

Temporal whitening by power-law adaptation in neocortical neurons

Christian Pozzorini, Richard Naud, Skander Mensi & Wulfram Gerstner

Spike-frequency adaptation (SFA) is widespread in the CNS, but its function remains unclear. In neocortical pyramidal neurons, adaptation manifests itself by an increase in the firing threshold and by adaptation currents triggered after each spike. Combining electrophysiological recordings in mice with modeling, we found that these adaptation processes lasted for more than 20 s and decayed over multiple timescales according to a power law. The power-law decay associated with adaptation mirrored and canceled the temporal correlations of input current received *in vivo* at the somata of layer 2/3 somatosensory pyramidal neurons. These findings suggest that, in the cortex, SFA causes temporal decorrelation of output spikes (temporal whitening), an energy-efficient coding procedure that, at high signal-to-noise ratio, improves the information transfer.

Neural signaling requires a large amount of metabolic energy¹. Consequently, neurons are thought to communicate using efficient codes in which redundant information is discarded². Theories of efficient coding³ successfully predict several features of sensory systems. At early stages of visual processing, inputs coming from the external world are decorrelated in both space and time^{4–7}; through sensory adaptation⁸, codes are dynamically modified so as to maximize information transmission^{9–12}; and sensory adaptation on multiple timescales^{11,13,14} could possibly reflect the statistics of the external world¹⁵.

Sensory adaptation is at least partially a result of intrinsic properties of individual neurons and, in particular, of SFA. SFA is not only observed at the early stages of sensory processing, but is also widespread in cortical neurons embedded in highly recurrent networks. Often modeled by a single process with one specific timescale^{16,17}, SFA also occurs on multiple timescales^{18–20}. In pyramidal neurons of the rat somatosensory cortex, three or more processing steps away from sensory receptors, SFA is scale free²¹, meaning that the effective speed at which individual neurons adapt is not fixed but depends on the input. Scale-free adaptation can be captured by simple threshold models with a power law–decaying spike-triggered process²² that possibly describes the combined action of Na⁺-channel inactivation^{23–25} and ion channels mediating adaptation currents^{26–28}.

Thus, three questions arise. First, can the temporal features of spike-triggered currents and spike-triggered changes in firing threshold, possibly spanning multiple timescales, be directly extracted from experimental data? Second, can SFA be explained by these spike-triggered effects? And finally, do the timescales of SFA match the temporal statistics of the inputs received by individual neurons? If temporal characteristics of inputs and SFA were matched, SFA could lead to a perfect decorrelation of the information contained in one spike with that of the previous one of the same neuron, a phenomenon known as temporal whitening²⁹. Temporal whitening in turn implies that, at a high signal-to-noise ratio (SNR), information transmission is enhanced³⁰.

RESULTS

The question of whether SFA is optimally designed for efficient coding can only be addressed if both the dynamics of SFA and the statistical properties of the inputs generated in biologically relevant situations are known. We used a combined theoretical and experimental approach to extract the dynamics of spike-triggered processes and SFA directly from *in vitro* recordings of cortical neurons. We then analyzed the synaptically driven membrane potential dynamics recorded *in vivo* from somatosensory neurons during active whisker sensation (data from ref. 31). Our overall goal was to study whether adaptation optimally removes the temporal correlations in the input to single neurons embedded in the highly recurrent network of the cortex.

SFA is mediated by two power-law spike-triggered processes

To reveal adaptation on multiple timescales, we stimulated layer 5 (L5) somatosensory pyramidal neurons with sinusoidal noisy currents of period T (500 ms to 16 s; Online Methods and Fig. 1). Single neurons responded with a firing rate $r(t)$ characterized by fast fluctuations around a sinusoidal mean $r_{\text{mean}}(t)$ given by the first-order approximation

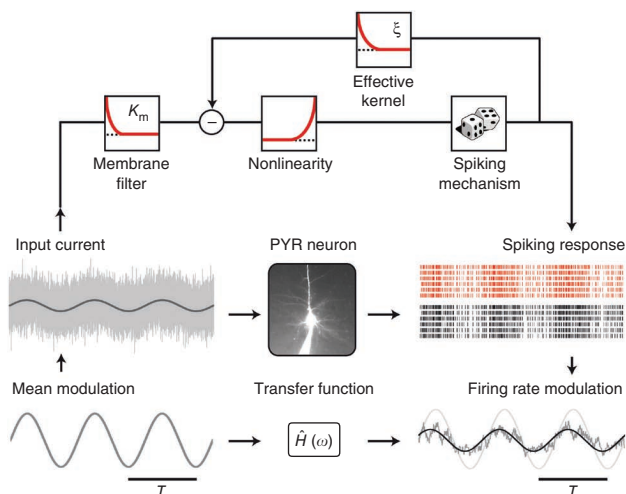
$$r_{\text{mean}}(t) = r_0 + \hat{H}_A(\omega) \cdot \Delta I_{\text{mean}} \sin(\omega t + \hat{H}_\Phi(\omega)) \quad (1)$$

where $\omega = 2\pi/T$ is the angular frequency of the input modulation, $r_0 \approx 4$ Hz is the average firing rate, $\hat{H}_A(\omega)$ is the amplitude response and $\hat{H}_\Phi(\omega)$ is the phase response. In the Fourier domain, the transfer function $\hat{H}(\omega) = \hat{H}_A(\omega) \exp(i\hat{H}_\Phi(\omega))$ constitutes a linear model for the modulation of the output firing rate (Fig. 1).

