

Contrasting Patterns of Receptive Field Plasticity in the Hippocampus and the Entorhinal Cortex: An Adaptive Filtering Approach

Loren M. Frank,¹ Uri T. Eden,¹ Victor Solo,² Matthew A. Wilson,³ and Emery N. Brown¹

¹Neuroscience Statistics Research Laboratory, Department of Anesthesia and Critical Care, Massachusetts General Hospital and Harvard University/Massachusetts Institute of Technology (MIT) Division of Health Sciences and Technology, Boston, Massachusetts 02114, ²School of Electrical Engineering and Telecommunications, University of New South Wales, Sydney, Australia, and ³Center for Learning and Memory, RIKEN/MIT Neuroscience Research Center and Department of Brain and Cognitive Sciences, MIT, Cambridge, Massachusetts 02139

Neural receptive fields are frequently plastic: a neural response to a stimulus can change over time as a result of experience. We developed an adaptive point process filtering algorithm that allowed us to estimate the dynamics of both the spatial receptive field (spatial intensity function) and the interspike interval structure (temporal intensity function) of neural spike trains on a millisecond time scale without binning over time or space. We applied this algorithm to both simulated data and recordings of putative excitatory neurons from the CA1 region of the hippocampus and the deep layers of the entorhinal cortex (EC) of awake, behaving rats. Our simulation results demonstrate that the algorithm accurately tracks simultaneous changes in the spatial and temporal structure of the spike train. When we applied the algorithm to experimental data, we found consistent patterns of plasticity in the spatial and temporal intensity functions of both CA1 and deep EC neurons. These patterns

tended to be opposite in sign, in that the spatial intensity functions of CA1 neurons showed a consistent increase over time, whereas those of deep EC neurons tended to decrease, and the temporal intensity functions of CA1 neurons showed a consistent increase only in the “theta” (75–150 msec) region, whereas those of deep EC neurons decreased in the region between 20 and 75 msec. In addition, the minority of deep EC neurons whose spatial intensity functions increased in area over time fired in a significantly more spatially specific manner than non-increasing deep EC neurons. We hypothesize that this subset of deep EC neurons may receive more direct input from CA1 and may be part of a neural circuit that transmits information about the animal’s location to the neocortex.

Key words: hippocampus; CA1; entorhinal cortex; plasticity; receptive field; adaptive filtering; place field; spatial; coding

Neural responses to stimuli can change over time as a result of experience (Merzenich et al., 1983; Weinberger, 1993; Edeline, 1999; Kaas et al., 1999). This receptive field plasticity has been observed in the CA1 region of the rat hippocampus in two conditions. First, the spatial receptive fields (“place fields”) observed in this region (O’Keefe and Dostrovsky, 1971) develop in approximately the first 5 min of exposure to a novel environment (Wilson and McNaughton, 1993). Second, place field plasticity occurs in familiar environments. Mehta et al. (1997) demonstrated that, over the course of an exposure to a previously visited environment, place fields expand and their centers of mass shift in the direction opposite the animal’s direction of motion. In subsequent work, Mehta et al. (2000) showed that the observed expansion and shifting could be explained by a skewing of the place fields resulting from feedforward plasticity in the connections from CA3 to CA1. This plasticity is reduced in aged animals (Shen et al., 1997) and blocked by the application of 3-(2-

carboxypiperazin-4-yl)-propyl-1-phosphonic acid, an NMDA receptor antagonist (Ekstrom et al., 2001), suggesting that it may be related to learning.

The deep layers of the entorhinal cortex (EC) receive the majority of neocortically bound hippocampal output (Amaral and Witter, 1995). Although CA1 cells tend to have compact place fields, we showed previously that deep EC cells tend to have place fields that cover long, contiguous sections of the environment and frequently show a pattern of activity we termed “path equivalence” in which a given neuron is active across locations associated with the same behavior both within and between environments (Frank et al., 2000). Little is known about the properties of spatial receptive field plasticity in the deep EC, however.

In both the rat and the monkey, cells in the EC and associated temporal cortices can show an enhanced or reduced response to the second presentation of a stimulus compared with the first (Miller et al., 1991; Fahy et al., 1993; Li et al., 1993; Miller and Desimone, 1994; Suzuki et al., 1997; Young et al., 1997; Xiang and Brown, 1998, 1999). Those findings suggest that deep EC neurons, like those in CA1, are likely to show spatial receptive field plasticity. Here we examine the patterns of receptive field plasticity present in simultaneously recorded neurons from CA1 and the deep EC of awake, behaving rats. We developed a spline-based adaptive point process filtering algorithm based on previous work (Brown et al., 2001) that allowed us to characterize the time course of changes in both the spatial structure of recep-

Received Oct. 9, 2001; revised Feb. 7, 2002; accepted Feb. 13, 2002.

This work was supported by the RIKEN/Massachusetts Institute of Technology (MIT) Neuroscience Center, the Center for Learning and Memory at MIT, National Institute of Mental Health Grants MH59733 and MH61637, and National Science Foundation Grant IBN 0081548. We thank Riccardo Barbieri and Ana Nathe for helpful discussions.

Correspondence should be addressed to Loren Frank, Department of Anesthesia and Critical Care, Massachusetts General Hospital, Clinics 3, 55 Fruit Street, Boston, MA 02114. E-mail: loren@neurostat.mgh.harvard.edu.

Copyright © 2002 Society for Neuroscience 0270-6474/02/223817-14\$15.00/0

tive fields and the interspike interval (ISI) structure of the spike trains.

Analyses of the same data focusing on static receptive field properties have been reported previously (Frank et al., 2000, 2001).

MATERIALS AND METHODS

Behavioral methods

The behavioral methods have been published previously (Frank et al., 2000, 2001), and those methods relevant to the current study will be summarized here. A total of four rats were used in this study. Animals were trained before surgery to run for liquid chocolate reward on a 150-cm-long U-shaped track (total length of track, ~300 cm). Once the animal performed at criterion levels, an 18 tetrode microdrive array was implanted, with six tetrodes targeting the CA1 region of the hippocampus and 12 tetrodes targeting the EC.

After recovery, recordings were made from both CA1 and the EC until neurons (putative single cells) could no longer be isolated. On each day of recording, the animal was first placed on the track, and then data collection was started. We therefore analyzed receptive field plasticity from the second to the final pass through the environment. Once neurons could no longer be isolated, the animal was killed, and between one and four lesions were made with each tetrode. The brain was then sectioned and stained. The locations of each tetrode on each day of recording were determined using the number of lesions, the relative locations of the tetrodes, and the recorded depths of each tetrode.

We separated both CA1 and EC neurons into putative excitatory and putative inhibitory subtypes using spike waveform widths and average firing rates as described previously (Frank et al., 2001). For these analyses, only putative excitatory neurons with a mean spike amplitude of at least 70 μ V and a peak spatial firing rate of at least 10 Hz (for the definition of peak spatial rate, see Frank et al., 2000) were considered. A total of 191 CA1 neurons and 56 deep EC neurons were included in the data set used for these analyses. To prepare the data for use with the algorithm, we generated a linear representation of the animal's position, converting each run from one end of the track to the other into a set of locations defined according to their distance from the starting point. We then determined the times when the animal was at each of the two ends of the track. We used those times to define the start and ending times of each pass from endpoint 1 to endpoint 2 and from endpoint 2 to endpoint 1. The start time of each pass was identical to the end time of the previous pass, so no data were excluded. The position data were then interpolated at a time step of 2 msec.

When examining the firing of place cells, it is important to distinguish between periods when the animal is actively exploring and periods when the animal is engaged in other, nonlocomotive behaviors, because these periods are associated with distinct patterns of EEG and neural activity in the hippocampus and EC. During exploration, EEG recorded from the hippocampus and the EC of the rat contains a large amount of power in the theta (6–12 Hz) band, and neuron activity in both CA1 and the EC is strongly modulated at that frequency (Buzsaki et al., 1983; Alonso and Garcia-Aust, 1987; Chrobak and Buzsaki, 1998; Frank et al., 2001). In contrast, when the animal is not actively exploring or attentive, little power in the theta band is present, and, instead, the EEG is characterized by large, irregular activity and the presence of sharp waves (Buzsaki et al., 1983; Buzsaki, 1986) and neural activity in hippocampus is characterized by the presence of synchronized bursts of high-frequency (100–200 Hz) activity (Buzsaki et al., 1992; Chrobak and Buzsaki, 1994). Because place-specific activity is generally associated with theta but not sharp wave EEG patterns, for each data set, we filtered the EEG from a tetrode located in the deep EC between 6 and 14 Hz and computed the rms amplitude in that frequency band. We then determined a cut-off value that we used to indicate the beginning of a non-theta period and a slightly higher cut-off that we used to indicate the end of that period. To ensure that only theta activity was included in our analyses, the values were set to exclude periods when the animal was moving very slowly or was stationary.

Algorithm

Accurately tracking receptive field plasticity requires a method that can produce estimates of the receptive field of a neuron on a short time scale. The most commonly used approach to estimating receptive field structure involves binning the spike data over time and calculating the firing

rate of the neuron in each bin, but these histogram-based approaches have some important limitations. First, histogram-based estimates are difficult to apply in situations in which one wishes to estimate the effects of multiple variables on firing rate. In those situations, it is necessary to sample every combination of the values of the variables sufficiently often to produce an estimate of firing rate for each combination. That requires very large data sets and becomes impractical as the number of variables increases.

Second, estimation using histograms is inefficient, in that a large number of parameters (the values of each bin) are estimated with relatively little data (the number of spikes in each bin) for each parameter. As a result, histogram estimates converge to the actual underlying function relatively slowly (Silverman, 1986; Brown et al., 2002). Third, these estimates will tend to be inaccurate when applied to small amounts of data, implying that histogram-based methods are ill-suited to estimating the plasticity in the structure of receptive fields on short time scales. For example, a rat running on a linear track can reach speeds of up to 1 m/sec, and an average CA1 or deep EC place cell fires at a peak rate of ~20 Hz (Frank et al., 2001). Thus, even with a large bin size of 5 cm, the animal would occupy each bin for only 1/20th of 1 sec, and, on a given pass through the environment, there would be many bins that contained no spikes and a few bins that contained a single spike. The histogram estimate would therefore indicate that the cell fired at 20 Hz in a few locations at 0 Hz everywhere else. The same problem arises when one attempts to estimate the interspike interval structure of the spike train on short time scales, because, on any given pass, a place cell may fire as few as five or 10 spikes, resulting in four to nine ISIs, and it is difficult to accurately estimate a complete interspike interval distribution from so few intervals.

Model-based adaptive estimation algorithms provide an alternate approach that avoids the problems associated with histogram-based estimates. Model-based approaches can easily handle multiple variables and do not require binning over time or space (Brown et al., 1998; Barbieri et al., 2001). When an adaptive algorithm is applied to such a model, the result is an estimate that, at each instant, combines the estimates from the previous time step with the new spiking data (Brown et al., 2001). That combination can produce accurate instantaneous estimates, even when only a few spikes are observed, and can therefore accurately track changes in receptive field structure over time.

