

How voltage-dependent conductances can adapt to maximize the information encoded by neuronal firing rate

Martin Stemmler¹ and Christof Koch²

¹ Innovationskolleg Theoretische Biologie, Invalidenstr. 43, D-10115 Berlin, Germany

² Computation and Neural Systems Program, California Institute of Technology 139-74, Pasadena, California 91125, USA

Correspondence should be addressed to M.S. (stemmler@itb.biologie.hu-berlin.de)

Information from the senses must be compressed into the limited range of responses that spiking neurons can generate. For optimal compression, the neuron's response should match the statistics of stimuli encountered in nature. Given a maximum firing rate, a nerve cell should learn to use each available firing rate equally often. Given a set mean firing rate, it should self-organize to respond with high firing rates only to comparatively rare events. Here we derive an unsupervised learning rule that continuously adapts membrane conductances of a Hodgkin-Huxley model neuron to optimize the representation of sensory information in the firing rate. Maximizing information transfer between the stimulus and the cell's firing rate can be interpreted as a non-Hebbian developmental mechanism.

Spiking neurons in any sensory system translate, or encode, the sights, sounds and smells of the outside world into trains of action potentials. Yet how does a neuron, through experience, 'learn' to represent the stimuli it receives? Furthermore, what general rules describe how a neuron should adapt to changes in its synaptic input environment?

Information measures provide a quantitative framework for investigating these questions. Analyses of spike trains from visual cortical neurons show that over 80% of the total information seems to be carried by the time-averaged firing rate^{1,2}, the simplest, although not the only, means of encoding information. If a neuron's range of available firing rates were unbounded, it could transmit infinite information. However, the firing rate can neither fall below zero nor exceed some maximal firing rate dictated by the action potential's duration and associated refractory period. As a result, sensory information must be compressed into a finite range of firing rates.

Under these limits, optimizing information transmission through a single neuron in the presence of uniform, additive noise has an intuitive interpretation, namely that the entire range of the neuron's available firing rates should be evoked by synaptic input, so that no particular firing rate occurs more frequently than any other. Such optimal neuronal representation of incoming sensory information takes full advantage of the regularity of natural stimuli. For vision, this regularity is reflected in the typical likelihoods of encountering particular visual contrasts, spatial orientations or colors³⁻⁵. Given these likelihoods, an optimized neural code would devote increased representation to more commonly encountered features. An analogous principle for non-spiking neurons has been tested experimentally by matching the statistics of naturally occurring visual contrasts to the response amplitudes of the blowfly's large monopolar cell⁶.

Alternatively, a neuron could seek to minimize the metabolic energy expenditure due to spiking⁷ while still transmitting as much

information as possible. In this case, a neuron should maintain a low average firing rate. If the mean firing rate is some set value, the optimal distribution of firing rates is no longer uniform, but decays exponentially, giving greater prominence to low firing rates. Such distributions have been observed in mammalian visual cortical neurons responding to natural scenes⁸.

In translating a synaptically induced conductance change x into a firing rate f (Fig. 1a), a neuron needs to map the probability distribution $p(x)$ of encountering any one input x onto a firing rate probability distribution $p(f)$. Each distribution is associated with an entropy, a number that reflects how many distinguishable synaptic inputs or firing rates, respectively, exist in the presence of noise, as weighted by the frequency of occurrence. These entropies are represented by the size of the circles in the Venn diagram (Fig. 1b).

The mutual information $I(x, f)$ (measured in bits) between the synaptic conductance changes x and the firing rates f specifies how much one can learn about one variable's value by knowing the other. In terms of the Venn diagram, increasing the mutual information $I(x, f)$ corresponds to increasing the overlap between the stimulus entropy and the entropy of the firing statistics. Whereas the former is fixed for a particular environment, the firing rate entropy can be influenced by the neuron's intrinsic properties. By changing these intrinsic properties, the neuron can maximize the latter entropy, which corresponds to achieving the optimal firing rate distribution described above.

In translating a range of different inputs into a set of distinct firing rates that maximize the mutual information $I(x, f)$, spiking neurons achieve a more difficult task than keeping the firing rate or excitatory postsynaptic potential amplitude constant under changing conditions, two tasks for which learning rules that adapt the membrane conductances have been proposed^{9,10}. Learning the proper representation of stimulus information goes beyond simply correlating input and output and must invoke the voltage-dependent properties of membrane conductances. This differs

articles

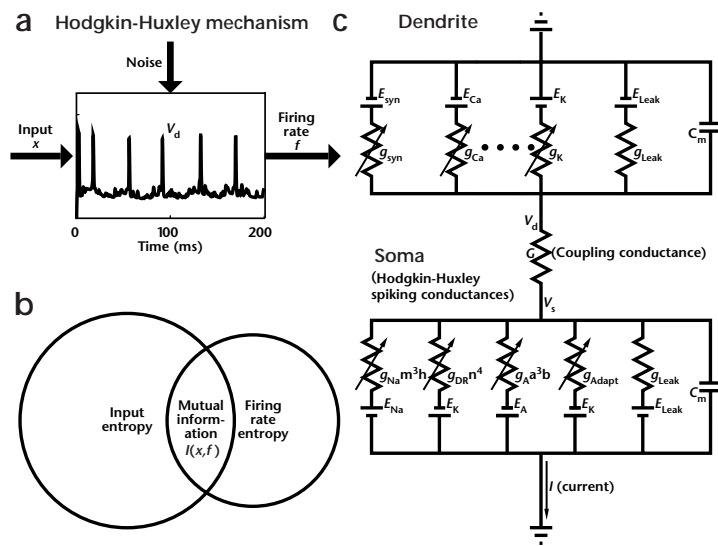


Fig. 1. The model neuron. **(a)** The sum of synaptic conductance input and synaptic noise was converted into voltage spikes by a Hodgkin-Huxley-like model. Spikes originate in the somatic compartment of the model in **(c)** and 'back-propagate' to the dendritic compartment. Inset, dendritic voltage trace. The set of all synaptic inputs x is associated with an entropy, illustrated by the Venn diagram in **(b)**, of fixed size for a given sensory environment. Likewise, the set of firing rates has a corresponding entropy, which is of adaptable size. The goal of mutual information maximization is to increase the overlap of these two entropies, the mutual information $I(x, f)$. **(c)** Circuit diagram of the model neuron. The two compartments represent the cell's soma and dendrites. The single calcium and single potassium conductance in the dendritic compartment represent multiple conductances of each type. To maximize the information transfer, the parameters for the calcium and potassium voltage-dependent conductances in the dendritic compartment were iteratively adjusted, whereas the somatic conductances responsible for the cell's spiking behavior remained fixed.

from the classic postulate of Hebb¹¹, in which synaptic learning in networks is a consequence of correlated activity between pre- and postsynaptic neurons. To optimize the firing rate distribution $p(f)$, the neuron must discover and learn the underlying structure of the inputs.