Given that SFA is at least partly a result of spike-triggered effects, the simple firing rate picture of equation (1) must be supported by a spike-based description. We therefore used intracellular recordings to fit a generalized leaky integrate-and-fire (GLIF) model (GLIF- ξ)

School of Computer and Communication Sciences and School of Life Science, Brain Mind Institute, École Polytechnique Fédérale de Lausanne, Lausanne, Switzerland. Correspondence should be addressed to C.P. (christian.pozzorini@epfl.ch).

Received 21 February; accepted 8 May; published online 9 June 2013; doi:10.1038/nn.3431



with escape-rate noise³² for stochastic spike emission (Fig. 1). To capture spike-triggered adaptation, the model featured an effective dynamic threshold, described by the function $\xi(s)$. This function (also known as the effective adaptation filter or kernel) summarizes the stereotypical sequence of biophysical events triggered by the emission of an action potential and accounts for both adaptation currents and physiological changes of the firing threshold. As the effects induced by consecutive spikes accumulate, the effective dynamic threshold produces SFA. Notably, the functional shape of $\xi(s)$, along with all other model parameters, was extracted from the intracellular recordings (see Online Methods and ref. 33).

Neocortical pyramidal neurons adapt their firing rates by means of two distinct biophysical mechanisms that increase the firing threshold and lower the membrane potential after each spike³³. To get an accurate estimation of the effective adaptation filter $\xi(s)$, we first fit a two-process GLIF model³³ (Supplementary Fig. 1) that explicitly featured both a dynamic threshold and an adaptation current, described by the filters $\gamma(s)$ and $\eta(s)$, respectively (Online Methods). In the model, the emission of action potentials depends only on the difference between the membrane potential and the firing threshold. Thus, spike-triggered currents $\eta(s)$ and movements of the firing threshold $\gamma(s)$ could be

Figure 1 Spiking neuron model GLIF- ξ and experimental protocol. To reveal SFA on multiple timescales, we repeatedly stimulated synaptically isolated L5 pyramidal neurons (PYR neuron) with fluctuating currents (input current) generated by adding filtered Gaussian noise to sinusoidal waves with different angular frequencies $\omega = 2\pi/T$ (mean modulation). The horizontal bars (bottom left and right) indicate the period T of modulation. The single neuron response (spiking response, black) was recorded intracellularly and the firing rate $r(t)$ was estimated by counting the number of spikes in each time bin (firing rate modulation, gray). The periodic oscillations of the firing rate $r_{\text{mean}}(t)$ (firing rate modulation, black) were related to the mean input (firing rate modulation, light gray) by a linear model defined in the Fourier domain by the transfer function $\hat{H}(\omega)$. We then used the intracellular recordings to fit the GLIF- ξ model (top). In this model, the input current is first low-pass filtered by the membrane filter $K_m(t)$ and then transformed into a firing intensity by an exponential nonlinearity. Spikes are emitted stochastically (spiking response, red) and trigger an adaptation process described by the effective adaptation kernel $\xi(s)$.

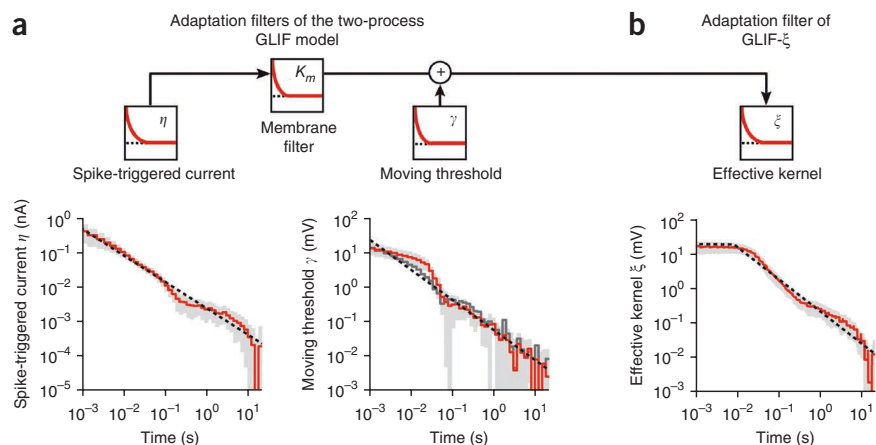
combined to obtain the effective adaptation filter $\xi(s)$ of the more parsimonious model GLIF- ξ (Online Methods and Fig. 2).

We found that, 22 s after the emission of an action potential, a small, but substantial, deflection remained in both the spike-triggered current $\eta(s)$ and the moving threshold $\gamma(s)$. Moreover, when displayed on log-log scales, the decay of both adaptation kernels was approximately linear over four orders of magnitude, meaning that both the adaptation current and the moving threshold were characterized by scale-free spike-triggered dynamics (Fig. 2a). Fitting $\eta(s)$ and $\gamma(s)$ with a power-law function (that is, $f_{\text{PL}}(s) = \alpha_f s^{-\beta_f}$) revealed that both spike-triggered processes have similar scaling exponents ($\beta_\eta = 0.76$, $\beta_\gamma = 0.87$). Consequently, the effective adaptation filter $\xi(s)$ was modeled by a truncated power law

$$\xi_{\text{PL}}(s) = \begin{cases} \alpha_\xi & , \quad \text{if } 0 < s < T_\xi \\ \alpha_\xi \cdot \left(\frac{s}{T_\xi}\right)^{-\beta_\xi} & , \quad \text{if } T_\xi < s < 22 \text{ s} \end{cases} \quad (2)$$

with parameters $\alpha_\xi = 19.2$ mV, $\beta_\xi = 0.93$ and $T_\xi = 8.3$ ms for the average kernel (Fig. 2b), and slightly different values for each individual cell (Supplementary Fig. 2), indicating that scale-free SFA is an intrinsic property of individual neurons.