Definitions. To construct an adaptive algorithm, it is first necessary to define a conditional intensity or rate function. This conditional intensity function takes the form $\lambda(t|\theta_t, H_t)$, where θ_t is a vector of parameter values at time t that relates H_t , the history of the process up to and including the current time, to the instantaneous firing rate of the neuron (Brown et al., 2002). Because the conditional intensity characterizes the firing rate of the process over time, $\lambda(t|\theta_t, H_t)\Delta t$ is the probability of a spike in $[t, t + \Delta t)$ when there is history dependence in the spike train. Our goal was to identify not only plasticity related to spatial receptive field structure but also plasticity related to the ISI structure of the spike train. We therefore set $\lambda(t|\theta_t, H_t) = \lambda^S(x(t))\lambda^T(t - \zeta_t)$ where $\lambda^S(x(t))$ is a function of the rat's position $x(t)$ at time t , and $\lambda^T(t - \zeta_t)$ is function of the time since the last spike where ζ_t is the time of the last spike. We call these separable functions the spatial and temporal intensity functions, respectively. This approach is similar to one used by Kass and Ventura (2001) for nonadaptive analyses of monkey supplementary eye field data.

In defining the intensity functions, we noted that CA1 and deep EC place fields have complex shapes that cannot be accurately described using Gaussian or other low-dimensional parametric surfaces (Muller et al., 1987; Frank et al., 2000; Mehta et al., 2000). In addition, the ISI histograms of CA1 and deep EC neurons are very different (Frank et al., 2001). To allow us to accurately capture the complex shape of place fields and the ISI structure of CA1 and deep EC spike trains, we defined $\lambda^S(x(t))$ and $\lambda^T(t - \zeta_t)$ as separate spline curves. A spline curve is a set of piecewise continuous polynomials that interpolate between a small number of given coordinates, known as control points (Hearn and Baker, 1996). A spline can approximate any well behaved function and can therefore capture the complex shapes of place fields and ISI functions.

In these analyses, we used cardinal splines to model the spatial and temporal intensity functions. Cardinal splines are two-dimensional functions whose slope at any control point is determined by the magnitudes of the two adjacent control points (Hearn and Baker, 1996). Thus, all cardinal splines are piecewise differentiable (C^3), and every point along a cardinal spline is completely specified by at most four adjacent control point magnitudes. The selection of a cardinal spline-based model provides an implicit set of smoothness and continuity conditions for the firing intensity function.

We assumed that the animal's position could vary continuously be-

tween locations x_{start} and x_{end} , and we defined a discrete set of spatial locations and magnitudes as the I control points $\{(x_i, \theta_i^S)\}_{i=1}^I$ such that $x_i < x_{i+1}$, $x_2 = x_{\text{start}}$, and $x_{I-1} = x_{\text{end}}$. The spatial intensity function is then defined by the cardinal spline as follows:

$$\lambda^S(x) = [u(x)^3 u(x)^2 u(x) 1] \begin{bmatrix} -0.5 & 1.5 & -1.5 & 0.5 \\ 1 & -2.5 & 2 & -0.5 \\ -0.5 & 0 & 0.5 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix} \begin{bmatrix} \theta_{i-1}^S \\ \theta_i^S \\ \theta_{i+1}^S \\ \theta_{i+2}^S \end{bmatrix},$$

for $x \in (x_i, x_{i+1}]$. Here, $u(x) = (x - x_i)/(x_{i+1} - x_i)$ is a fractional measure of the distance between the animal's position $x(t)$ and the two adjacent control points. Note that any point on the spline is fully determined by the four nearest spatial control point magnitudes.

The time locations and associated magnitudes for the control points of the temporal spline were defined similarly as the J pairs: $\{(t_j, \theta_j^T)\}_{j=1}^J$ such that $t_j < t_{j+1}$, where t_{J-1} is a time longer than the maximum interspike interval. The temporal intensity function is then given by the following:

$$\lambda^T(t) = [v(t)^3 v(t)^2 v(t) 1] \begin{bmatrix} -0.5 & 1.5 & -1.5 & 0.5 \\ 1 & -2.5 & 2 & -0.5 \\ -0.5 & 0 & 0.5 & 0 \\ 0 & 1 & 0 & 0 \end{bmatrix} \begin{bmatrix} \theta_{j-1}^T \\ \theta_j^T \\ \theta_{j+1}^T \\ \theta_{j+2}^T \end{bmatrix},$$

for $t \in (t_j, t_{j+1}]$. Here, $v(t) = (t - t_j)/(t_{j+1} - t_j)$, similar to $u(x)$ above. Under this model, the temporal component of the intensity function depends only on the time since the last spike. In addition, because neural firing rates cannot be negative, we defined $\lambda^S(x(t)) = \max(\lambda^S(x(t)), 0)$ and $\lambda^T(t - \zeta_i) = \max(\lambda^T(t - \zeta_i), 0)$.

In previous work, (Brown et al., 2001), we constructed a general adaptive estimation algorithm using point process observations. For any parametric neural spike train model, this algorithm can describe the time evolution of the set of parameters that best explains the observed firing activity. We applied a similar algorithm to the spline-based place field model described above by adaptively altering the spline control point magnitudes based on the instantaneous spiking of each cell. At each instant in time, the update algorithm determines the number of spikes that occurred and compares that with the number expected given the current value of the parameter estimates. This difference, known as the innovation, takes the form $N(t_{k-1}, t_k) - \lambda(t_{k-1}|\theta_{t_{k-1}}, H_{t_k})\Delta_k$, where $N(t_{k-1}, t_k)$ is the number of spikes observed from time t_{k-1} to time t_k , and Δ_k is the length of the current time step. The algorithm then updates the control point magnitudes in such a way as to locally increase or decrease the estimated firing rate depending on that innovation. This algorithm can be run using arbitrarily small time steps, and, for these analyses, we chose a time step of 2 msec.

We constructed a modified adaptive updating algorithm with a constant learning rate with respect to the spline derivatives of the following form:

$$\theta_{ik}^S = \theta_{ik-1}^S + \epsilon_S \frac{d\lambda^S}{d\theta^S} [N(t_{k-1}, t_k) - \lambda(t_{k-1}|\theta_{t_{k-1}}, H_{t_k})\Delta_k] \quad (1.1)$$

$$\theta_{ik}^T = \theta_{ik-1}^T + \epsilon_T \frac{d\lambda^T}{d\theta^T} [N(t_{k-1}, t_k) - \lambda(t_{k-1}|\theta_{t_{k-1}}, H_{t_k})\Delta_k],$$

where $\theta_{ik}^S = [\theta_{1,t_k}^S, \theta_{2,t_k}^S, \dots, \theta_{I,t_k}^S]$ and $\theta_{ik}^T = [\theta_{1,t_k}^T, \theta_{2,t_k}^T, \dots, \theta_{J,t_k}^T]$ are the spatial and temporal control point magnitude vectors at time t_k , ϵ_S and ϵ_T are the spatial and temporal learning rates, $d\lambda^S/d\theta^S$ is the vector of derivatives of the spatial intensity function with respect to each spatial control point magnitude, and $d\lambda^T/d\theta^T$ is the vector of derivatives of the temporal intensity function with respect to each temporal control point magnitude. We derive our adaptive algorithm from our previous work described by Brown et al. (2001) as follows. The previous instantaneous steepest descent algorithm for point processes gave the updating rule for the parameters as follows:

$$\theta_{ik} = \theta_{ik-1} + \epsilon \nabla \log(\lambda(t_{k-1}|\theta_{t_{k-1}}, H_{t_k})) [N(t_{k-1}, t_k) - \lambda(t_{k-1}|\theta_{t_{k-1}}, H_{t_k})\Delta_k],$$

where ∇ denotes the gradient with respect to θ_{ik} , and

$$\epsilon = \begin{bmatrix} \epsilon_1 & 0 \\ 0 & \epsilon_2 \end{bmatrix}$$

is a vector of learning rates. The algorithm in Equation 1.1 follows by setting $\epsilon_S = \epsilon_1 \lambda^S(x(t_{k-1}))$ and $\epsilon_T = \epsilon_2 \lambda^T(t - \zeta_i)$. Thus,

this adaptive estimation algorithm is based on the same instantaneous log-likelihood criterion function described previously but uses a modified learning rate.

Using the definition of $\lambda^S(x(t))$, we can show that the instantaneous derivative of the spatial spline with respect to any spatial control point magnitude is given by the following:

$$\frac{d\lambda^S(x(t))}{d\theta_j^S} = \begin{cases} 0.5u(t)^3 - 0.5u(t)^2 & x(t) \in [x_{i-2}, x_{i-1}] \\ -1.5u(t)^3 + 2u(t)^2 + 0.5u(t) & x(t) \in [x_{i-1}, x_i] \\ 1.5u(t)^3 - 2.5u(t)^2 + 1 & x(t) \in [x_i, x_{i+1}] \\ -0.5u(t)^3 + u(t)^2 - 0.5u(t) & x(t) \in [x_{i+1}, x_{i+2}] \\ 0 & \text{otherwise} \end{cases}.$$

The derivative is nonzero at exactly four control points. The same is true for the derivative of the temporal spline with respect to all temporal control point magnitudes. Thus, if the animal were located at a position between control points x_j and x_{j+1} , the algorithm would update the values of spatial control points $\{x_{j-1}, x_j, x_{j+1}, x_{j+2}\}$. The choice of cardinal splines therefore causes updates to the intensity functions to be local, in that the presence or absence of a spike at a given time will change the spatial and temporal intensity functions over only a small region.

Control point spacing. The control point spacings regulate the minimum width of features within the spline functions that the model can estimate. If the spacings are set too large, the spline will over-smooth the data, but if they are set too small, the spline will fit the spike train rather than the underlying process. The spacing size should be determined by some natural feature size of the underlying process. We therefore chose control point spacings consistent with our experience with the spatial receptive field and ISI structure of CA1 and deep EC neurons (Frank et al., 2001).

The spatial control points were positioned along the track every 10 cm, allowing for spatial features on the scale of ~ 5 cm. To account for both position and direction of movement in the spatial component of the intensity function, we expressed the rat's motion along a 300 cm track by a proportion of its complete periodic trajectory. A full back-and-forth run along the track represents a single spatial period of the animal's movement, so $x(t)$ represents the total distance from the starting point to the current location. It is therefore a one-dimensional linear measure. Additionally, connections were made between the beginning and the end of the periodic path by equating the magnitudes of each of the following pairs of control points: x_1 and x_{J-2} , x_2 and x_{J-1} , and x_3 and x_J . When the magnitude of one element of the pair was changed, then so too was its counterpart. The resulting control point set thus describes a circular spline representing the animal's back-and-forth motion along the track.

The temporal control point locations were positioned along an abscissa representing a continuous range of times since the last spike. A first subset of control points contained regularly spaced locations between 1 and 25 msec at even stretches of 4 msec. The remaining control points were spaced between 25 msec and the longest observed ISI at even stretches of 25 msec. This specific temporal control point spacing was designed to be the largest possible spacing that could extract fine features in the short interspike intervals, corresponding to the refractory period and the tendency of the CA1 neurons to fire in bursts, and broader features in the longer ISIs, corresponding to features such as theta rhythmicity. Larger control point spacings help prevent over-fitting, as is explained in Discussion.

Implementation. Because the spatial and temporal splines are separable, a doubling in the magnitude of all of the control points of the spatial spline and a halving of the magnitude of all of the control points of the temporal spline would result in exactly the same firing rate (conditional intensity) profile over time. To ensure that the relative magnitudes of the two splines remained constant over time, we used a cyclic descent method to implement the adaptive algorithm. Given an initial estimate of the spatial and temporal intensity functions, we fixed the temporal function and ran the algorithm forward, allowing the spatial function to change over time. Beginning with the same initial estimates, we then fixed the evolution of the spatial function and reran the algorithm with an adapting temporal function. Thus, for the first (second) half of an iteration, the changes in the spatial (temporal) function were fixed to be those found for the previous iteration, and the values of the temporal (spatial) function were updated according to the adaptive algorithm. We continued to iterate the model until the patterns of change remained stable from one run to the next using the following convergence criterion:




Figure 1. Examples of changing spatial and temporal intensity functions used for simulation studies. The *left* shows the initial and final values for the spatial intensity functions from S2, and the *right* shows the initial and final values for the temporal intensity functions from T2. The *dots* superimposed on the *curves* represent control points.

a change of no more than the greater of 3 or 10% of any control point magnitude of the spatial intensity function and a change of no more than the greater of 0.3 or 10% of any control point magnitude of the temporal intensity function. These criteria were chosen based on the relative magnitudes of the spatial and temporal intensity functions derived from CA1 and deep EC neurons as described below. We tested this approach using simulations in which the spatial and temporal splines remained constant over time and found that the relative magnitudes of the spatial and temporal control point magnitudes did not change over time.