The possibility that neurons can indeed learn an 'optimized neural code' is made plausible by evidence that neurons regulate synaptic efficiency¹² as well as intrinsic membrane conductances. For instance, blocking spiking activity *in vitro* increases sodium current and decreases TEA-sensitive potassium current, resulting in a marked increase in the cell's excitability¹³. In addition, conditioning causes changes in the postsynaptic excitability of neurons^{14,15}, and the homeostatic regulation of cellular rhythmic activity following the reduction of synaptic input is accomplished through the adjustment of membrane conductances^{16,17}. Finally, during development, the nature and frequency of incoming stimuli regulate both anatomical structure and the cellular distribution of ionic conductances in

neurons^{18,19}. We propose a theory to describe how voltage-dependent membrane conductances can adjust to match a neuron's firing rate behavior to its stimulus environment. This theory provides predictions for the developmental time course of the Na^+ , Ca^{2+} and K^+ conductances present in the cell membrane²⁰.

RESULTS

To explore how neurons adapt their intrinsic properties to represent stimuli, we constructed a simplified model of a neuron consisting of two electrotonic compartments (Fig. 1c). The cell body contains the classic Hodgkin-Huxley sodium and delayed-rectifier potassium conductances, with the addition of a transient potassium A current and a slowly acting potassium current that models firing rate accommodation. The soma is coupled through a conductance G to a dendritic compartment containing the synaptic input conductance as well as six adjustable voltage-dependent Ca^{2+} and K^+ membrane conductances. Standard Hodgkin-

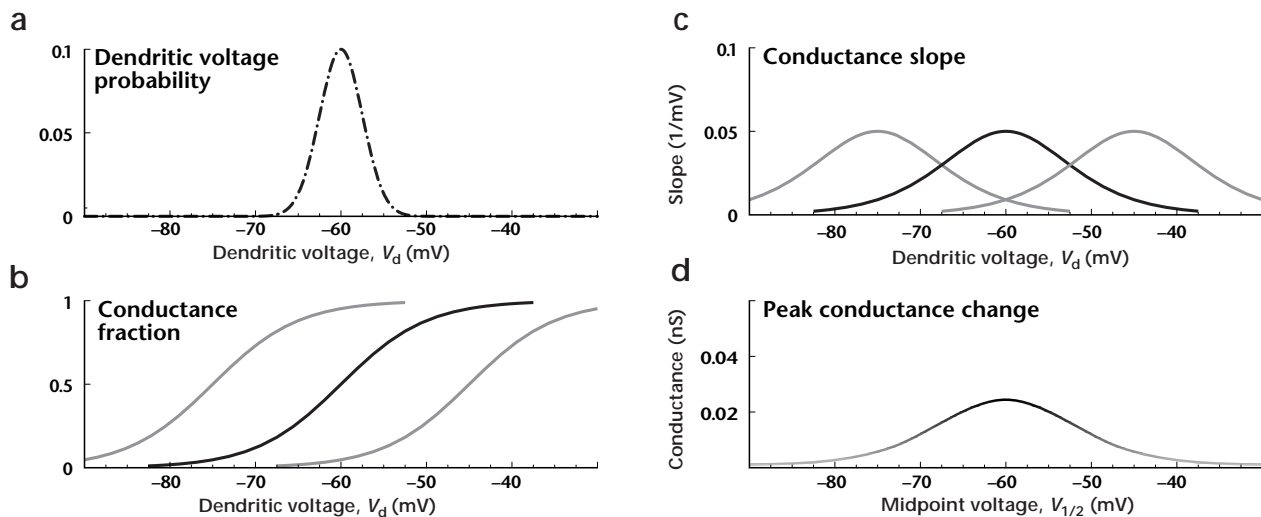


Fig. 2. The first term in the learning rule for changing the peak conductance of voltage-dependent conductances in the dendritic compartment. **(a)** Synaptic inputs to the cell give rise to a distribution of voltages in the dendritic compartment. **(b)** Three S-shaped Boltzmann functions describe the average conductance fractions for three non-inactivating calcium conductances in this example. **(c)** Peak conductances adapt in proportion to the activation functions' slopes. The conductance plotted in black increases the most, the gray conductances less. **(d)** Weighting the slope by the probability distribution in **(a)** yields the expected change $\langle \Delta \bar{g} \rangle$ in the peak conductance, displayed as a function of the activation midpoint voltage.

