centre," *J. Physiol.*, vol. 386, pp. 219–246, 1987.

[23] G. B. Stanley, F. F. Li, and Y. Dan, "Reconstruction of natural scenes from ensemble responses in the lateral geniculate nucleus," *J. Neurosci.*, vol. 19, no. 18, pp. 8036–8042, 1999.

Nicholas A. Lesica (S'01) received the B.S. degree in electrical engineering and computer science from the Massachusetts Institute of Technology, Cambridge, in 2000, and the Ph.D. degree in engineering sciences from Harvard University, Cambridge, in 2005.

He is currently a Postdoctoral Fellow with the Division of Engineering and Applied Sciences, Harvard University. His research interests include signal processing in sensory systems, sensory adaptation, system identification, image processing, and neurophysiology of the early visual pathway.

Garrett B. Stanley (S'96–A'97) received the B.M.E. degree with highest honors from the Georgia Institute of Technology, Atlanta, in 1992, and the M.S. and Ph.D. degrees in mechanical engineering from the University of California at Berkeley in 1995 and 1997, respectively.

From 1995 to 1997, he was an American Heart Association Predoctoral Fellow. From 1997 to 1999, he was a Postdoctoral Fellow with the Neuroscience Division, Department of Molecular and Cell Biology, the University of California at Berkeley, and a National Institutes of Health Postdoctoral Fellow. In 1999, he joined the faculty of the Division of Engineering and Applied Sciences, Harvard University, Cambridge, MA, where he is currently an Associate Professor of bioengineering and an active member of the Harvard-MIT Division of Health Sciences and Technology (HST). His research interests include biological signal processing, experimental and theoretical approaches for understanding neural coding in sensory systems, neuronal point processes, parameter estimation, and the development of devices for recording from and stimulating the nervous system.