Initial estimates. It was also necessary to select a method to generate an initial spline estimate at which to start the adaptive update algorithm. Mehta et al. (2000) showed previously that changes in the place field structure can be detected as early as the second pass through the environment. For that reason, it is essential to develop a good estimate for the initial values of the spatial and temporal intensity functions. We investigated defining starting values for the spatial and temporal intensity functions using only the data from the first pass from one end of the environment to the other. Simulation studies showed that estimate often did not accurately capture the true underlying structure of the receptive field. We generated a more accurate initial estimate by repeatedly running the adaptive estimation algorithm backward in time. We initialized the spatial intensity function to have a value of the mean firing rate of a neuron for all positions and initialized the temporal intensity function to have a value of one at all points. Running the algorithm backward in time, we then used the cyclic descent method described above and iterated until the first pass estimates converged using the same criteria as for the forward estimation. This approach yielded reasonable starting values for the simulation studies discussed below. In addition, this approach is conservative, because the estimate for the first pass is obtained by combining the spiking from the first pass with the estimate derived from the subsequent passes. As such, if the first pass estimate is in error, the error will lead to an underestimate rather than an overestimate of the amount of change that occurred.

Simulation studies

Learning rates. To select and evaluate the learning rates that could accurately track changes in receptive field structure we devised a set of simulated random spike processes based on predetermined time-varying spatial intensity and temporal ISI modulation driving functions. The set of spatial driving functions from which simulations were drawn was designed to produce a very wide range of changes. These simulations consisted of the following. S1 was a Gaussian curve with a fixed center located 250 cm from the start of the track, a fixed SD of 20 cm, and a maximum firing rate that increased linearly from 20 to 50 Hz. S2 was a slowly skewing Gaussian produced by the product of a Gaussian curve and a sigmoid with identical centers. The center of the Gaussian curve moved backward linearly from 250 to 233 cm from the beginning of the track. The height of the Gaussian component increased from 15 to 20 Hz. The SD of the Gaussian grew linearly from 20 to 30 cm. The sigmoid function had the form $(1 + e^{-\gamma(x-\mu)})^{-1}$, where the γ parameter varied linearly from 0 to 1, producing a left skew (for the initial and final spatial intensity functions, see Fig. 1). S3 was a rapidly skewing Gaussian similar to that described in S2. Both the height of the Gaussian and the γ parameter of the sigmoid increased as in S2. The center moved backward linearly from 250 to 150 cm, and the SD of the Gaussian increased linearly from 20 to 50 cm.

The set of temporal intensity driving functions used in the simulation studies consisted of the following. T1 was a fixed driving function with two peaks representing a propensity to fire in bursts centered at an ISI of

9 msec with a maximum modulation height of 3.6, and a theta rhythm modulation centered at an ISI of 125 msec with a maximum modulation height of 5.5. This modulation profile was obtained by running the adaptive algorithm on a single CA1 pyramidal cell and using the modulation function estimated for the starting value. T2 was characterized by rapidly increasing bursting and decreasing theta regions. The burst peak increased linearly from no modulation to the maximum value in T1, whereas the theta modulation peak decreased linearly from the peak value in T1 to nothing. T3 was characterized by slowly increasing bursting and decreasing theta regions. The dynamics were as in T2, with both peak heights halved (for initial and final temporal intensity functions, see Fig. 1). T4 was characterized by rapidly decreasing bursting and increasing theta. The burst peak decreased linearly from the peak value in T1 to no modulation, whereas the theta modulation peak increased linearly from zero to the peak value in T1. T5 was characterized by slowly decreasing bursting and increasing theta with dynamics as in T4, with both peak heights halved.

There are a total of 15 different combinations of spatial and temporal driving functions. For each of these, we calculated the instantaneous firing intensity function that would be generated as a rat ran back and forth at a constant speed of 25 cm/sec for 800 sec and simulated five instances of a spike train using the time-rescaling algorithm described by Brown et al. (2002). For each instance, the simulated spike train was then fed back into the adaptive algorithm across a range of learning rates. The spatial learning rates varied from 0.5 to 3.5 by increments of 0.5, and the temporal learning rates varied from 0.0 to 0.35 by increments of 0.05. For every combination of those spatial and temporal learning rates, we calculated the sum over time of the total mean squared error between the estimated control point magnitudes and the values for the actual driving functions. Thus, for each simulation, we produced two error surfaces. The first related the spatial and temporal learning rates to the error in the spatial intensity function, and the second related the spatial and temporal learning rates to the error in the temporal intensity function. We examined each of these error surfaces to determine a range of spatial and temporal learning rates that would minimize both the spatial and temporal errors.

Tracking. To quantify the ability of this adaptive estimation algorithm to track the spatial and temporal features of interest, we examined its accuracy in tracking each of these features individually. The spatial features that we were interested in tracking included the total spatial area, mean, SD, and localized skewness. These statistics were computed as follows: total spatial area $A = \int_{x_1}^{x_2} \lambda^S(x) dx$, center (mean) $\mu = \int_{x_1}^{x_2} x \lambda^S(x) dx$, scale (SD)

$$s = \sqrt{\int_{x_1}^{x_2} (x - \mu)^2 \lambda^S(x) dx},$$

and a localized skewness measure $\text{skew} = (\int_{x_a}^{x_b} (x - \mu)^3 \lambda^S(x) dx) / (As^{-3/2})$, where x_a and x_b were obtained by extending the left and right interquartile whiskers by three times their original magnitude (Velleman and Hoaglin, 1981). This modified version of the skewness measure was used because we found that the skewness computed across all positions was very sensitive to small outliers in the intensity function. These spatial statistics were calculated from the estimated spatial intensity spline at each moment in time. It is important to note that, whereas an increase or decrease in the total area of the spatial intensity function will be associated with an increase or decrease in average firing rate, respectively, the actual firing rate at each time will be the product of the magnitude of the appropriate points of the spatial and temporal intensity functions. As such, the spatial intensity function can be thought of as the place field of the neuron with the temporal structure (bursting, theta modulation, etc.)




Figure 2. A comparison of the estimated and actual spatial and temporal intensity functions. In the *left* (spatial intensity) and *right* (temporal intensity) plots, the gray lines represent the underlying spatial and temporal driving functions, respectively, used to generate a set of observed spike data as described above. In this example, the driving functions were a slowly skewing shifting Gaussian spatial intensity (a shifted version of S2) and fixed ISI modulation functions (T1). The spline estimates of these intensity

functions, obtained by applying our adaptive estimation algorithm to this simulated spike train data, are displayed in the corresponding plot as *black lines*. Overall, the agreement between the actual and estimated functions is very good. The largest apparent deviations are in the temporal intensity function at long interspike intervals. This is attributable to the very small number of long ISIs and the corresponding lack of sampling of these intervals.

factored out. All of the temporal statistics of interest were computed as the area under the estimated temporal ISI modulation spline, computed in the same manner as the spatial area above, localized to specific ISI regions. They included a burst area from 1 to 21 msec, a burst-to-theta area from 21 to 75 msec, a theta modulation area from 75 to 150 msec (corresponding to a modulation of between ~ 7 and 13 Hz), and a theta to two-theta area from 150 to 300 msec.

For the spatial measures of mean, area, and SD and for the temporal measures of the burst, burst-to-theta, and theta areas, we simulated 100 examples of cells in which the value of the measure increased or decreased linearly over the course of the 800 sec simulation by 1, 5, 10, or 50%. All other spatial and temporal characteristics were held fixed. The initial control point magnitudes used for these simulations were chosen to be resemble those seen in real data (see Results) and were as follows: the spatial intensity function was initialized to be a Gaussian with a mean location of 150 cm, a peak firing rate of 20 Hz, and an SD of 15 cm; the temporal intensity function was the same as that used for the static temporal function T1 for the learning rate simulation studies described above.

Application of algorithm to experimental data

Neural data: goodness-of-fit. When applying a set of analyses to data, whether the analyses involve histogram-based estimates or an explicit statistical model, it is important to determine whether the estimates produced by the analyses accurately describe the data. For the simulations, the goodness-of-fit could be evaluated directly by comparing the adaptive estimate with the actual driving function, but because we do not have access to the actual driving function for CA1 and deep EC neurons, another approach is required. In this situation, we determined how well the adaptive model captured the structure of the data by constructing Komologorov–Smirnov (KS) plots for each neuron (Barbieri et al., 2001; Brown et al., 2002). Briefly, to construct each plot, we computed, for each ISI, $\tau_k = \int_{l_{k-1}}^{l_k} \lambda(u|\theta_u, H_u) du$, where l_k is the time of the k -th spike, and $\lambda(u|\theta_u, H_u)$ is the conditional intensity function. According to the time rescaling theorem (Brown et al., 2002), if $\lambda(u|\theta_u, H_u)$ correctly describes the rate process underlying the spike train, the τ_k values will be exponentially distributed with unit rate. Thus, setting $z_k = 1 - e^{-\tau_k}$ produces a set of independent uniform random variables on the interval [0, 1]. A measure of the agreement between the cumulative distribution of the z_k values and the quantiles of the uniform distribution directly evaluates how well the original model agrees with the spike train data. We applied that analysis to the model with adapting spatial and temporal components, as well as to a model with an adaptive spatial component and a static temporal component set to one for all ISIs, to determine whether the addition of an adapting temporal component significantly improved the fit to the data.

We analyzed the changes in the spatial and temporal intensity functions by computing the same measures used above for the simulations. To allow for comparison across different animals and data sets, we sampled these measures at 20 points, evenly spaced in time along each pass from one end of the U track to the other. For the measures applied to the spatial receptive field, each direction of motion on the track was considered separately. As mentioned above, a single temporal intensity function was used for both directions of motion along the track. We defined a “pass” for the temporal function to be a single forward and backward pass in each direction along the track so that the spatial and temporal functions would be displayed on the same relative scale. To produce a meaningful estimate of the mean and skewness of the spatial intensity function, only those passes with a single place field (for definition, see

Frank et al., 2000) were included in the analysis of center location and skewness (CA1, $n = 232$ place fields; deep EC, $n = 62$ place fields).

Experimental data: relationship between plasticity and spatial coding. We also examined the relationship between patterns of plasticity in the spatial intensity function and the precision of spatial coding in CA1 and deep EC neurons. For each neuron, we fit a linear regression to the area of the spatial receptive field of the neuron and determined whether the line, fit to the area, changed height by $>20\%$ over the course of the run. Examining the CA1 and the deep EC separately, we divided the neurons into two groups, one whose area increased by at least 20%, and the other whose area either increased by $<20\%$ or decreased. We computed the average position information coefficient (Skaggs et al., 1993), which measures the relative position specificity of each neuron averaged over the course of the entire run. We also divided the neurons into groups based on 10, 15, and 25% changes to ensure that our results were not attributable to the specific choice of the 20% cutoff.