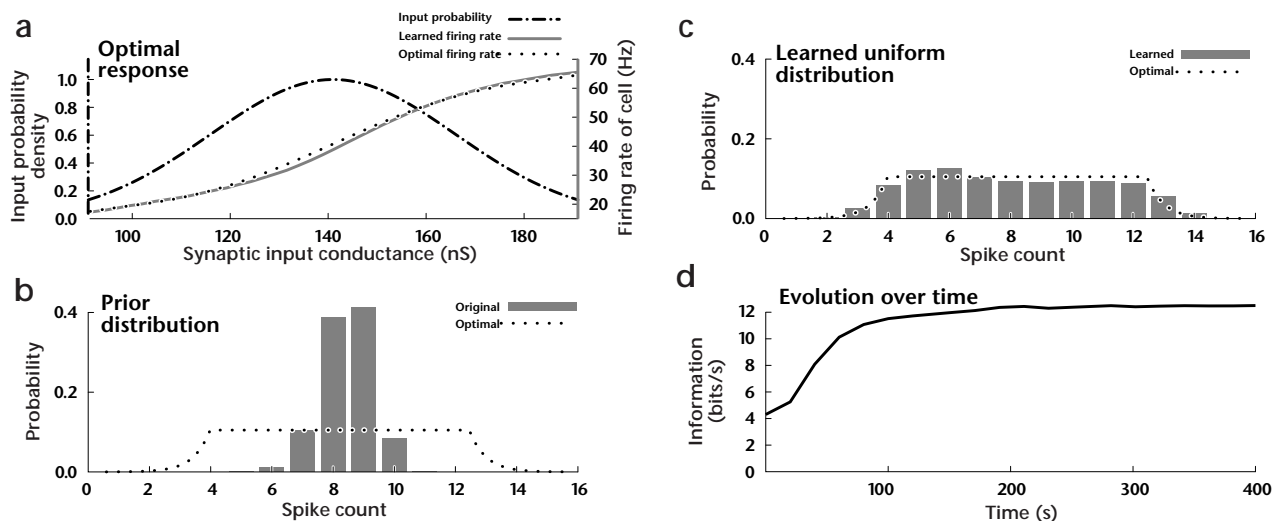


Fig. 3. Learning the optimal firing rate response with the constraint that the firing rate must lie between 20 and 60 spikes per second. **(a)** The inputs to the model are synaptic conductances, drawn randomly from a non-negative Gaussian distribution (dash-dotted line). The learning rule—maximizing the information in the cell's firing rate—was used to adjust the peak conductances, midpoint voltages and slopes of the 'dendritic' Ca^{2+} and K^{+} conductances over the course of 10.9 (simulated) minutes. The learned firing rate response curve (gray line) is asymptotically proportional to the cumulative probability distribution of inputs (dotted line), which is simply the integral under the Gaussian. **(b)** Histograms of spike counts recorded within the 200-ms stimulus time intervals before learning. Dotted line, optimal distribution of spike counts, given that these spike counts can be approximated by continuous firing rates. **(c)** Learning shifts the spike count distribution from the peaked distribution in **(b)** to a much broader one, so that, in response to randomly selected synaptic inputs, a neuron produces each spike count between 4 and 12 equally often. **(d)** The information rate during adaptation, as estimated from ten-second windows and averaging over 60 trials of stochastic learning from the same initial state.

Huxley-like equations govern the membrane potential and a set of activation and inactivation variables (see Methods).

We allow the neuron to continually modify the voltage-dependent membrane conductances in the direction that maximizes information transmission. A prescription for changing any parameter associated with a voltage-dependent conductance (such as the associated channel density) is termed a 'learning rule' for that parameter. However, no internal teacher exists

that can tell the neuron whether it is performing correctly; the learning must occur in an unsupervised manner.

To be plausible, a principle for the self-organization of cellular properties should not require that a neuron explicitly calculate the mutual information $I(x, f)$ and store the result of this calculation, say, in the intracellular calcium concentration. Learning theory^{21–23} neatly obviates this problem by allowing the neuron to 'guess' at how it should change its properties (see Methods).

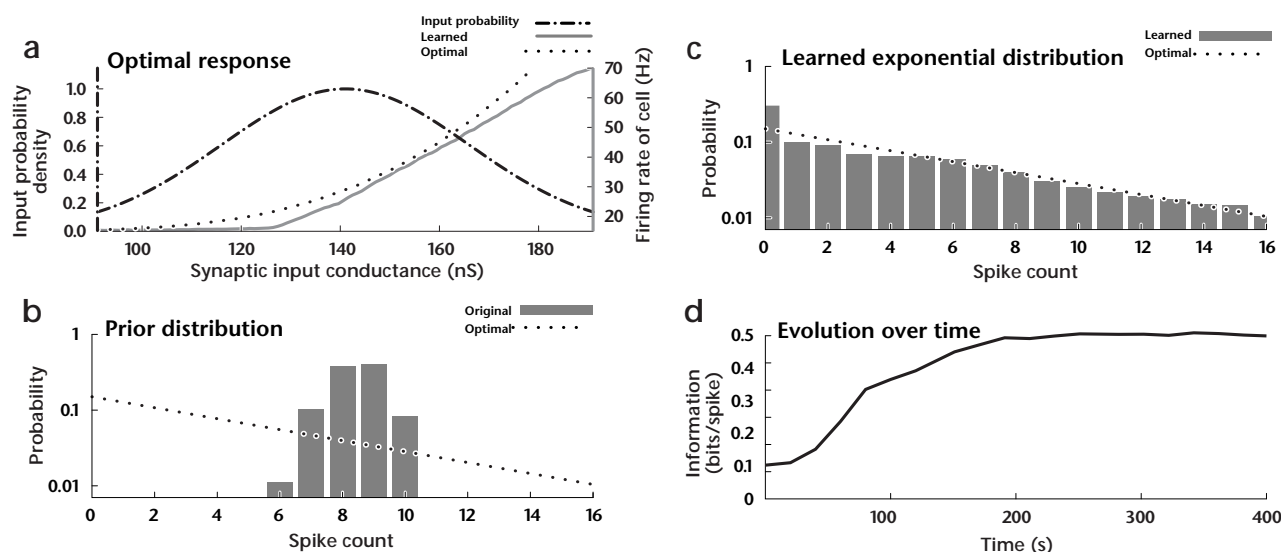


Fig. 4. Learning the optimal firing rate response curve assuming an average firing constraint of 30 Hz. The same synaptic input distribution was used as in Fig. 3. With this constraint, the target distribution (dotted line in **b** and **c**) is exponential, indicated by the straight line in the semilogarithmic graph, instead of being uniformly flat.