Figure 2 Power-law adaptation filters extracted from *in vitro* recordings. (a) Adaptation filters of the two-process GLIF model. Left, mean spike-triggered current $\eta(s)$ (red) obtained by averaging the results of different cells ($n = 14$). The dashed black line shows the fit of a power-law function $\eta_{\text{PL}}(s) = \alpha_\eta s^{-\beta_\eta}$ with parameters $\alpha_\eta = 0.44$ nA, $\beta_\eta = 0.76$ and s in milliseconds. Right, mean moving threshold $\gamma(s)$ (red) obtained by averaging the results of different cells ($n = 14$). The dashed black line shows the fit of a power-law function $\gamma_{\text{PL}}(s) = \alpha_\gamma s^{-\beta_\gamma}$ with parameters $\alpha_\gamma = 24.4$ mV, $\beta_\gamma = 0.87$ and s in milliseconds. The dark gray line is a control showing an independent estimation of the average moving threshold $\gamma(s)$ obtained with an alternative fitting procedure (Online Methods). (b) The spike-triggered current $\eta(s)$ and the moving threshold $\gamma(s)$ were combined (block diagram) to obtain the effective adaptation filter $\xi(s)$ of the GLIF- ξ model. The mean adaptation filter $\xi_{\text{L}}(s)$ (red, GLIF- ξ_{L}) obtained by averaging the effective spike-triggered adaptation measured in individual cells ($n = 14$ neurons; Supplementary Fig. 2) is shown in red. The optimal fit of a truncated power-law $\xi_{\text{PL}}(s)$ (dashed black, GLIF- ξ_{PL}) yielded an exponent $\beta_\xi = 0.93$ (see equation (2)). In all panels, the gray area indicates 1 s.d. for the distribution of filters across different cells (asymmetric errors are a result of log scales).



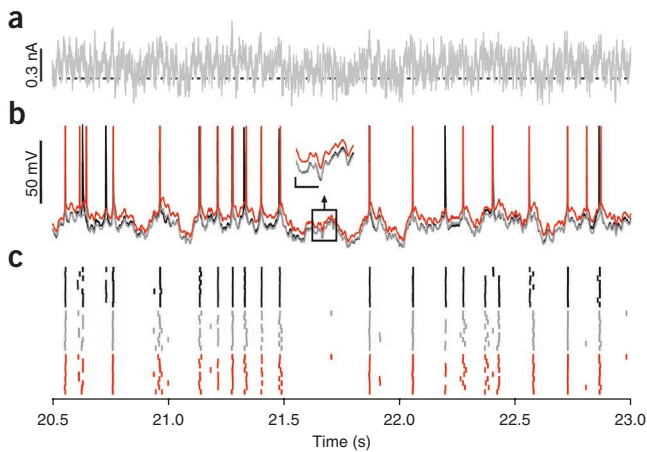


Figure 3 The GLIF- ξ_L model predicts the occurrence of single spikes with millisecond precision. (a) Typical 2.5-s segment of injected current. The same fluctuating current is presented several times. The dashed black line represents 0 nA. (b) The spiking response, but not the subthreshold membrane potential, predicted by the GLIF- ξ_L model (red) is consistent with the experimental data (black). In the two-process GLIF model (gray), where spike-triggered currents and threshold movements are modeled by two distinct processes, the dynamics of the subthreshold membrane potential predicted by the model is in excellent agreement with the experimental data. Inset, comparison of subthreshold membrane potential (scale bars represent 40 ms, 5 mV). (c) Spiking response of both the neuron (black) and the GLIF- ξ_L model (red) to repetitive presentation of the same current. By construction, the spiking response of the GLIF- ξ_L model is identical to that of the two-process GLIF model (gray).

We refer to the model with a single spike-triggered adaptation filter as GLIF- ξ_L , where ξ_L indicates that SFA is implemented by a 22-s-long filter obtained by combining the moving threshold and the spike-triggered current extracted from the experimental data. With the same logic, we use GLIF- ξ_{PL} to refer to a model in which the effective filter $\xi(s)$ is described by the truncated power law $\xi_{PL}(s)$ defined by equation (2) (a list with all GLIF- ξ_{PL} model parameters is given in **Supplementary Table 1**).

Power law SFA explains neural activity on short timescales

Valid single neuron models should predict the occurrence of individual spikes with millisecond precision³⁴. In response to a single injection of a fluctuating current (**Fig. 3a**), neurons emitted spikes that the GLIF- ξ_L model was able to predict with a high degree of accuracy (**Fig. 3b**). When the same current was injected repetitively, the spiking responses revealed the stochastic nature of single neurons: certain action potentials were emitted reliably with a high temporal precision, whereas others did not occur at each repetition or were characterized by larger temporal jitters. The GLIF- ξ_L model also captured this aspect (**Fig. 3c**). To validate our model, we quantified its predictive power using a similarity measure denoted M_d^* (Online Methods and ref. 35). On average, GLIF- ξ_L was able to predict more than 80% of the spikes ($M_d^* = 0.807$, s.d. = 0.04) with a precision of ± 4 ms (**Supplementary Fig. 3**). Statistically indistinguishable performance ($n = 12$ cells, paired t test, $t_{11} = 0.30$, $P = 0.77$) was achieved by GLIF- ξ_{PL} ($M_d^* = 0.804$, s.d. = 0.05), indicating that spike-triggered processes were well described by a truncated power law (**Supplementary Fig. 4**).