RESULTS

Simulation results

Learning rate

A broad grid search over spatial and temporal learning rates resulted in both spatial and temporal error surfaces with single minima. Thus, for each error surface, a single combination of spatial and temporal learning rates was found that minimized the differences between the adaptive estimate and the actual driving function. There was some variability in the location of that minimum across 15 simulations and the five instances of each simulation, but in all cases in which the simulation involved changing driving functions, the minimum was close to the point corresponding to a spatial learning rate of 2.0 and a temporal learning rate of 0.15. We therefore chose those values for our simulation studies.

Tracking

An example of the instantaneous actual and estimated intensity functions is shown in Figure 2. This example shows the capacity of the algorithm to accurately estimate both the spatial and temporal intensity functions. Video 1 similarly shows the entire time course of an example of the algorithm tracking simulated data [Videos 1–4 are available on the *Journal of Neuroscience* website (www.jneurosci.org)]. To quantify the accuracy of tracking for the spatial and temporal area measures and for the spatial scale measure, we first normalized the trajectory of the measure (the set of values across time) for each simulated or real neuron by its mean. We then fit a regression line to the trajectory of the measure produced by the algorithm and compared the amount of change predicted by that regression line with the actual amount of change used in the simulation. The errors are reported in terms of absolute percentages extrapolated to 40 total passes through the environment to permit easy comparison of simulated and real data. Thus, an error of 2% on a simulation with a change of 10%

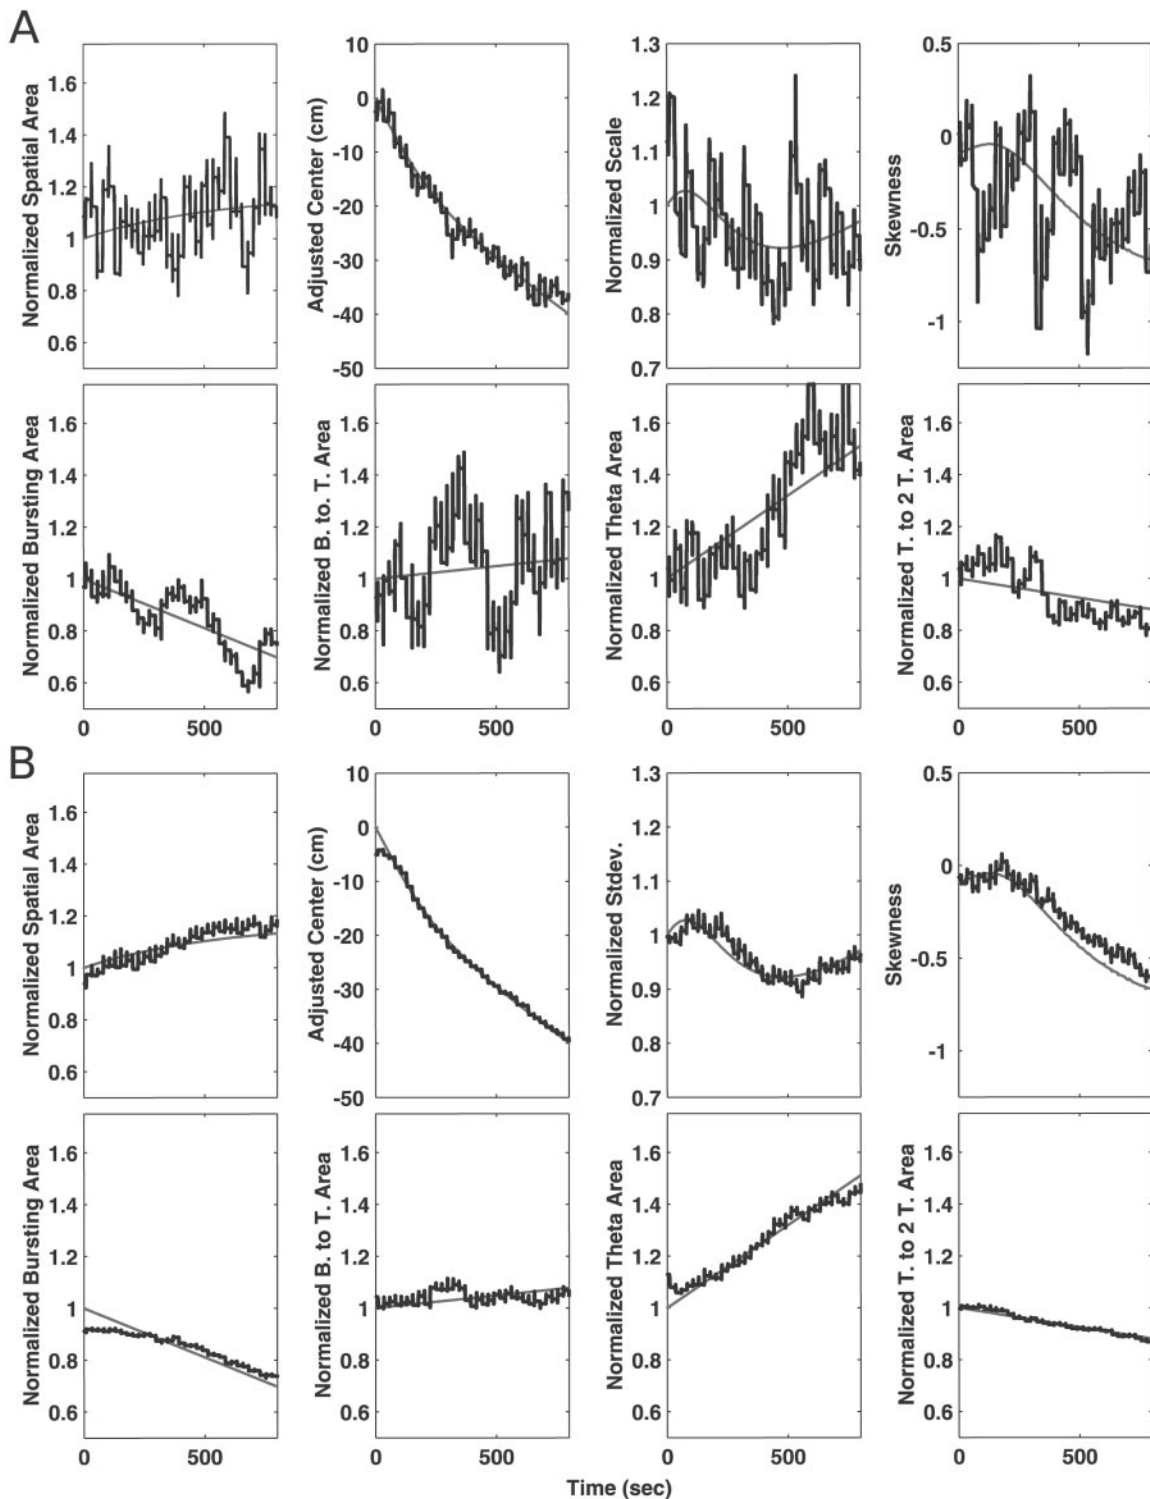


Figure 3. Examples of tracking from simulation S2–T2. In each plot, the gray line shows the actual trajectory of the measure, and the thicker black line shows the measure derived from the algorithm. Because our interest was in accurately capturing the trends in the data, we plotted the spatial area and scale measures, as well as all of the temporal measures, on a normalized scale. To construct these plots, we normalized the actual statistic by its mean and shifted it to start at one. We similarly normalized the estimated statistic by its mean and shifted it by the same amount as the actual statistic. This produces plots in which changes are expressed as a proportion of the initial value. To construct the adjusted center plots, we also shifted the actual center curve to begin at zero and shifted the estimated curve by the same amount as the actual curve, so the two could be easily compared. *A*, Examples of the tracking of measures of the spatial and temporal intensity functions from a single simulated cell. The top row shows the spatial measures, and the bottom row shows the temporal measures. The tracking of the spatial area was somewhat noisy because of variability of the spike train and our choice of a learning rate that could track very fast changes, but nevertheless the trend captured by the estimate is similar to that present in the real data. The tracking of the center was very accurate. The tracking of the scale and the skewness were also, like that of the area measure, somewhat noisy, but once again the trends in the estimate follow the actual trends. Similarly, the trends in the measures of area under different sections of the temporal intensity curve were similar to the actual trends. *B*, The same measures averaged over 100 simulations. The trajectory of the normalized area, adjusted (Figure legend continues.)

indicates that the algorithm produced an estimated trend whose trend led to a change between 8 and 12%. We quantified the movement of the center in the same manner, except we did not normalize the trajectory.

Here we summarize the data by reporting the maximum deviation over the increasing and decreasing simulations. For the 1, 5, and 10% levels of changes of the simulated areas and SDs, the magnitude of the estimated trend was within 2% of the magnitude of the actual trend. For the 50% change, the magnitude of the estimated trend differed by up to 6% for the area and up to 9% for the SD. These deviations were always conservative, in that the algorithm underestimated the actual change. The errors in the mean were at most 1 cm for levels of change of 10 cm or less. The error for movements of 50 cm was ~3 cm.

The tracking of the burst, burst-to-theta, and theta area measures showed similar levels of deviation from the true trajectory. These deviations increased from between 1 and 6% for the 1, 5, and 10% changes to as much as 20% for the 50% change. As for the large deviations between the estimated and actual spatial intensity functions, the larger deviations were always conservative in that, when there was a large error, the algorithm consistently underestimated the actual magnitude of change. An analysis of these large errors showed that an increase or decrease in one region of the temporal modulation function produces an increase or decrease in firing rate over all positions; the algorithm was distributing the change into both the temporal and the spatial intensity functions. In effect, these simulated data can be well fit by a number of different combinations of spatial and temporal intensity functions.

We then examined the tracking for simulations such as those used for the learning rate estimation, in which multiple measures were changing simultaneously. We found that the algorithm tracked well for both individual examples (Fig. 3A) and populations of simulated neurons (Fig. 3B). The estimated trajectory for the measures of the spatial and temporal intensity functions derived from a single example were variable because of the stochasticity of the spike train. Nonetheless, the trends from the estimated measures were similar to those of the actual driving functions. When these trends were averaged across 100 simulated cells, the agreement between the estimates and the actual function was very good for all spatial and temporal measures. There was a small deviation between the trends for the spatial center and the burst and theta area because of the conservative nature of the first pass estimate. These results were consistent across the 15 types of simulations, in that the algorithm generally tracked the slower changes very well. The faster changes were also generally tracked well, and, when the tracking was less accurate, the estimated trends indicated a smaller change than was actually present.

Experimental data results

To ensure that our results were robust to the choice of learning rates, we ran the algorithm with all combinations of spatial learning rates of 1.0, 2.0, and 3.0 and temporal learning rates of

0.075, 0.15, and 0.225. If the learning rates were too low to accurately follow changes in the firing rate, then the adaptive model would smooth these changes, suggesting that the changes occur over a longer period than was actually the case. As the learning rates were increased, the adaptive model would more and more closely approximate the changes in the actual firing rate until it accurately captured the true trend. Increasing the learning rate beyond that point would result in a model capturing the true trend but was more variable than at lower learning rates. Thus, if the observed rate of change of a measure is the same across different learning rates, we are more confident that we have accurately captured the true rate of change.