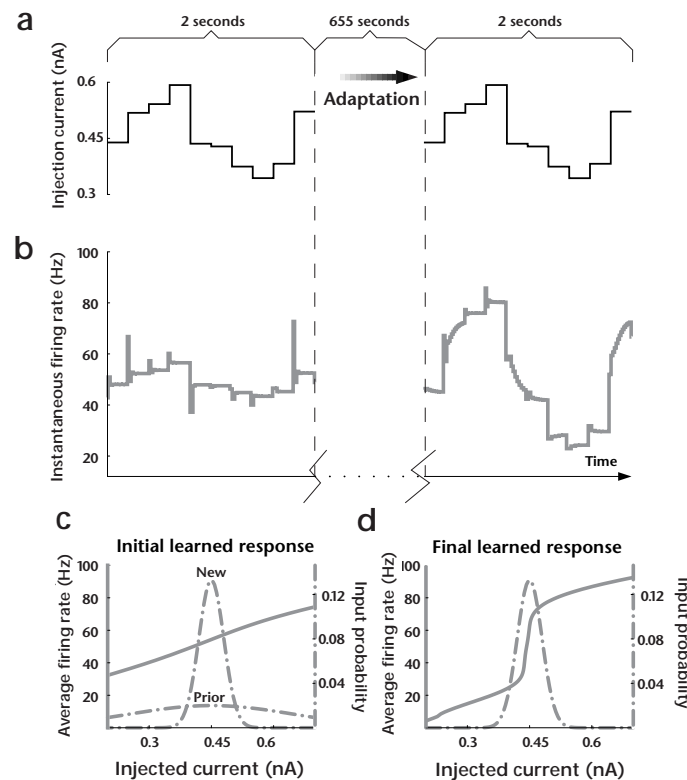


Fig. 5. An experimental prediction. **(a)** The model neuron was given a new random current injection step every 200 ms over the course of an hour-long experiment, to which the neuron learned to respond as shown in the left half of **(b)**. **(c)** A sudden six-fold decrease in the standard deviation of the applied current found the model neuron mismatched to the new 'stimulus environment'. The adaptation rule, however, enabled the neuron to adjust its dendritic conductances within 11 minutes, as can be seen by comparing the neuron's firing rate in response to the same stimulus current sequence before and after adaptation in **(b)**. **(d)** The neuron's average response curve as a function of current became considerably steeper, reflecting the increased output bandwidth accorded to the narrower input distribution.

Each guess is based solely on the stimulus experienced at that instant and can, therefore, be wildly inaccurate. Averaged over time, however, the sum of these guesses is guaranteed to point in the right direction of change.

Given that a neuron need not compute the mutual information explicitly, does it need to 'know' the arrival rates of photons impinging on the retina or the frequencies of sound waves hitting the ear's tympanic membrane to maximize the information transfer? As the voltage-dependent conductances do not directly sense the original external stimuli, but rather only sense voltages, the answer to this question is, fortunately, no. In the limit of vanishing multiplicative noise, maximizing the information between the firing rates and the dendritic voltages is equivalent to maximizing sensory information, as long as we can guarantee that transformation from stimulus to firing rate is always one-to-one.

In general, the learning rule for any one conductance depends on the values of all the other conductances, because the mutual information is a global property of the stimulus set. Dendritic voltage and ion conductances are nonlinearly coupled; conductances affect voltage, which, in turn, sets conductances. This coupling satisfies the physical requirement of charge conservation; when the neuron fires periodically, the average current into the neuron flowing

through the synaptic and the voltage-dependent dendritic conductances must equal the average current discharged by the neuron. This constraint fixes the sum of any set of nonlinear input currents to the dendritic compartment to a simple (and, in some cases, linear) function of the voltage. As the learning mechanism must incorporate how the current to the somatic compartment changes with dendritic voltage, charge conservation results in a learning mechanism that is strictly local, so that the mechanism for changing one conductance is independent of all other conductances.

These considerations allow us to derive a learning rule from first principles that decreases or increases the maximal dendritic calcium or potassium conductance \bar{g}_j by an increment $\Delta\bar{g}_j$ each time a stimulus is presented:

$$\Delta\bar{g}_j = \frac{\eta}{\bar{g}_j} \left\langle \frac{\delta I_j}{\delta V_d(t)} + c \langle V_d \rangle I_j \right\rangle \quad (1)$$

Here η is a learning rate, angular brackets indicate an average over the stimulus duration, $V_d(t)$ is the voltage in the dendritic compartment, I_j is the inward current through the j th conductance, $\delta I_j / \delta V_d(t)$ is the variational derivative of this current with respect to the instantaneous voltage, $V_d(t)$, and $c \langle V_d \rangle$ is a simple function that implements the constraints on the firing rates. Biophysically, $\Delta\bar{g}_j$ can be thought of as a change in the maximal conductance of the underlying ion channels or as a change in the membrane density of these channels. The learning rule uncouples the different conductance types; peak conductances for different types change simultaneously, yet independently, reflecting the local nature of the learning rule.

To illustrate the learning rule, imagine that a neuron contains three types of calcium conductances in the dendritic compartment, each one activating in a different voltage range (**Fig. 2b**). For didactic purposes, we neglect conductance inactivation present in the model of **Fig. 1c**. At any given time, any sustained, suprathreshold stimulus x that causes the cell to fire will lead to antidromic spike propagation (**Fig. 1a**) from the somatic to the dendritic compartment. The time-averaged voltage $\langle V_d \rangle$ is an increasing function of the stimulus x , and the full range of synaptic inputs produces a distribution of time-averaged voltages (**Fig. 2a**).

To simplify the explanation of the learning rule, suppose that the variational derivative $\delta I_j / \delta V_d(t)$ can be replaced by the usual derivative dI_j / dV_d , the slope of the j th current as a function of the average voltage. Such an approximation ignores fast fluctuations in the dendritic voltage caused by spiking in the somatic compartment but otherwise leaves the learning rule unchanged. If the average voltages encountered are far from the calcium reversal potential, the ionic current is the product of \bar{g}_j , the calcium activation function, and the reversal potential; hence, the current slope is proportional to the slope of the calcium activation function (**Fig. 2c**). In response to any particular voltage $V_d(t)$ resulting from synaptic input, the first term in the learning rule emphasizes the most voltage-sensitive conductances. In the simple scheme of **Fig. 2**, calcium conductances with the highest positive conductance slopes are increased maximally, thereby increasing the slope of the voltage response to the particular input x and, hence, the slope of the firing rate. The higher the slope, the greater the difference in firing rate between input x and its neighbor $x + \Delta x$. In other words, a greater portion of the range of firing rates is devoted to representing input x and its immediate neighbors. The change in the peak conductance, averaged over all inputs, is a function of the calcium conductance's midpoint voltage (**Fig. 2d**). Because conductance modification depends on the sign of the ionic current in the equation, potassium conductances

with activations similar to those in Fig. 2b generally change in the direction opposite to the change in calcium conductances.