As expected, the subthreshold response observed *in vitro* was systematically overestimated by GLIF- ξ_L (**Fig. 3b**). This is explained by the fact that GLIF- ξ_L artificially translates spike-triggered currents

into effective threshold movements. In a two-process GLIF model, where adaptation currents and threshold movements are described as two distinct features (that is, when $\eta(s)$ and $\gamma(s)$ are not combined in a single effective kernel), model prediction of the membrane voltage and experimental data were in good agreement (**Fig. 3b**), confirming the validity of our fitting procedure. In terms of mere spike-timing prediction, the two-process GLIF model and the more parsimonious GLIF- ξ model were equivalent (**Fig. 3c**). Thus, we continued with just the single-process model GLIF- ξ . Overall, the spike time prediction procedure demonstrates the ability of both GLIF- ξ_L and GLIF- ξ_{PL} to capture the spiking activity of single neurons on the timescale of milliseconds.

Power law SFA explains neural activity on slow timescales

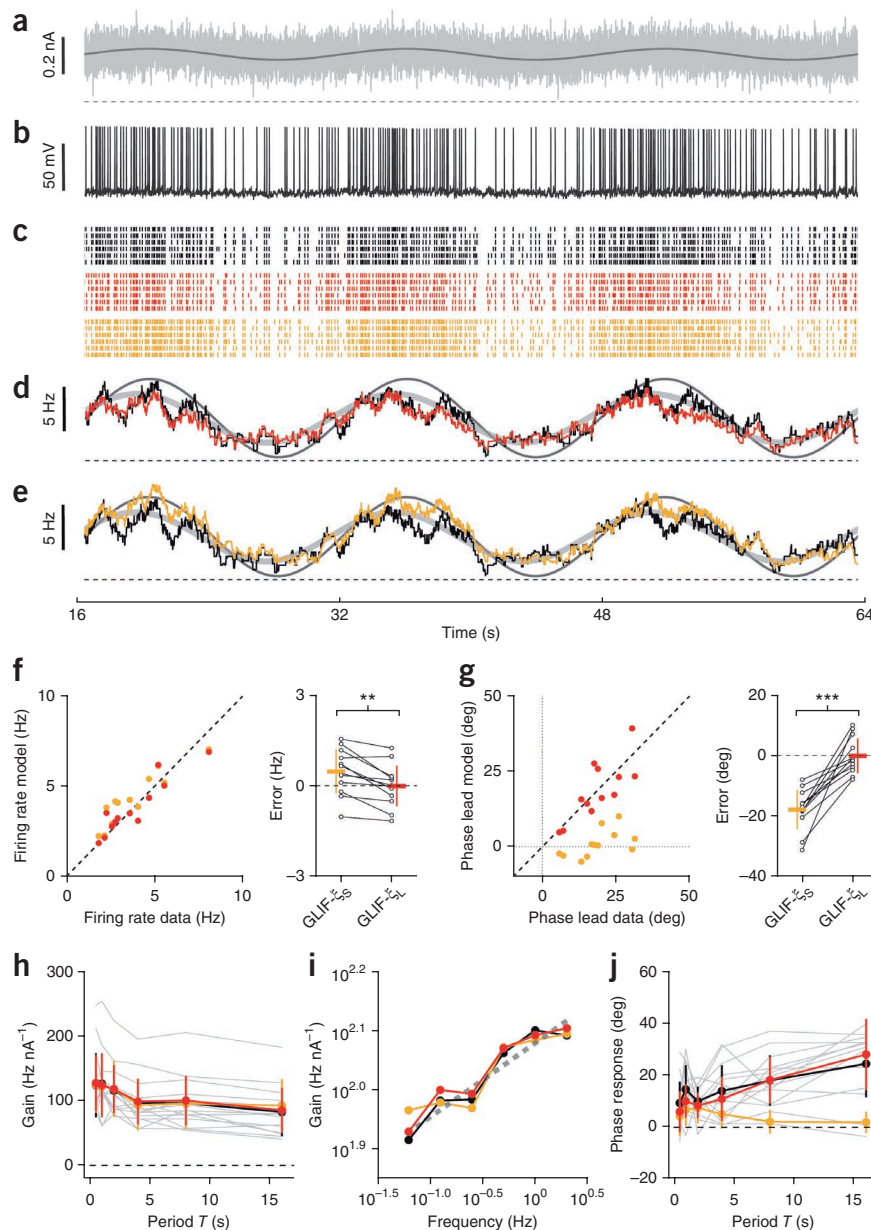
We asked whether the 22-s-long adaptation filter $\xi(s)$ could also predict the firing rate modulation on the much slower timescale of seconds. To this end, we used the GLIF- ξ_L model fitted on responses to different frequencies of modulation ($0.5 \leq T \leq 16$ s) to predict the firing rates recorded in the second part of the experiment, where one of the two slowest modulations ($T = 8$ or 16 s) was chosen and repetitively presented to the cell. A comparison of the raster plots obtained by injecting the same current in both the neuron and the GLIF- ξ_L model revealed that the spiking activity of the real neuron closely resembled the one predicted by the model (**Fig. 4a–c**). Furthermore, the match between the running-mean peristimulus time histograms (PSTHs) constructed for the model and the experimental data revealed that both responses shared a similar phase advance (**Fig. 4d**), indicating that our GLIF- ξ_L model is sufficient to capture the characteristic signature of SFA under slow sinusoidal stimulation²¹.