We also ran the algorithm with a spatial learning rate of 2.0 and a temporal learning rate of 0.0 to determine whether the adaptive temporal component made a substantial contribution to the model goodness-of-fit. Any results that were not consistent across all combinations are mentioned in the text. In applying the algorithm to the data, we used the same control point spacings and the same implementation described above. To determine whether trends seen in the estimated temporal and spatial intensity functions reflected actual changes in the data, we compared the magnitude of the trend with the magnitude of the errors for each of the measures determined above (Simulation results, Tracking). If the trend for the real data was larger than the errors, we can conclude that the trend is real. As an example, area changes of 10% or less were associated with errors of 2% or less. Thus, if we observed a change of 10% in the estimated spatial intensity function, we would conclude that there was a real trend whose magnitude was between ~8 and 12%.

Goodness-of-fit

An examination of the KS plots revealed that the inclusion of an adapting temporal function made a substantial difference in the fit of the model to the data (Fig. 4). When we examined the model without an adapting temporal component, we found that, for the 95% (99%) confidence bounds, the model predictions were within the bounds for 0 of 191 (3 of 191) CA1 neurons and 2 of 56 (4 of 56) deep EC neurons. When an adaptive temporal component was added, the agreement between the model and the data improves substantially, and, for the 95% (99%) confidence bounds, the model predictions were within the bounds for 71 of 191 (115 of 191) CA1 neurons and 25 of 56 (41 of 56) deep EC neurons.

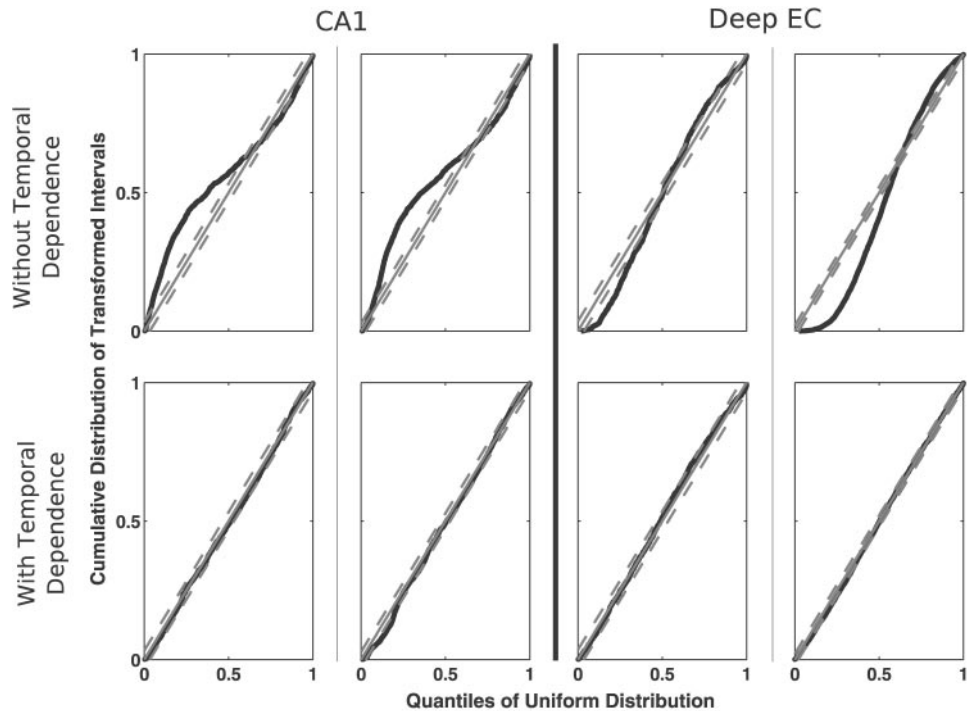
CA1 receptive field plasticity

Examples of the instantaneous estimates of spatial and temporal intensity functions for two CA1 neurons are shown in Figure 5. In both cases, only motion in one direction along the U track is shown. The complete evolution of the spatial and temporal intensity functions for a representative CA1 neuron is shown in Video 2. When we examined the averaged trajectories for the spatial and temporal measures computed for CA1 neurons, we found that the area of the spatial receptive field increases and the center of the spatial receptive field moved backward over time

←

center, and normalized scale are all very close to the actual trajectories. The small deviation visible at the start of the area and center measures is attributable to the first pass estimate. The skewness also tracked well, although there were slightly larger errors for the skewness than for the area or the mean. The algorithm was also able to simultaneously track the trends in the different regions of the temporal function. The estimates of the burst-to-theta (21–75 msec) and theta to two-theta (150–300 msec) area were essentially identical to those for the actual function, and the estimates for the burst (1–21 msec) and theta (75–150 msec) regions follow the true values very closely, with a slight deviation at the beginning that was, as with the area and mean, related to the algorithm used to determine the first pass estimate. Thus, the algorithm is able to track changes in both the spatial receptive field and in the interspike interval structure of spike train data.

Figure 4. Examples of Komologorov–Smirnov plots showing the agreement between the statistical model and the data for CA1 and deep EC neurons. The plots were constructed as described in Materials and Methods. The *top row* shows the results when the model contained only an adaptive spatial component (e.g., the initial temporal intensity function was everywhere set to one and the temporal learning rate was set to zero), and the *bottom row* shows the results when the model contains both adaptive spatial and temporal components with a temporal learning rate of 0.15. In both cases, the spatial learning rate was 2.0. Each *plot* shows the correct distribution (*solid gray line*), the 95% confidence intervals for the correct distribution (*gray dashed lines*), and the distribution resulting from the adaptive model (*black line*). Each *column* shows the results for a single cell in two conditions. Although the spatial function changes over time, the agreement between the model and the data is poor, suggesting that the model does not accurately capture the firing rate of the spike train. The plots for the CA1 neurons suggest that the lack of fit is most pronounced for low z_k values, which correspond primarily to shorter ISIs (Barbieri et al., 2001), suggesting that the model predicts fewer short intervals than are present in the data. In contrast, the plots for the deep EC neurons show lack of fit for the short intervals, which suggests that the model predicts more short intervals than were present in the data. The relative scarcity of short intervals is most likely to be attributable to the pronounced refractory period of most deep EC neurons (Frank et al., 2001). The *bottom row* shows the plots when the model includes adaptive spatial and temporal components. When the model includes both components, the fit to the data is much improved, and, even when the model distribution lies outside of the 95% confidence intervals, it is nevertheless much closer to them than without the temporal intensity function.



(Fig. 6A). The time courses of these changes were different, in that, although the area increased approximately linearly over time and had a relatively small SE, the center reached its most extreme values by approximately the 15th pass and appeared to be relatively more variable. In addition, there was a significant decrease in the scale of the spatial receptive fields. Finally, the modified measure of skewness had consistent negative value throughout the run but did not appear to change over time.

We then examined the changes in the temporal intensity function and found that there were specific changes in the intensity associated with different intervals over time (Fig. 6B). The area of the bursting region was approximately constant, whereas the burst-to-theta (20–75 msec) region showed a brief increase in area over the first 10 passes and then a return to the starting levels. The theta region (75–150 msec) showed by far the largest and most consistent change, increasing by 10% within the first 10 passes and then remaining constant. That increase occurred at a much faster rate than the increase in the spatial area. Finally, the area of the theta to two-theta (150–300 msec) region remained approximately constant.

Deep EC receptive field plasticity

Examples of the instantaneous spatial and temporal intensity functions for two deep EC neurons are shown in Figure 7. The spatial intensity functions of deep EC neurons were less spatially concentrated than those of CA1 neurons, reflecting the differences in spatial specificity between the regions. The temporal intensity functions of deep EC neurons also tended to be very different from their CA1 counterparts in that there was generally very little area below 10 msec. The temporal intensity functions of deep EC neurons also frequently had a peak near 100 msec, but

the structure beyond that theta peak was not clearly consistent across multiple cells.

Just as the individual spatial and temporal intensity functions of deep EC neurons differ substantially from those of CA1 neurons, the patterns of receptive field plasticity found in the activity of deep EC neurons also differ from those of CA1 neurons. There was no clear tendency for the center of the spatial receptive fields of the deep EC neurons to move over time (Fig. 8A). In addition, on average, the area of the spatial intensity actually decreased by ~6% over 40 passes. That slow decrease was significantly greater than the errors in tracking found for simulations. There was no trend in the movement of the center, but the decrease in area was accompanied by a small but significant decrease in the scale of the spatial intensity function of ~7%. Finally, there was no tendency for the spatial intensity functions to be positively or negatively skewed.

When we examined the intervals of the temporal intensity function (Fig. 8B), we found a slight but nonsignificant decrease in the burst, theta, and theta to two-theta regions. The burst-to-theta region decreased by ~9% over 40 passes, a change that was significantly greater than zero. We also found that, when the model was run without an adapting temporal component (e.g., with a temporal learning rate of zero), the decrease in the spatial component was ~14%, indicating that, on average, the firing rate of deep EC neurons decreased by 14% over the course of 40 passes.

Relationship between plasticity and spatial coding

When we examined groups of CA1 neurons with increasing (>20%) or non-increasing spatial intensity function areas, we found that these groups did not differ in terms of either their

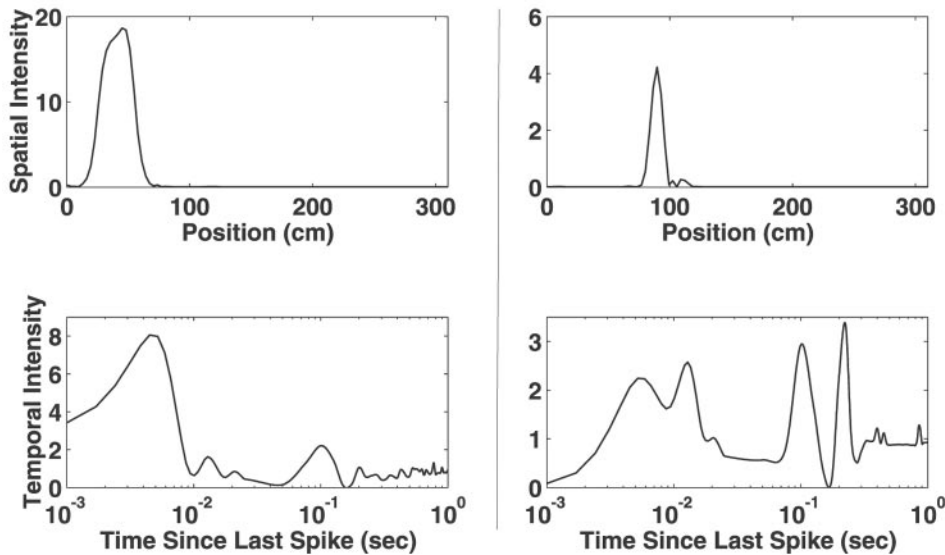


Figure 5. Examples of instantaneous estimates of the spatial and temporal intensity functions from two CA1 neurons. Each *column* represents the functions for a single neuron. The spatial receptive fields of these two neurons are both sharply peaked, although their widths differ. The temporal intensity functions for the two neurons are very different, suggesting that the interspike interval structure of their spike trains differ. The neuron whose intensity function is shown on the *left*, for example, has a very large peak at ~ 5 msec, indicating that this neuron has a strong tendency to fire in bursts. The intensity function on the *right* has two peaks at low intervals, suggesting that this neuron fires in three spike bursts. In addition, this intensity function has sharp peaks at ~ 110 and 220 msec, indicating that the neuron fires frequently at the period of the theta rhythm and on alternating periods of the theta rhythm. These peaks correspond to interspike intervals that occurred frequently, and the peaks were present through-

out the run, suggesting that they are consistent features of the temporal intensity function. These *plots* provide an easily interpretable summary of the firing properties of the neurons at each moment in time, and videos that show the evolution of these functions provide a powerful tool for examining the evolution of receptive fields.

temporal intensity functions or their average position information measures. Thus, groups of CA1 neurons whose spatial intensity functions increased were not, on average, either more or less position specific than the other CA1 neurons whose spatial intensity functions did not increase.