The second term in the learning rule is the product of the average current through the j th conductance and the control term $c(\langle V_d \rangle)$, which enforces the firing rate constraints. Given the assumptions in the Methods section, the constraints on the firing rate translate into constraints on average voltage $\langle V_d \rangle$. Imposing a maximum and minimum firing rate implies that $c(\langle V_d \rangle)$ is zero for most commonly encountered voltages, a positive constant for voltages below some minimum and a negative constant for voltages above some maximum.

Different inputs x lead to different conductance changes. As the learning rule has two terms that balance each other, a conductance change driven by the first term and elicited by one input can be negated by the change due to another input that causes the second term to become dominant. In the end, only the most frequent inputs 'win' and thereby gain the largest representation in the output firing rate.

This learning rule—generalized to also change the midpoint voltage and the steepness of the activation and inactivation curves—was used to train the model neuron as it responded to random, 200-ms amplitude modulations of a synaptic conductance to the dendritic compartment (Fig. 3a). Even though the synaptic input was corrupted by additional background noise, the cell learned the input's statistical structure, matching its steady-state firing rate to the cumulative distribution function of the conductance inputs. The spike count distribution broadened to allow the nearly equal use of all firing rates (Fig. 3b and c), which more than doubled the total firing rate information (Fig. 3d). The information in the timing of spikes rose as well, increasing from 22.1 to 35.6 bits per second, as estimated by linear signal reconstruction analysis^{24–26}, while the average firing rate stayed nearly constant, decreasing from 42.3 Hz to 40.5 Hz. Learning proceeded swiftly; the model neuron approached optimal performance within two minutes from the start of learning (Fig. 3d). An animated illustration that shows the firing rate probability distribution changing in time can be found at <http://www.klab.caltech.edu/infomax>.

A similar learning rule maximized the mutual information given an average firing rate constraint of 30 Hz (Fig. 4). In contrast to a constraint on the peak firing rate, an average firing rate constraint should lead to an exponential distribution of firing rates²⁷, as was indeed the case (Fig. 4c). This restriction is associated with a constraint function $c(\langle V_d \rangle)$ that switches between a positive and a negative constant at a single, fixed value of V_d .

DISCUSSION

Using first principles, we have derived learning rules for changing the nonlinear conductances in a Hodgkin-Huxley neuronal model to maximize the information carried by firing rate. The rules for changing a conductance depend only on the voltage and the state of the conductance itself and hence do not violate the fundamental precepts of biophysical plausibility. However, unlike standard Hebbian learning or its variants, learning rules to maximize information have not been the subject of extensive experimental investigation so far.

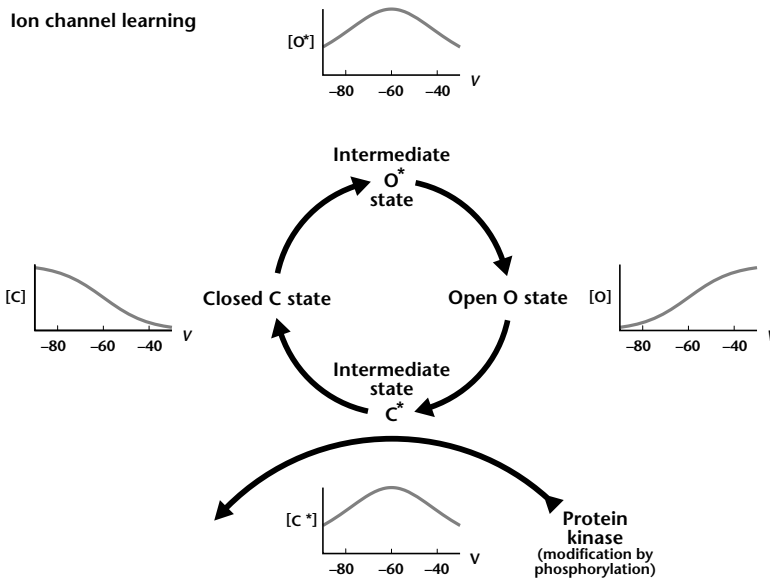


Fig. 6. Mechanism for state-dependent channel modification. In a standard symmetric energy-barrier model for an ion channel, two stable voltage-dependent states exist, labeled O and C. In moving from one state to the other, the ion channel passes through intermediate, metastable states marked by asterisks in the diagram. The fraction of ion channels in any particular state is a function of voltage, such that $[O]$ scales as $\langle m[V(t)] \rangle$ and $[C]$ as $\langle 1 - m[V(t)] \rangle$, and the occupancy of the intermediate states scales as the square root of $\langle \delta m[V(t)] / \delta V(t) \rangle$, as displayed in the inset diagrams of the state 'concentrations'. One possible mechanism for voltage-dependent, channel-by-channel modification is to restrict the interaction between the ion channel and protein kinases, G proteins or other second messengers. Such a restriction occurs if the interaction domain of the ion channel protein becomes accessible to a phosphorylating kinase, for instance, only when the protein is in a particular configuration, that is, in one and only one voltage-dependent state. The graph shows one such possibility, namely that of a kinase that can only interact with an intermediate state of the ion channel. Using such a mechanism, the rate of conductance change can be made very nearly equal to the slope of the current through this particular ionic conductance, as required by the learning mechanism of Fig. 2.

Our learning rule gives rise to experimental predictions that can be evaluated by varying the statistics of afferent synaptic input to the cell. If the changes in membrane properties can be initiated by changes in the membrane potential and do not require synaptic input to trigger a second-messenger cascade, then such an experiment can be carried out using standard current-clamp methods *in vitro*. In Fig. 5, after the neuron had adjusted to one particular Gaussian distribution of injection currents, the distribution was made narrower. The neuron responded to this change in its input distribution by increasing its gain. Retinal ganglion cells similarly adapt not only to the stimulus mean, but also to the variance of the stimulus²⁸.