To study the role of the 22-s-long adaptation filter of GLIF- ξ_L , we then fitted the same single-process model under the assumption that the adaptation filter $\xi(s)$ had a duration of only 1 s (GLIF- ξ_S , for short adaptation filter). Compared with GLIF- ξ_L , the firing rate predicted by GLIF- ξ_S was in phase with the input and not with the spike output of the cells (**Fig. 4e**), indicating that GLIF- ξ_S was unable to capture the slow components of SFA (that is, the model with a short adaptation filter predicted an incorrect phase advance). To provide stronger evidence, we systematically quantified the ability to predict both the mean firing rate r_0 (**Fig. 4f**) and the phase lead \hat{H}_Φ (**Fig. 4g**). Although the GLIF- ξ_L model was capable of very good predictions that were in statistical agreement with the experimental data (errors $\Delta r_0 = -0.01$ Hz, s.d. = 0.67; $n = 12$ cells, Student t test, $t_{11} = -0.04$, $P = 0.97$; $\Delta \hat{H}_\Phi = -0.17$ deg, s.d. = 5.7; $n = 12$ cells, Student t test, $t_{11} = -0.10$, $P = 0.92$), GLIF- ξ_S tended to both overestimate the average firing rate and underestimate the phase advance (errors $\Delta r_0 = 0.47$ Hz, s.d. = 0.72; $n = 12$ cells, Student t test, $t_{11} = 2.16$, $P = 0.05$; $\Delta \hat{H}_\Phi = -17.9$ deg, s.d. = 6.5 deg; $n = 12$ cells, Student t test, $t_{11} = -9.16$, $P < 10^{-6}$), demonstrating that an adaptation filter of 1 s was not sufficient.

Finally, we measured the transfer function $\hat{H}(\omega)$ for both real neurons and spiking models by fitting equation (1) on the firing rates observed in response to six frequencies of modulation (**Fig. 4h–j**). For both real neurons and GLIF- ξ_L , the amplitude response $\hat{H}_A(\omega)$ was stronger at higher frequencies than at lower ones, revealing high-pass filtering, a characteristic feature of SFA (**Fig. 4h**). Consistent with observations in L2/3 pyramidal neurons²¹, plotting the amplitude response $\hat{H}_A(f)$ as a function of the input frequency $f = T^{-1}$ on log-log scales revealed that the gain of L5 pyramidal neurons was approximately power law (**Fig. 4i**). Moreover, the phase response $\hat{H}_\Phi(\omega)$ was always positive, indicating that, for all of the frequencies that we tested, the output firing rate led the input modulation (**Fig. 4j**). GLIF- ξ_L

Figure 4 The GLIF- ξ_L model accurately predicts the firing rate response on multiple timescales.

(a) Input current (gray) with slow mean modulation (dark gray) of period $T = 16$ s. (b) Membrane potential recorded in a single trial. (c) The firing activity (black) obtained by repetitive presentation of the same input current is compared with predictions of GLIF- ξ_L (red) and GLIF- ξ_S (orange). (d) Data from c were used to build two PSTHs (black, data; red, GLIF- ξ_L). The two sinusoidal functions represent the input modulation (dark gray) and the best fit of the experimental data (light gray). (e) Data are presented as in d, but with the prediction of GLIF- ξ_S (orange). (f) Performance in predicting the average firing rate r_0 of new stimuli. Left, model predictions are plotted against experimental data. Each dot represents a different cell. Right, each couple of open circles shows the prediction errors on the same cell. GLIF- ξ_L (red) was slightly more accurate than GLIF- ξ_S (orange) ($n = 12$, paired t test, $t_{11} = -4.09$, $P = 0.002$). (g) Performance in predicting the phase response \hat{H}_ϕ to inputs at $T = 8-16$ s. GLIF- ξ_L (red) outperformed GLIF- ξ_S (orange) ($n = 12$ cells, paired t test, $t_{11} = 6.31$, $P = 6.0 \times 10^{-5}$). Data are presented as in f. (h) Gain $\hat{H}_A(T)$ as a function of the period $T = 2\pi/\omega$. (i) Log-log plot of the gain $\hat{H}_A(f)$ as a function of the input frequency $f = T^{-1}$. Experimental data were fitted by a power law with scaling exponent $\beta_H = 0.12$ (dashed gray). (j) Phase response $\hat{H}_\phi(T)$ as a function of the period T . In h-j, data from individual cells ($n = 14$, gray lines) were averaged (black) and compared with the predictions of GLIF- ξ_L (red) and GLIF- ξ_S (orange). In all panels, error bars indicate 1 s.d. and horizontal dashed lines indicate zero. ** $P < 0.01$, *** $P < 0.001$.



was able to capture the features of the transfer function observed in L5 pyramidal neurons. Similar results were obtained with GLIF- ξ_{PL} (Supplementary Fig. 4), confirming that the spike-triggered processes observed *in vitro* were correctly modeled by a truncated power law lasting 22 s. Our experimental results (Fig. 4h-j) are very similar to those obtained in L2/3 pyramidal neurons²¹ and provide independent evidence for multiple timescales of adaptation. Overall, accounting for long-lasting spike-triggered effects with an appropriate adaptation filter is crucial for capturing the response of L5 pyramidal neurons on multiple timescales.