In contrast, the deep EC group, whose spatial intensity function increased in area by at least 20% ($n = 14$), differed from the non-increasing group ($n = 42$) in two respects. First, the temporal intensity functions differed (Fig. 9). We found that there was a substantial difference in the trends for the burst-to-theta region of the temporal intensity functions but no difference in the other regions. The burst-to-theta region of the increasing area group increased by an average of $\sim 3\%$ over the course of 40 passes. Although that increase was small and close to the error levels seen in the simulations, it was very different than the 14% decrease seen in the non-increasing group. We verified that this difference was significant by examining the error from groups of 14 simulated cells drawn from the 100 simulations in which the burst-to-theta region changed. We found that the error levels for these smaller groups were generally small ($<5\%$), and, when the errors were larger, they once again resulted from the algorithm underestimating the magnitude of the real trend. Thus, the increasing and non-increasing groups differed significantly in the amount of change seen in the burst-to-theta region of the temporal intensity function.

Second, there was a relationship between patterns of receptive field plasticity and position information for deep EC neurons. Figure 10 shows examples of two deep EC neurons whose spatial intensity functions increased in area by at least 20% over the course of 40 passes and two deep EC neurons whose spatial intensity functions did not increase. Those neurons with an increasing spatial intensity function had smaller and more concentrated spatial receptive fields than non-increasing neurons. We quantified the spatial specificity by computing the mean and SEs for the position information for each group. The mean position information coefficient was significantly higher for the neurons in the increasing group (mean \pm SE, increasing group, 1.50 ± 0.17 bits per spike; non-increasing group, 0.87 ± 0.10 bits per spike;

$p < 0.005$; paired t test). These differences, and the differences in the temporal intensity function discussed above, were significant for all percentages tested (10, 15, 20, and 25%) across all combinations of learning rates, except when both the smallest spatial and smallest temporal learning rates were used. In that case, the relationship between spatial area change and position information approached but did not reach significance, suggesting that the lowest learning rates may not be high enough to accurately track receptive field changes. To ensure that the observed difference was not related to the quality of the recordings, we also examined the distributions of spike waveform amplitude for the two groups. There was no difference between the groups (median amplitudes, increasing group, $86 \mu\text{V}$; non-increasing group, $97 \mu\text{V}$; $p > 0.2$; Wilcoxon rank sum test). Videos 3 and 4 show two deep EC neurons, one with a spatial intensity function showing increasing spatial area and the other showing a decreasing spatial area.

DISCUSSION

We constructed a point process adaptive filter to estimate a dynamic firing rate function from a neural spike train and to therefore track rapid changes in neural receptive fields. At every step of the recursive update algorithm, we compare the expected probability of observing a spike with whether or not one was observed. At observation times when a spike occurs, the firing estimate is increased significantly, and, at all other observation times, it is decreased by a much smaller amount. With reasonable choices for learning rates, the end result is that the algorithm produces an accurate dynamic estimate of receptive field structure.

Simulation results

Our analysis of the behavior of the algorithm suggested that it was consistently well behaved across a wide range of learning rates, but that, not surprisingly, higher learning rates were required to track fast changes in the spatial and temporal learning rates. Higher values of the learning rate make the algorithm more “local” in that the estimate of the intensity functions at any given

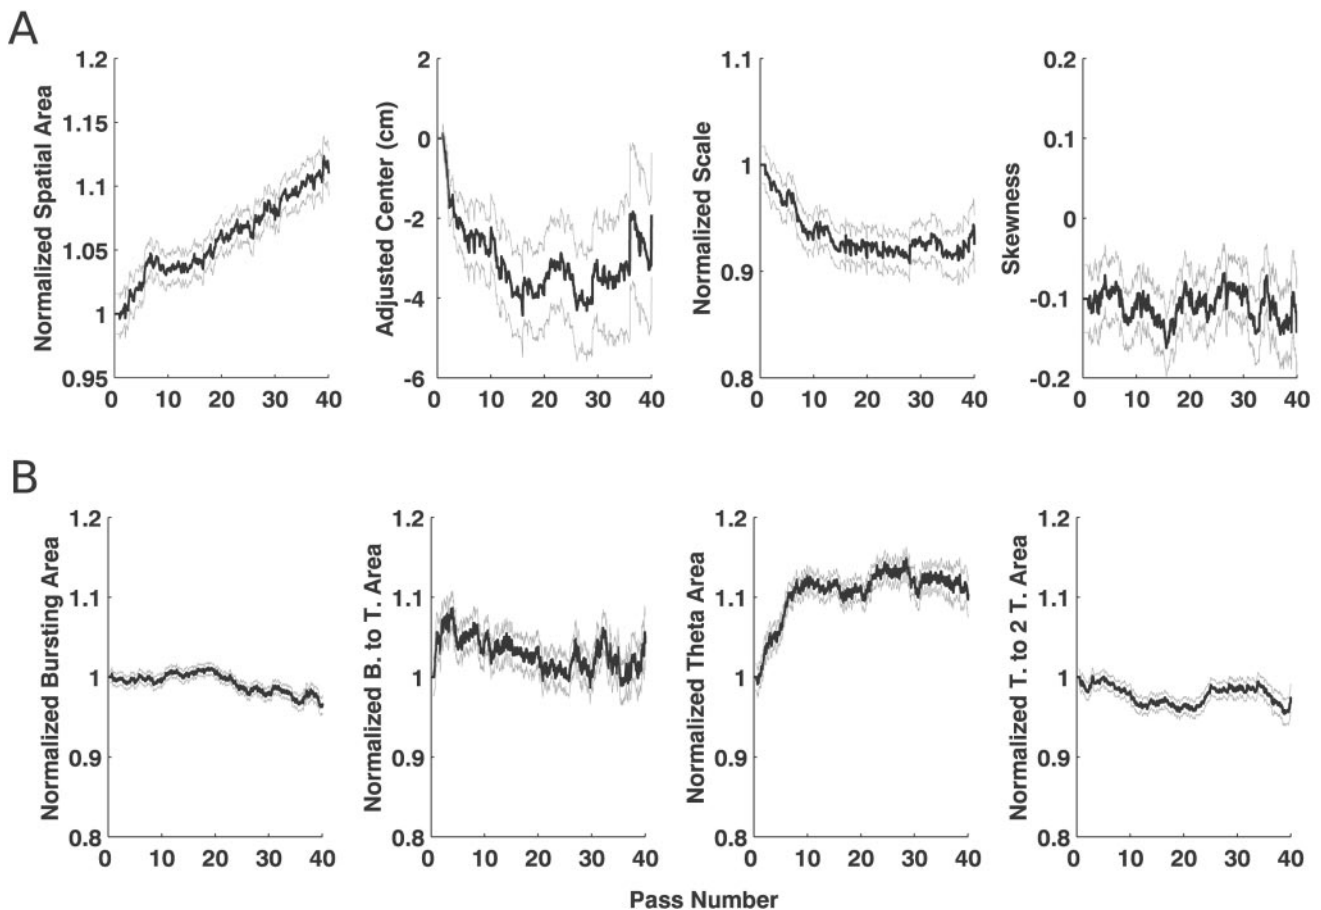


Figure 6. Mean and SEs for the estimates of the spatial and temporal statistics for CA1 neurons. The *black line* in each *plot* represents the trajectory of the mean for the statistics, and the *gray lines* represent the mean \pm SE. The units on the *x-axis* are passes, in which one pass corresponds to the animal moving from one end of the U track to the other. The values on the *y-axis* are, for all statistics except the spatial mean and skewness, normalized to express the change in terms of the proportion of the initial value. The spatial means have been adjusted to start at 0 cm. *A*, The trajectories of the statistics for the spatial function. The area shows a clear and consistent increasing trend while the mean moves backwards over time. The scale also shows a clear decreasing trend, indicating that, even as the spatial intensity function increases in area, it decreases in extent, suggesting that, on average, CA1 activity becomes more spatially localized over time. Finally, the skewness measure had a consistently negative value but did not show a trend. These results are generally consistent with previous findings for CA1 cells. *B*, The trends for the temporal statistics. The burst (1–21 msec) region and the theta to two-theta regions (150–300 msec) showed little change, and the burst-to-theta region (21–75 msec) showed a small initial increase, followed by a return to baseline. In contrast, the theta region (75–150 msec) showed a rapid increase of $\sim 10\%$. That increase is consistent with the predictions of the model proposed by Mehta et al. (2000).

time will depend more on the very recent observations than on the more remote observations. In our analyses, we found that, even with a learning rate high enough to track very fast changes, the algorithm was nevertheless able to track smaller changes.

Accurately estimating dynamic spatial and temporal intensity functions from a spike train is a nontrivial operation, and our algorithm performed well for a wide variety of simulated patterns of plasticity. The tracking of the spatial intensity function was very accurate in that the deviations between the actual and estimated trends were generally small. Thus, this algorithm can be used to estimate the instantaneous structure of spatial receptive fields. The tracking of the temporal statistics was slightly less accurate, but, importantly, when the estimate of the algorithm reflected a change of 10% or more, a change of that magnitude or larger was always present in the actual temporal intensity function.

When we examined simulations in which multiple variables were changing simultaneously, we found that the tracking of the temporal functions was very accurate, with errors that were consistently in the direction of conservative estimates. We believe

that these complex situations effectively disambiguated spatial and temporal changes and that the algorithm was therefore able to accurately estimate the trends in both functions. Overall, the simulation results indicate that the algorithm can track simultaneous changes in both the spatial and temporal intensity functions and that, when the changes are large, any errors tend to be conservative. Thus, this indicates that our algorithm can be used to estimate dynamic spatial and temporal intensity functions from real neural data.

Goodness-of-fit tests

Using the KS plots, we determined that an adaptive model with only a spatial component was not able to accurately describe the evolution of the firing rate of CA1 or deep EC neurons. Adding an adaptive temporal component greatly improved the fit, indicating that the spike trains that would be produced by the model were very similar to those generated by the neurons. That temporal component was therefore able to capture a great deal of the temporal structure of the spike trains, including the refractory period, bursting, and theta rhythmicity. This result is consistent

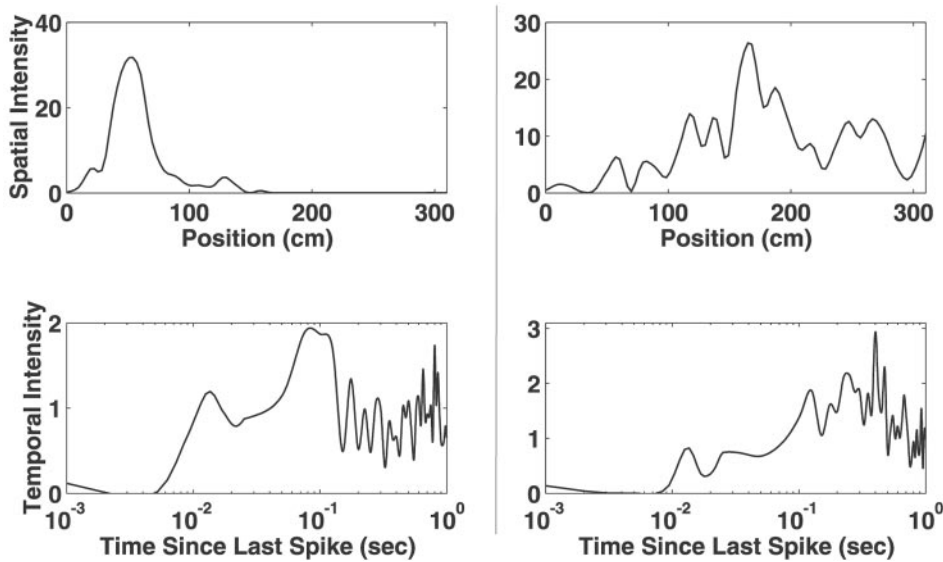


Figure 7. Examples of instantaneous estimates of the spatial and temporal intensity functions from two deep EC neurons. Each column represents the functions for a single neuron. The spatial intensity functions illustrate the variety of spatial receptive field shapes found in the deep EC. The temporal intensity functions of deep EC neurons reflected the very small number of intervals <10 msec and generally had a peak near 120 msec, the period of the theta rhythm. Once again, the variability present in the portion of the temporal intensity function corresponding to larger ISIs (>300 msec) is attributable to the infrequent sampling of this region.

with our previous findings for nonadapting or static models, in which including non-Poisson temporal structure improved the model fit to the data (Barbieri et al., 2001; Kass and Ventura, 2001). We should note that to obtain a better model it may be necessary to include other variables, such as velocity (McNaughton et al., 1983), and spatiotemporal interactions, such as phase precession (O'Keefe and Recce, 1993). Nevertheless, because it would be very difficult to estimate the changing spatial and temporal structure of CA1 and deep EC neural spike trains using histogram-based methods, we suggest that adaptive models may provide a more accurate account of the nature of receptive field plasticity in these regions.