As the slope of the firing rate response curve increased, the pattern of conductances in the model's dendritic compartment rearranged itself to give greater prominence to the high-threshold, non-inactivating Ca^{2+} conductances and low-threshold, transient K^+ conductances; this mix of conductances diminished net inward current at low voltages, whereas it augmented net inward current at high voltages, thereby forcing the response slope to increase. In contrast, as the slope of the response curve decreased, low-threshold inactivating Ca^{2+} conductances emerged, accompanied by high-threshold K^+ conductances that activate upon depolarization, reminiscent of the delayed rectifier conductance in the original Hodgkin-Huxley model. A shift, as opposed to a slope change, occurred in the firing rate response curve as the mean level of synap-

tic input increased, but the input variance remained constant. The neuron reduced its overall excitability either by increasing low-threshold K^+ conductances or by decreasing high-threshold Ca^{2+} conductances that previously had boosted the responses.

As only the summed modulatory current is uniquely determined, not the individual ionic currents, the set of conductances that optimize the firing rate responses for a given distribution of synaptic inputs is not unique. The contributions of different currents therefore depend on the idiosyncratic environmental history of the cell. Yet, for a fixed synaptic input distribution, the exact sequence in which these synaptic inputs arrive is not important in determining the pattern of conductances that emerges. Numerical simulations reveal that the conductance learning rule is general and powerful enough to handle exponential, bimodal and other non-Gaussian synaptic input distributions.

Adaptation to the variance of the input distribution using the learning rule mechanism can occur on the time scale of minutes, whereas adaptation to the mean input only requires several seconds. These time scales, which must be at least an order of magnitude longer than the typical stimulus durations, are simply lower limits on what is possible; the time scale of adaptation, inherent in the time-dependent learning rate η of the learning rule, can be much slower, depending on the underlying adaptation mechanism. For instance, the expression of new channel proteins is likely to require at minimum a few hours, whereas the phosphorylation of channels already present in the membrane can occur in seconds.

Adaptively learning to generate the proper mix of conductances implies that the cell can individually address the underlying ion channels that give rise to the conductances. Although the detailed substrate for maximizing information at both single-cell and network levels awaits investigation, it is known that cells can differentially distribute ion channels between the apical dendrites and cell bodies²⁹, and that certain forms of signal transduction are tied to the calcium influx through one particular calcium channel type, but are completely insensitive to the calcium influx through other channel types³⁰. Thus, neurons can selectively control different ion channel types.

The terms in the learning rule to maximize the transmitted information can be shown to map onto simple biophysical correlates. Consider, for instance, the simplest model of an ion channel that could underlie the conductances in Fig. 2b: stable open and closed states separated by a symmetric energy barrier and connected kinetically by intermediate, metastable states (Fig. 6). For such an ion channel, the $\delta I_j / \delta V_d(t)$ term is mirrored in the equilibrium rate of transitions between the open and closed state; in other words, the transition rate is proportional to the slope of the activation function raised to an exponent. Notably, the first term in the learning rule preferentially changes those conductances whose ion channels open and close most frequently.

To change the information transfer properties of the cell, a neuron could use state-dependent phosphorylation of ion channels or gene expression of particular ion channel subunits, possibly mediated by a second messenger cascade initiated by a G protein, to modify voltage-dependent conductances, as suggested in Fig. 6. Thus, single neurons possess the subcellular tools necessary to adaptively compress sensory information.

Our unsupervised learning algorithm enabled neurons to reflect the statistics of sensory input in an ongoing and dynamic manner. The theory behind the learning rule predicts that neural adaptation need not be restricted to purely synaptic mechanisms, but involves the regulation of dendritic or somatic voltage-dependent conduc-

tances at the single-cell level. These new possibilities for learning await experimental elucidation.

METHODS

The dynamics of the model are given by Hodgkin-Huxley-like equations that govern the membrane potential and a set of activation and inactivation variables, m_j and h_j , respectively. In each compartment of the neuron, the voltage V evolves as

$$C_m \frac{dV}{dt} = \sum_j \bar{g}_j m_j^{p_j} h_j^{q_j} (E_j - V) \quad (2)$$

where C_m is the membrane capacitance, \bar{g}_j is the peak value of the j th conductance, p_j and q_j are integers, and E_j is the ion-specific reversal potential. The variables m_j and h_j obey first-order kinetics of the type $dm/dt = (m_\infty(V) - m)/\tau(V)$, where $m_\infty(V)$ denotes the steady-state activation when the voltage is clamped to V and $\tau(V)$ is the voltage-dependent time constant. In general, Hodgkin-Huxley models can exhibit complex behaviors on several timescales, such as 'burst' firing patterns—sequences of multiple spikes interspersed with periods of silence. Choosing a standard model³¹ for the somatic compartment parameters, however, restricts the model to that of a regularly spiking cell, which adapts to a sustained stimulus by spiking periodically. The slowly inactivating A-type potassium current in this model allows spiking to occur at arbitrarily low firing rates. Firing rate adaptation is modeled by introducing a variable to describe accumulation of intracellular calcium during action potentials; this calcium concentration drives a potassium conductance whose maximal value scales roughly with the cell's firing rate. The dendritic compartment's calcium and potassium conductances have simple activation and inactivation functions described by distinct Boltzmann functions. Along with the peak conductance values, the Boltzmann functions' midpoint voltages $V_{1/2}$ and slopes s adapt to the statistics of stimuli. The midpoint voltages were initially distributed between -60 and $+20$ mV, with a difference of 20 mV between the midpoint for activation and inactivation; all slopes had an initial value of 10 mV. For simplicity, the dendritic conductances' time constants are set to a constant 5 ms. The reversal potential for the synaptic conductance in the dendritic compartment is set to 5 mV; for calcium, the reversal potential is set to 70 mV. Noise enters the model solely through the addition of Gaussian fluctuations to the synaptic input conductance; these fluctuations are white up to a cut-off frequency of 0.5 kHz and have the same net variance as the synaptic input conductance signal. Additional details and final, adapted parameter values are given at <http://www.klab.caltech.edu/infomax>.

All results presented in Figs. 3, 4 and 5 can also be reproduced in a minimal model with a single adjustable potassium conductance and a single adjustable calcium conductance. In general, additional conductances improve the flexibility of the model to adapt to different synaptic input patterns, but are not essential.