Power law SFA is optimally tuned for temporal whitening

Our model describes how the net current resulting from dendritic integration is encoded into a spike train at the somata of neocortical pyramidal neurons. To investigate the implications of power-law adaptation, we considered a situation in which a population of N uncoupled GLIF- ξ_{PL} neurons had to encode a common input $I(t) = I_0 + \Delta I(t)$ in the instantaneous firing rate $A(t)$, also known as the population activity. Note that, as the neurons in our population were all identical and received the same input, the population activity $A(t)$ is identical to the PSTH measured by repetitively injecting the same current into one cell. For relatively small fluctuations around a mean activity A_0 , we can assume that the population operates in a linear

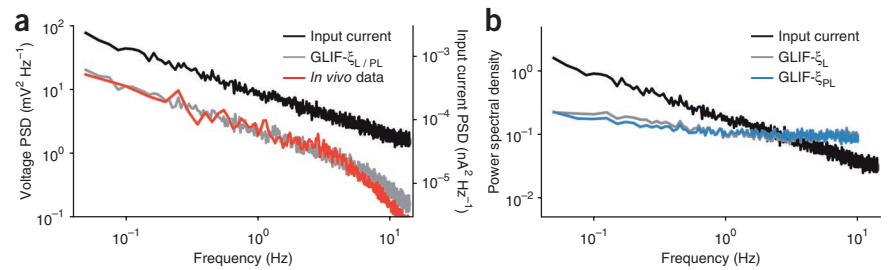
regime and responds to an external input fluctuation $\Delta I(t)$ according to the first-order approximation

$$A(t) = A_0 + \int_0^t \Delta I(t-s)H(s)ds + n(t) \quad (3)$$

where the impulse response $H(s)$ is the inverse Fourier transform of $\hat{H}(\omega)$, and the noise $n(t)$ is a result of stochastic firing in a finite population. Both terms depend on the intrinsic properties of the individual neurons and, in particular, on the precise shape of the adaptation filter $\xi(s)$.

For large populations, the noise term in equation (3) becomes negligible and optimal coding is achieved by the removal of temporal correlations potentially present in the input^{30,36}. This encoding strategy is known as temporal whitening and requires the population activity to have a flat power spectrum $A(f) = \text{constant}$ (Supplementary Modeling).

Figure 5 Power-law adaptation is near-optimally tuned to perform temporal whitening. **(a)** Power spectral density of the intracellular membrane potential fluctuations recorded *in vivo* from L2/3 pyramidal neurons (voltage PSD, red). The power spectrum was computed using 20-s-long recordings ($n = 57$) obtained from seven different cells (data from ref. 31). Fitting a power law (data not shown) on the frequency band $0.05 < f < 2$ Hz yielded a scaling exponent $\beta_1 = 0.67$. The power spectrum of the scale-free input used to stimulate a population of GLIF- ξ_{PL} neurons ($N = 100$) is shown in black (input current PSD). The power spectrum of the subthreshold response of individual GLIF- ξ_{PL} neurons (voltage PSD, gray) was in good agreement with the *in vivo* recordings. **(b)** The population activity of a group of GLIF- ξ_{PL} neurons in response to an *in vivo*-like input (black, copied from **a**) had a nearly flat spectrum $A(f)$ (blue). Similar results were obtained with GLIF- ξ_L neurons (gray). To allow a direct comparison between input and output powers, all the spectra shown in **b** were normalized to have the same total power.



SFA is known to implement high-pass filtering of the input current^{37,38}. In the case of power-law adaptation, the population response is characterized by a power-law gain (Fig. 4h,i and ref. 21), suggesting that, in neocortical pyramidal neurons, spike-triggered processes might be optimally tuned to efficiently encode scale-free signals (that is, signals that are temporally correlated across multiple timescales). However, the question of whether the functional role of power-law adaptation is to implement temporal whitening can only be answered if the statistical properties of the input received *in vivo* by neocortical pyramidal neurons are known.

To this end, we analyzed the synaptically driven membrane potential dynamics recorded from somatosensory pyramidal neurons during active whisker sensation (Online Methods). A spectral analysis performed on data from ref. 31 revealed that, at low frequencies, the power spectrum of the subthreshold membrane potential fluctuations was characterized by a power-law decay (Fig. 5a), indicating that natural stimuli received by somatosensory pyramidal neurons were indeed scale free.

To provide further evidence, we simulated the activity of a population of GLIF- ξ_{PL} neurons in response to an *in vivo*-like input characterized by a scale-free spectrum (Fig. 5a). The statistics of the subthreshold responses obtained in individual GLIF- ξ_{PL} neurons were consistent with the ones observed *in vivo* (Fig. 5a). Moreover, we found that the power spectrum of the population activity $A(f)$ (Fig. 5b) was much closer to a horizontal line than that of the input, indicating that a population of GLIF- ξ_{PL} neurons efficiently encodes *in vivo*-like signals by removing temporal correlations present in the input. Similar results were obtained with a population of GLIF- ξ_L neurons, where the adaptation filter $\xi(s)$ was not an idealized power law, but the average kernel extracted from intracellular recordings (Fig. 5b).

Overall, our results suggest that, in neocortical pyramidal neurons, power-law spike-triggered adaptation mirrors and approximately cancels the temporal correlations of signals generated in a biologically relevant situation. This result provides evidence for efficient coding at the level of single neurons embedded in the highly recurrent network of the cortex.

DISCUSSION

Neocortical pyramidal neurons are known to adapt their firing rate on multiple timescales^{20,21}. We found that SFA is a result of two separable spike-triggered mechanisms: each time an action potential is fired, both an adaptation current and a movement of the firing threshold are induced. Our results indicate that these spike-triggered effects are long (more than 20 s) and decay with a power law (Fig. 2), highlighting

the fact that SFA does not have a specific timescale. A GLIF model with an effective power-law spike-triggered process simultaneously captured both the fast dynamics critical for the prediction of individual spikes (Fig. 3) and the slow processes that modulate the firing rate (Fig. 4 and Supplementary Fig. 4). Notably, we found that, in behaving mice, the currents resulting from dendritic integration and received as input at the somata of pyramidal neurons were characterized by long-range temporal correlations that were partially removed by power-law spike-triggered adaptation (Fig. 5). This final observation indicates that, in cortex, power law SFA is near-optimally tuned for efficient coding.