In considering the goodness-of-fit of this model, it is important to consider the possibility that we over-fit the spike train. To avoid that problem, we chose to widely space the control points. If we had spaced the control points closely together and used high learning rates, we could have produced a model that over-fit the data, in that there would be a sharp peak in the spatial and temporal intensity functions at the location and interspike interval corresponding to each spike. That type of model could accurately describe the spike train by specifying a very high firing rate at the exact time and location of each spike and a very low rate elsewhere, but the resulting rate curve would not provide any insights into the changing relationship between firing rate and the variables of interest: location and interspike interval. Fortunately, even in that case, an analysis of the dynamics of the model will not tend to produce misleading results, because it would be essentially the same as analyzing the original spike train.

Receptive field plasticity in CA1

We found that not only did the spatial firing properties of CA1 neurons change over time, but so too did the interspike interval properties of the neurons. Our findings for CA1 receptive field plasticity were generally consistent with those reported previously (Mehta et al., 1997, 2000; Shen et al., 1997; Ekstrom et al., 2001). We found that the spatial intensity function tended to shift backward along the direction of motion and increased in size with experience. We also found that the scale decreased substantially over time, suggesting that the overall distribution of firing narrows over time, even as the overall width of the place field as measured in previous studies increases. These results suggest that

the precision of spatial coding in CA1 increases with experience. In addition, we found that the spatial intensity functions were negatively skewed and that skew remained approximately constant over time.

Mehta et al. (2000) reported that place field skew developed over the first few passes through the environment. The difference between those results and the findings presented here is likely to be attributable to two factors. First, Mehta et al. reported that the majority of the skewness change occurred between the first and second passes through the environment. We did not have data from the animal's first pass through the environment, so those changes were not available to us. Second, we restricted the region over which we measured skewness because we found that small outliers had a large effect on the skewness measure. That restriction may also have contributed to this result. Because the data analyzed here are not directly comparable with those analyzed by Mehta et al. (2000), these results should not be seen as contrary to previous findings.

The changes in the spatial intensity function were accompanied by an increase in the theta region of the temporal intensity function while all other regions remained essentially constant, indicating that the neurons fired progressively more often at ISIs close to the period of the theta rhythm. That is consistent with the model of Mehta et al. (2000), in that the model predicted that the expansion and skewing of the place field in the direction opposite the animal's direction of movement would result in a greater tendency to fire at theta intervals near the beginning of the place field. We also found that the time course of those changes was faster than the changes in the spatial area but slower than the changes in skewness reported previously, indicating that there are different time scales for different aspects receptive field plasticity in CA1 and suggesting that there may be different mechanisms as well.

Receptive field plasticity in the deep EC

Our findings concerning receptive field plasticity in the deep EC contrast with those for CA1. While CA1 spatial receptive fields tended to increase in area, deep EC receptive fields generally decreased in area. That decrease was seen both in the spatial intensity function and most dramatically in the 21 to 75 msec region of the temporal intensity function. Thus, not only were the

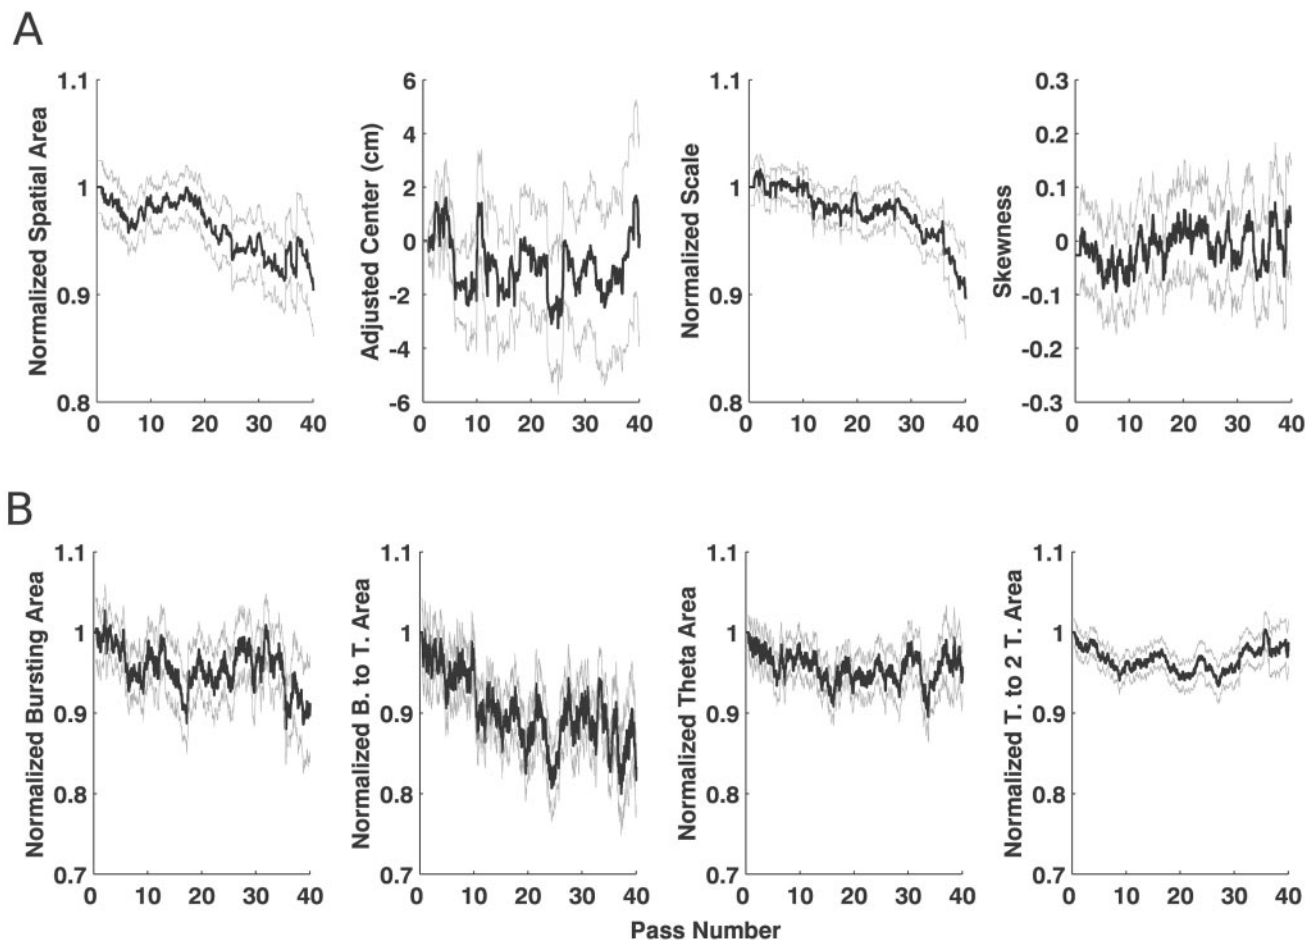


Figure 8. Mean and SEs for the estimates of the spatial and temporal statistics for CA1 neurons. The *black line* in each *plot* represents the trajectory of the mean for the statistics, and the *gray lines* represent the mean \pm SE. *A*, The trajectories for the statistics of the spatial function. Unlike CA1 neurons, the spatial intensity functions of deep EC neurons tended to decrease in area by $\sim 6\%$ over the course of 40 passes. The mean of the spatial intensity function also showed no tendency to move backward, and no trends were found for the skewness. There was also a decrease of $\sim 7\%$ in the scale, suggesting that the fields become somewhat more concentrated over time. *B*, The trajectories for the statistics of the temporal function. All regions of the temporal intensity function showed decreases, but the decrease was significant only for the burst-to-theta region in which the decrease averaged 9%. Thus, the decrease in the spatial area of the intensity function was associated with a relatively large decrease in the burst-to-theta region of the intensity function. When the model was run without a temporal intensity function, the spatial area decreased by $\sim 15\%$, indicating that the combination of spatial and temporal function declines results in the decrease of $\sim 15\%$ in the firing rate of these neurons over 40 passes through the environment.

dynamics of the spatial intensity function different between CA1 and the deep EC, but so too were the dynamics of the temporal intensity function. Those differences, when combined with the contrasting patterns of spatial coding seen in CA1 and the deep EC (Frank et al., 2000, 2002), indicate that substantial processing takes place in the circuitry between CA1 and the deep EC.

The overall decrease seen in the activity of deep EC cells is consistent with previous observations from other types of tasks in which repeated presentations of the same stimulus lead to a reduction in response in the EC and other temporal lobe regions in both rats and monkeys (Miller et al., 1991; Li et al., 1993; Fahy et al., 1993; Miller and Desimone, 1994; Zhu et al., 1995; Suzuki et al., 1997; Young et al., 1997; Xiang and Brown, 1998, 1999). Thus, similar processes may mediate the reduction in response in spatial, as well as nonspatial, tasks.

Furthermore, Miller and Desimone (1994), in an experiment in which recordings were made from primate inferotemporal (IT) cortex in animals performing a serial match-to-sample task, showed that, although the majority of IT neurons showed decremental responses, a subset of them showed incremental responses

to the second presentation if it was a match to the sample stimulus. That suggests that there is a correlation between receptive field plasticity and neural coding properties. We therefore examined our data for a similar correlation.

We found that those deep EC neurons that showed a substantial increase in spatial area differed from other deep EC neurons in two ways. First, whereas the temporal intensity functions of the non-increasing group showed a dramatic decline in the burst-to-theta (21–75 msec) region, those of the increasing area group did not show that decline. Very little is known about the deep EC, so it is difficult to attribute that change to a specific mechanism, but, nonetheless, these findings suggest that receptive field plasticity in the deep EC is associated with a mechanism that selectively changes the propensity of deep EC neurons to fire at those ISIs.