In the model, stimuli are taken to be maintained synaptic input conductances g_{syn} lasting 200 ms, drawn randomly from a fixed, continuous probability distribution, such as the Gaussian distribution in Fig. 3a with a mean and standard deviation of 141 ± 25 nS. All conductances in the model, including the synaptic input, are always restricted to remain non-negative.

After an initial transient phase, the dendritic membrane potential $V_d(t)$ oscillates with a period T dictated by the somatic spiking conductances. Thus, the input conductance g_{syn} maps onto a periodic voltage waveform $V_d(t)$, then onto an averaged current to the soma as shown in Fig. 1c and then finally onto an output firing rate f . Each step in this mapping is posited to be invertible.

Assuming that the uncertainty of encountering the time-averaged rate f in response to the sustained stimulus x can be adequately described by the standard deviation $\sigma_f(x)$ (the 'noise') in the firing rate in response to x , a lower bound for the mutual information³² is given by

$$I_{LB}(x, f) = -\int \ln(p(f) \sigma_f(x)) p(x) dx - \ln(\sqrt{2\pi e}) \quad (3)$$

To incorporate the constraints on the firing rate mathematically, we define a function $F = \beta \int E(f) p(f) df - I_{LB}(x, f)$. Setting $E(f) = f$ and taking β to be the energy cost per spike is equivalent to imposing a constraint on the mean firing rate; a minor change in the function $E(f)$, namely setting it

to zero for intermediate firing rates, approximates a constraint on the minimal and maximal firing rate. Given these constraints, minimizing F maximizes the mutual information bound $I_{LB}(x, f)$.

Each time a stimulus x occurs, a learning algorithm should continually modify the parameters for each of the voltage-dependent membrane conductances, such as the density of channels, in the direction that most decreases F . The neuron, however, has no biophysically plausible mechanism to compute the exact function F ; the only available option is a stochastic algorithm. Stochastic approximation²¹ adjusts a parameter, such as the peak conductance \bar{g}_j in equation (2), by $\Delta\bar{g}_j = -\eta\partial/\partial\bar{g}_j [\delta F/\delta p(x)]$ each time a stimulus is presented; this particular change will, by definition, occur with probability $p(x)$. If the only noise source is additive noise in the output firing rate, we can replace x by the dendritic voltage $\langle V_d \rangle$, based on a fundamental property of mutual information³³: information shared between firing rate and any invertible function of the stimulus equals that shared with the stimulus itself. In a biophysically more detailed model, a learning mechanism that increased channel density of voltage-dependent conductances would, of course, also increase the current noise. In that case, the information between the firing rate and the voltage time-course constitutes an upper bound on the information transmitted about the stimuli.

Calculating the partial derivatives in stochastic approximation requires an analytical expression for the steady-state current–discharge relationship at the soma. Because the somatic adaptation current is significant, subtractive and slow (with an intrinsic time constant of 50 ms), the current–discharge relationship possesses the self-consistent solution $f(I) = (I - I_0)/k$, where I is the current reaching the soma, I_0 is the threshold current, and k , the constant of proportionality, scales linearly with the strength of adaptation. The method of averaging applied to the adapted firing rate, current conservation and the chain rule of calculus allow one to compute the required variational derivatives, such as $\delta I/\delta V(t)$, that lead to the equation for adjusting the peak conductances \bar{g}_j and the corresponding equations for changing the midpoints and steepness of the activation and inactivation functions.

Five assumptions lead to the particularly simple form of the learning rule in the text. First, the model's firing rate is determined by the average current $\langle I \rangle$ reaching the somatic compartment during the interspike interval. Furthermore, the function relating the firing rate to the time-averaged input current, $f(\langle I \rangle)$, can be approximated by a linear polynomial in $\langle I \rangle$. This assumption happens to correspond to the measured current–discharge relationships of mammalian cortical neurons in experimental current injection studies, both in slice and in anesthetized animals, which tend to be linear or nearly linear^{34–37}. Second, the current $\langle I \rangle$ discharged through the coupling conductance G and the Hodgkin–Huxley spiking mechanism is proportional to the mean voltage $\langle V_d \rangle$ in the dendritic compartment. This assumption holds when the coupling conductance G is small compared to the somatic Hodgkin–Huxley conductances or, alternatively, in the presence of strong synaptic noise. Third, the time scale of each stimulus is long enough so that averaging over the steady-state interspike interval is equivalent to averaging over the stimulus interval. In other words, the perturbation introduced by averaging over the transient neuronal dynamics at stimulus onset is assumed to be weak. Fourth, the variance in the firing rate in response to any and all stimuli can be treated as though it had arisen from independent, additive noise. Fifth, the time constants of the dendritic modulatory conductances can be treated as constant over the range of voltages typically experienced in the dendritic compartment.

The timescale of learning can be varied by changing η in the learning rule. For Figs. 3 and 4, the learning rate decayed with time: $\eta(t) = \eta_0 \exp(-t/\tau_{\text{learning}})$, with $\eta_0 = 4.3 \times 10^{-3}$ and $\tau_{\text{learning}} = 4.4$ minutes in Fig. 3, whereas $\eta_0 = 1.0 \times 10^{-3}$ and $\tau_{\text{learning}} = 44$ minutes in Fig. 4. For Fig. 5, the initial Gaussian distribution of injected currents had its mean at 0.45 nA and a standard deviation of 0.2 nA.

ACKNOWLEDGEMENTS

This work was supported by the Alexander v. Humboldt Foundation, the Howard Hughes Medical Institute, the Deutsche Forschungsgemeinschaft, NIMH, ONR, NSF and the NSF-ERC Program at Caltech and was carried out in part at Caltech. We thank V. Lucic, F. Gabbiani, D. Schmitz and R. Stemmler for comments on the manuscript.