Extent of spike-triggered effects

According to our results, an individual spike can still affect the firing activity of a neuron 20 s after its emission. It is possible that spike-triggered effects have an even longer duration. After 22 s, however, the magnitudes of both the moving threshold and the spike-triggered current were too small to be measured by our method (for $s > 20$ s, $\eta(s) < 0.1$ pA and $\gamma(s) < 0.01$ mV). As the effects of consecutive spikes accumulate, these small contributions shaped the single neuron response in a substantial way (Fig. 4f,g).

Although power-law adaptation was necessary to capture the firing rate fluctuations, a model with spike-triggered processes that only last for 1 s (GLIF- ξ_S) achieved very high performances ($M_d^* = 0.80$, s.d. = 0.03) in predicting the occurrence of individual spikes (Supplementary Fig. 3). This fact probably explains why power-law adaptation has not been observed in previous studies in which model validation was only based on spike time prediction.

Biophysical implementation of power-law adaptation

Our fitting procedure enabled us to discriminate between adaptation processes implemented by spike-triggered currents and physiological changes of the firing threshold. However, the biophysical details concerning the implementation of power-law dynamics were not part of our model. In principle, power-law relaxations can be approximated by a sum of exponentials covering a wide range of timescales^{20,22}. It is therefore likely that the spike-triggered current $\eta(s)$ that we found resulted from the combined action of multiple ion channels operating on different timescales, such as calcium-dependent, sodium-dependent and high voltage-activated potassium channels. Note, however, that a match of the relative strength of different currents implies a fine-tuned regulation of gene expression levels. Consistent with this hypothesis, multiple timescales of SFA have been previously modeled by biophysical models with several channels mediating adaptation currents^{20,21,29}. Alternatively, scale-free dynamics could

also be an intrinsic property of single channels. In particular, the power-law decay that we found in the moving threshold $\gamma(s)$ might reflect the scale-free dynamics observed during sodium-channel de-inactivation³⁹. In this alternative view, scale-free dynamics is likely to emerge from the presence of multiple inactivated states of ion channels^{19,40}.

All of our *in vitro* results are from mouse L5 neurons. We also investigated SFA in L2/3 and obtained very similar results (data not shown). In particular, we found that L2/3 pyramidal neurons adapted by means of power-law filters that closely resembled the ones observed in L5 and caused positive phase lead of the firing rate response to slow sinusoidal currents. These preliminary results suggest that L2/3 and L5 somatosensory pyramidal neurons share similar adaptation mechanisms. We also fitted GLIF models to the data of ref. 21 and found that both L2/3 and L5 pyramidal neurons of the rat somatosensory cortex adapted by means of spike-triggered power-law processes (data not shown), indicating that this mechanism is conserved across species and could be a common feature of cortical pyramidal neurons.

Functional implications

Both the moving threshold and the spike-triggered current extracted from *in vitro* recordings were characterized by power-law decays with very similar scaling exponents. This suggests that the particular shape of the adaptation filters is important. Neural signaling consumes a large amount of metabolic energy^{1,2}. The brain should therefore represent information using codes in which redundant information is discarded. According to efficient coding theory, optimality is achieved by adapting to the stimulus statistics and, at high SNR, by completely removing correlations that are potentially present in the signals to be encoded³. Efficient coding theory has been used to explain neural processing at early stages of the visual system. In the retina, center-surround receptive fields coupled with nonlinear processing strongly attenuate spatial correlations of natural images^{4,7}. Similarly, in primary visual cortex (V1), spatial decorrelation of features has been found¹⁵. In the temporal domain, neural firing was found to be decorrelated in the lateral geniculate nucleus of the cat⁶. Moreover, pyramidal neurons of V1 adapt on multiple timescales, providing further temporal decorrelation²⁹. However, it has remained unclear whether SFA serves for temporal redundancy reduction in the cortex.

To solve this issue, we estimated the statistics of input currents received *in vivo* at the somata of L2/3 somatosensory pyramidal neurons during active whisker sensation (data from ref. 31). This current, which reflects spatio-temporal statistics of spike arrivals at the synapses as well as subsequent filtering in the dendritic tree, can be seen as the driving current for spike generation. We found that input currents of pyramidal neurons did not have a preferred timescale, but were characterized by scale-free dynamics. Moreover, our numerical simulations showed that power-law spike-triggered processes were near optimally tuned to completely remove the temporal correlations revealed by the power-law decay of the input spectrum (Fig. 5). Overall, these results provide evidence for efficient coding in single neocortical neurons stimulated with behaviorally relevant signals.

The GLIF- ξ model implements a form of predictive coding. Indeed, the sum of adaptation processes, $\xi(s)$, triggered by past spikes can be interpreted as a linear predictor of the future input and further spiking only occurs when the real input exceeds the prediction. Consistent with our results, it has been shown that predictive coding of scale-free inputs by means of power-law spike-triggered kernels reduces the number of action potentials required to achieve a certain SNR⁴¹.

Temporal whitening versus noise shaping

For deterministic signals encoded in the absence of noise, efficient coding theory states that redundancy reduction is the optimal solution. However, in presence of noise, complete decorrelation can be detrimental. Redundancy can indeed improve the robustness of a code³⁰. To assess optimal coding in small populations of neurons, we must consider the noise term $n(t)$ associated with stochastic firing (equation (3)).