Second, we found that the activity of deep EC cells whose spatial intensity function tended to increase in area was significantly more place specific than that of deep EC cells with a non-increasing spatial intensity function. It is important to note that there is no reason that increasing spatial area should be associated with position specificity; an increase in area can result




Figure 9. The temporal intensity functions for the increasing and non-increasing groups of deep EC neurons. Deep EC neurons were split into two groups based on the change observed in their spatial intensity function, one whose spatial area increased by at least 20% and another whose spatial area increased by <20% or decreased. *A*, The temporal intensity functions for the increasing group. *B*, The temporal intensity functions for the non-increasing group. Once again, the burst-to-theta (21–75 msec) region showed by far the greatest changes, with a small but nonsignificant increase in the case of the increasing group and a large, significant decrease of ~14% for the decreasing group. Thus, the difference in spatial receptive field change between these two groups are associated primarily with changes in the propensity for these neurons to fire in intervals of 21 to 75 msec. Unfortunately, little is known about the physiology of neural circuitry in the deep EC, so these results are, at present, difficult to relate to a specific mechanism.

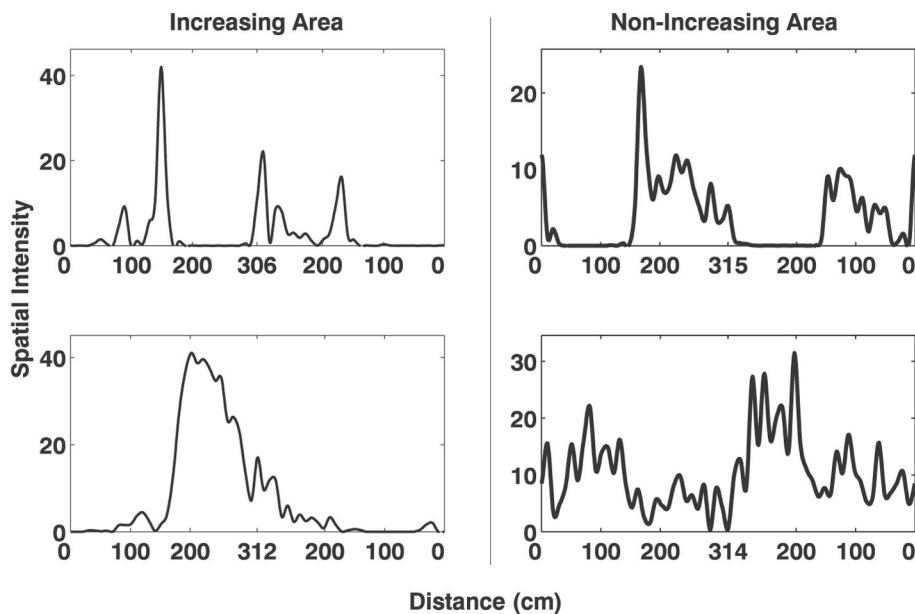


Figure 10. Examples of instantaneous receptive fields from deep EC neurons with increasing (*left column*) and non-increasing (*right column*) spatial intensity function areas. Each *plot* represents the spatial intensity function of a single deep EC neuron. On average, the neurons with increasing spatial intensity functions had more restricted spatial receptive fields than neurons with non-increasing spatial intensity functions.

in either diminished or increased spatial specificity depending on where the increase occurs. Thus, these deep EC neurons resemble CA1 neurons in terms of both their relatively high position specificity and the plasticity seen in their spatial receptive fields.

Because cells in the subiculum, the other major hippocampal region input to the deep EC, show only weakly place-specific firing (Barnes et al., 1990; Sharp and Green, 1994; Sharp, 1997, 1999), the relationship between spatial coding and receptive field

plasticity in the deep EC may reflect differences in the amount of direct CA1 input that these deep EC cells receive. This subpopulation may be important in transmitting information about the animal's spatial location to other brain regions as opposed to other deep EC neurons that represent relationships among spatially distinct positions (Frank et al., 2000).

REFERENCES

- Alonso A, Garcia-Austt E (1987) Neuronal sources of theta rhythm in the entorhinal cortex of the rat. II. Phase relations between unit discharges and theta field potentials. *Exp Brain Res* 67:502–509.
- Amaral DG, Witter MP (1995) Hippocampal formation. In: *The rat nervous system* (Paxinos C, ed), pp 443–493. New York: Academic.
- Barbieri R, Quirk MC, Frank LM, Wilson M, Brown EN (2001) Construction and analysis of non-Poisson stimulus response models of neural spike train activity. *J Neurosci Methods* 105:25–37.
- Barnes CA, McNaughton BL, Mizumori SJ, Leonard BW, Lin LH (1990) Comparison of spatial and temporal characteristics of neuronal activity in sequential stages of hippocampal processing. *Prog Brain Res* 83:287–300.
- Brown EN, Frank LM, Tang D, Quirk MC, Wilson MA (1998) A statistical paradigm for neural spike train decoding applied to position prediction from ensemble firing patterns of rat hippocampal place cells. *J Neurosci* 18:7411–7425.
- Brown EN, Nguyen DP, Frank LM, Wilson MA, Solo V (2001) An analysis of neural receptive field dynamics by point process adaptive filtering. *Proc Natl Acad Sci USA* 98:12261–12266.
- Brown EN, Barbieri R, Ventura V, Kass RE, Frank LM (2002) The time-rescaling theorem and its application to neural spike train data analysis. *Neural Comput* 14:325–346.
- Buzsaki G (1986) Hippocampal sharp waves: their origin and significance. *Brain Res* 398:242–252.
- Buzsaki G, Leung LW, Vanderwolf CH (1983) Cellular bases of hippocampal EEG in the behaving rat. *Brain Res* 287:139–171.
- Buzsaki G, Horvath Z, Urioste R, Hetke J, Wise K (1992) High-frequency network oscillation in the hippocampus. *Science* 256:1025–1027.
- Chrobak JJ, Buzsaki G (1994) Selective activation of deep layer (V–VI) retrohippocampal cortical neurons during hippocampal sharp waves in the behaving rat. *J Neurosci* 14:6160–6170.
- Chrobak JJ, Buzsaki G (1998) Gamma oscillations in the entorhinal cortex of the freely behaving rat. *J Neurosci* 18:388–398.
- Edeline JM (1999) Learning-induced physiological plasticity in the thalamo-cortical sensory systems: a critical evaluation of receptive field plasticity, map changes and their potential mechanisms. *Prog Neurobiol* 57:165–224.
- Ekstrom AD, Meltzer J, McNaughton BL, Barnes CA (2001) NMDA receptor antagonism blocks experience-dependent expansion of hippocampal “Place fields.” *Neuron* 31:631–638.
- Fahy FL, Riches IP, Brown MW (1993) Neuronal activity related to visual recognition memory: long-term memory and the encoding of recency and familiarity information in the primate anterior and medial inferior temporal and rhinal cortex. *Exp Brain Res* 96:457–472.
- Frank LM, Brown EN, Wilson MA (2000) Trajectory encoding in the hippocampus and entorhinal cortex. *Neuron* 27:169–178.
- Frank LM, Brown EN, Wilson MA (2001) A comparison of the firing properties of putative excitatory and inhibitory neurons from CA1 and the entorhinal cortex. *J Neurophysiol* 86:2029–2040.
- Frank LM, Brown EN, Wilson MA (2002) Entorhinal place cells: trajectory encoding. In: *The neural basis of navigation: evidence from single cell recording* (Sharp PE, ed), pp 97–116. New York: Kluwer.
- Hearn D, Baker MP (1996) *Computer graphics, C version*. Upper Saddle River, NJ: Prentice Hall.
- Kass RE, Ventura V (2001) A spike-train probability model. *Neural Comput* 13:1713–1720.
- Kaas JH, Florence SL, Jain N (1999) Subcortical contributions to massive cortical reorganizations. *Neuron* 22:657–660.
- Li L, Miller EK, Desimone R (1993) The representation of stimulus familiarity in anterior inferior temporal cortex. *J Neurophysiol* 69:1918–1929.
- McNaughton BL, Barnes CA, O'Keefe J (1983) The contributions of position, direction, and velocity to single unit activity in the hippocampus of freely-moving rats. *Exp Brain Res* 52:41–49.
- Mehta MR, Barnes CA, McNaughton BL (1997) Experience-dependent, asymmetric expansion of hippocampal place fields. *Proc Natl Acad Sci USA* 94:8918–8921.
- Mehta MR, Quirk MC, Wilson MA (2000) Experience-dependent asymmetric shape of hippocampal receptive fields. *Neuron* 25:707–715.
- Merzenich MM, Kaas JH, Wall JT, Sur M, Nelson RJ, Felleman DJ (1983) Progression of change following median nerve section in the cortical representation of the hand in areas 3b and 1 in adult owl and squirrel monkeys. *Neuroscience* 10:639–665.
- Miller EK, Desimone R (1994) Parallel neuronal mechanisms for short-term memory. *Science* 263:520–522.
- Miller EK, Li L, Desimone R (1991) A neural mechanism for working and recognition memory in inferior temporal cortex. *Science* 254:1377–1379.
- Muller RU, Kubie JL, Ranck JBJ (1987) Spatial firing patterns of hippocampal complex-spike cells in a fixed environment. *J Neurosci* 7:1935–1950.
- O'Keefe J, Dostrovsky J (1971) The hippocampus as a spatial map. Preliminary evidence from unit activity in the freely-moving rat. *Brain Res* 34:171–175.
- O'Keefe J, Recce ML (1993) Phase relationship between hippocampal place units and the EEG theta rhythm. *Hippocampus* 3:317–330.
- Sharp PE (1997) Subicular cells generate similar spatial firing patterns in two geometrically and visually distinctive environments: comparison with hippocampal place cells. *Behav Brain Res* 85:71–92.
- Sharp PE (1999) Subicular place cells expand or contract their spatial firing pattern to fit the size of the environment in an open field but not in the presence of barriers: comparison with hippocampal place cells. *Behav Neurosci* 113:643–662.
- Sharp PE, Green C (1994) Spatial correlates of firing patterns of single cells in the subiculum of the freely moving rat. *J Neurosci* 14:2339–2356.
- Shen J, Barnes CA, McNaughton BL, Skaggs WE, Weaver KL (1997) The effect of aging on experience-dependent plasticity of hippocampal place cells. *J Neurosci* 17:6769–6782.
- Silverman BW (1986) *Density estimation for statistics and data analysis*. New York: Chapman and Hall.
- Skaggs WE, McNaughton BL, Gothard K, Markus EJ (1993) An information-theoretic approach to deciphering the hippocampal code. In: *Advances in neural information processing 5* (Hanson SJ, Cowan JD, Giles CL, eds), pp 1030–1037. San Mateo, CA: Kaufman.
- Suzuki WA, Miller EK, Desimone R (1997) Object and place memory in the macaque entorhinal cortex. *J Neurophysiol* 78:1062–1081.
- Velleman PF, Hoaglin DC (1981) *Applications, basics, and computing of exploratory data analysis*. Boston: Duxbury.
- Weinberger NM (1993) Learning-induced changes of auditory receptive fields. *Curr Opin Neurobiol* 3:570–577.
- Wilson MA, McNaughton BL (1993) Dynamics of the hippocampal ensemble code for space. *Science* 261:1055–1058.
- Xiang JZ, Brown MW (1998) Differential neuronal encoding of novelty, familiarity and recency in regions of the anterior temporal lobe. *Neuropharmacology* 37:657–676.
- Xiang JZ, Brown MW (1999) Differential neuronal responsiveness in primate perirhinal cortex and hippocampal formation during performance of a conditional visual discrimination task. *Eur J Neurosci* 11:3715–3724.
- Young BJ, Otto T, Fox GD, Eichenbaum H (1997) Memory representation within the parahippocampal region. *J Neurosci* 17:5183–5195.
- Zhu XO, Brown MW, Aggleton JP (1995) Neuronal signalling of information important to visual recognition memory in rat rhinal and neighbouring cortices. *Eur J Neurosci* 7:753–765.