RECEIVED 3 DECEMBER 1998; ACCEPTED 31 MARCH 1999

- Tovée, M. J., Rolls, E. T. & Treves, A. Information encoding and the responses of single neurons in the primate temporal visual cortex. *J. Neurophysiol.* **70**, 640–654 (1993).
- Heller, J., Hertz, J. A., Kjaer, T. W. & Richmond, B. J. Information flow and temporal coding in primate pattern vision. *J. Comput. Neurosci.* **2**, 175–193 (1995).
- Baddeley, R. J. & Hancock, P. J. A statistical analysis of natural images matches psychophysically derived orientation tuning curves. *Proc. R. Soc. Lond. B Biol. Sci.* **246**, 219–223 (1991).
- Atick, J. J. Could information theory provide an ecological theory of sensory processing? *Network* **3**, 213–251 (1992).
- Ruderman, D. L. Statistics of natural images. *Network* **5**, 517–548 (1995).
- Laughlin, S. A simple coding procedure enhances a neuron's information capacity. *Z. Naturforsch.* **36**, 910–912 (1981).
- Laughlin, S. B., de Ruyter van Steveninck, R. R. & Anderson, J. C. The metabolic cost of neural information. *Nat. Neurosci.* **1**, 36–41 (1998).
- Baddeley, R. *et al.* Responses of neurons in primary and inferior temporal visual cortices to natural scenes. *Proc. R. Soc. Lond. B Biol. Sci.* **264**, 1775–1783 (1997).
- LeMasson, G., Marder, E. & Abbott, L. F. Activity-dependent regulation of conductances in model neurons. *Science* **259**, 1915–1917 (1993).
- Bell, A. J. Self-organisation in real neurons: Anti-Hebb in 'channel space'? *Neural Information Processing Systems* **4**, 59–67 (1992).
- Hebb, D. O. *The Organization of Behavior* (Wiley, New York, 1949).
- Davis, G. W. & Goodman, C. S. Synapse-specific control of synaptic efficacy at the terminals of a single neuron. *Nature* **392**, 82–86 (1998).
- Desai, N. S., Rutherford, L. C. & Turrigiano, G. G. Plasticity in the intrinsic excitability of cortical pyramidal neurons. *Nat. Neurosci.* **2**, 515–520 (1999).
- Moyer, J. R. Jr., Thompson, L. T. & Disterhoft, J. F. Trace eyeblink conditioning increases CA1 excitability in a transient and learning-specific manner. *J. Neurosci.* **16**, 5546–5546 (1996).
- Schreurs, B. G., Gusev, P. A., Tomsic, D., Alkon, D. L. & Shi, T. Intracellular correlates of acquisition and long-term memory of classical conditioning in Purkinje cell dendrites in slices of rabbit cerebellar lobule HVI. *J. Neurosci.* **18**, 5498–5507 (1998).
- Turrigiano, G., Abbott, L. F. & Marder, E. Activity-dependent changes in the intrinsic properties of cultured neurons. *Science* **264**, 974–977 (1994).
- Turrigiano, G., LeMasson, G. & Marder, E. Selective regulation of current densities underlies spontaneous changes in the activity of cultured neurons. *J. Neurosci.* **15**, 3640–3652 (1995).
- Purves, D. *Neural Activity and the Growth of the Brain* (Cambridge Univ. Press, New York, 1994).
- Gu, X. & Spitzer, N. C. Distinct aspects of neuronal differentiation encoded by frequency of spontaneous Ca²⁺ transients. *Nature* **375**, 784–787 (1995).
- Koch, C. *Biophysics of Computation: Information Processing in Single Neurons* (Oxford Univ. Press, 1998).
- Tsypkin, Y. Z. *Adaptation and Learning in Automatic Systems* (Academic, New York, 1971).
- Linsker, R. Local synaptic learning rules suffice to maximize mutual information in a linear network. *Neural Comput.* **4**, 691–702 (1992).
- Bell, A. J. & Sejnowski, T. J. An information maximization approach to blind separation and blind deconvolution. *Neural Comput.* **7**, 1129–1159 (1995).
- Bialek, W., Rieke, F., de Ruyter van Steveninck, R. & Warland, D. Reading a neural code. *Science* **252**, 1854–1857 (1991).
- Gabbiani, F. & Koch, C. Coding of time-varying signals in spike trains of integrate-and-fire neurons. *Neural Comput.* **8**, 44–66 (1996).
- Rieke, F., Warland, D., de Ruyter van Steveninck, R. & Bialek, W. *Spikes: Exploring the Neural Code* (MIT Press, Cambridge, Massachusetts, 1997).
- Reif, F. *Fundamentals of Statistical and Thermal Physics* (McGraw-Hill, 1965).
- Smirnakis, S. M. *et al.* Adaptation of retinal processing to image contrast and spatial scale. *Nature* **386**, 69–73 (1997).
- Hoffman, D. A., Magee, J. C., Colbert, C. & Johnston, D. K⁺ channel regulation of signal propagation in dendrites of hippocampal pyramidal neurons. *Nature* **387**, 869–875 (1997).
- Deisseroth, K., Heist, E. K. & Tsien, R. W. Translocation of calmodulin to the nucleus supports CREB phosphorylation in hippocampal neurons. *Nature* **392**, 198–202 (1998).
- Connor, J. A., Walter, D. & McKown, R. Neural repetitive firing: modifications of the Hodgkin–Huxley axon suggested by experimental results from crustacean axons. *Biophys. J.* **18**, 81–102 (1977).
- Stein, R. B. The information capacity of nerve cells using a frequency code. *Biophys. J.* **7**, 797–826 (1967).
- Pinsker, M. S. *Information and Information Stability of Random Variables and Processes* (Holden-Day, San Francisco, 1964).
- Granit, R., Kernell, D. & Shortess, K. S. Quantitative aspects of repetitive firing of mammalian motoneurons, caused by injected currents. *J. Physiol. (Lond.)* **168**, 911–931 (1963).
- Mason, A. & Larkman, A. Correlations between morphology and electrophysiology of pyramidal neurons in slices of rat visual cortex. *J. Neurosci.* **10**, 1415–1428 (1990).
- Jagadeesh, B., Gray, C. M. & Ferster, D. Visually evoked oscillations of membrane potential in cells of cat visual cortex. *Science* **257**, 552–554 (1992).
- Ahmed, B., Allison, J. D., Douglas, R. J. & Martin, K. A. An intracellular study of the contrast-dependence of neuronal activity in cat visual cortex. *Cereb. Cortex* **7**, 559–570 (1997).