Previous studies have shown that non-renewal firing activity with negatively correlated interspike intervals can achieve higher information rates by noise shaping^{42–44}. In this coding strategy, the SNR is increased in the frequency band of the input signal by transferring the effective noise power to other frequencies (Supplementary Modeling). As already hypothesized⁴⁵, we found that, at low frequencies, spike-triggered adaptation resulted in a reduction of noise, which was completely counterbalanced by a similar modification of the gain controlling the amplitude of the signal, so that the SNR remained unchanged. Consequently, modifying the adaptation filter $\xi(s)$ did not affect the power spectrum of the effective noise (Supplementary Fig. 5), indicating that noise shaping is probably not the functional role of power-law adaptation. The question of how this result generalizes to different stimulation procedures is beyond the scope of this study.

In computational studies of memory and learning in neural networks, SFA is often neglected and, when considered, it is usually assumed to operate on short timescales. From our perspective, the power law of spike-triggered adaptation could be helpful in bridging the gap between the millisecond timescale of spike timing and behavioral timescales. Moreover, our results suggest that power-law adaptation causes temporal decorrelation of output spikes, a procedure that, at high SNR, improves information transfer.

METHODS

Methods and any associated references are available in the [online version of the paper](#).

Note: Supplementary information is available in the [online version of the paper](#).

ACKNOWLEDGMENTS

We thank C.C.H. Petersen, B.N. Lundstrom, G. Hennequin and A. Seeholzer for helpful discussions. We are also grateful to S. Crochet for sharing *in vivo* recordings and to B.N. Lundstrom for sharing the data that inspired this work. Finally, we thank S. Naskar for his help with the *in vitro* recordings. This project was funded by the Swiss National Science Foundation (grant no. 200020 132871/1; C.P. and S.M.) and by the European Community's Seventh Framework Program (BrainScaleS, grant no. 269921; S.M. and R.N.).

AUTHOR CONTRIBUTIONS

C.P. and R.N. conceived the study. C.P. designed the experiments, analyzed the data, performed the modeling and wrote the initial draft of the manuscript. S.M. contributed to data analysis and modeling. W.G. supervised the project. All of the authors worked on the manuscript.

COMPETING FINANCIAL INTERESTS

The authors declare no competing financial interests.

Reprints and permissions information is available online at <http://www.nature.com/reprints/index.html>.

1. Attwell, D. & Laughlin, S.B. An energy budget for signaling in the grey matter of the brain. *J. Cereb. Blood Flow Metab.* **21**, 1133–1145 (2001).
2. Laughlin, S.B. Energy as a constraint on the coding and processing of sensory information. *Curr. Opin. Neurobiol.* **11**, 475–480 (2001).
3. Barlow, H. Possible principles underlying the transformation of sensory messages. in *Sensory Communication* (ed. Rosenblith, W.A.) 217–234 (MIT Press, Cambridge, Massachusetts, 1961).
4. Srinivasan, M.V., Laughlin, S.B. & Dubs, A. Predictive coding: a fresh view of inhibition in the retina. *Proc. R. Soc. Lond. B Biol. Sci.* **216**, 427–459 (1982).
5. Dong, D. & Atick, J. Temporal decorrelation: a theory of lagged and nonlagged responses in the lateral geniculate nucleus. *Network* **6**, 159–178 (1995).

6. Dan, Y., Atick, J. & Reid, R.C. Efficient coding of natural scenes in the lateral geniculate nucleus: experimental test of a computational theory. *J. Neurosci.* **16**, 3351–3362 (1996).
7. Pitkow, X. & Meister, M. Decorrelation and efficient coding by retinal ganglion cells. *Nat. Neurosci.* **15**, 628–635 (2012).
8. Wark, B., Lundstrom, B.N. & Fairhall, A.L. Sensory adaptation. *Curr. Opin. Neurobiol.* **17**, 423–429 (2007).
9. Wainwright, M.J. Visual adaptation as optimal information transmission. *Vision Res.* **39**, 3960–3974 (1999).
10. Brenner, N., Bialek, W. & de Ruyter van Steveninck, R. Adaptive rescaling maximizes information transmission. *Neuron* **26**, 695–702 (2000).
11. Fairhall, A.L., Lewen, G.D., Bialek, W. & de Ruyter Van Steveninck, R.R. Efficiency and ambiguity in an adaptive neural code. *Nature* **412**, 787–792 (2001).
12. Maravall, M., Petersen, R.S., Fairhall, A.L., Arabzadeh, E. & Diamond, M.E. Shifts in coding properties and maintenance of information transmission during adaptation in barrel cortex. *PLoS Biol.* **5**, e19 (2007).
13. Baccus, S.A. & Meister, M. Fast and slow contrast adaptation in retinal circuitry. *Neuron* **36**, 909–919 (2002).
14. Ulanovsky, N., Las, L., Farkas, D. & Nelken, I. Multiple time scales of adaptation in auditory cortex neurons. *J. Neurosci.* **24**, 10440–10453 (2004).
15. Simoncelli, E.P. & Olshausen, B. Natural image statistics and neural representation. *Annu. Rev. Neurosci.* **24**, 1193–1216 (2001).
16. Izhikevich, E.M. Simple model of spiking neurons. *IEEE Trans. Neural Netw.* **14**, 1569–1572 (2003).
17. Brette, R. & Gerstner, W. Adaptive exponential integrate-and-fire model as an effective description of neuronal activity. *J. Neurophysiol.* **94**, 3637–3642 (2005).
18. Spain, W.J. & Schwindt, P. Two transient potassium currents in layer V pyramidal neurones from cat sensorimotor cortex. *J. Physiol. (Lond.)* **434**, 591–607 (1991).
19. Gilboa, G., Chen, R. & Brenner, N. History-dependent multiple-timescale dynamics in a single-neuron model. *J. Neurosci.* **25**, 6479–6489 (2005).
20. La Camera, G. *et al.* Multiple time scales of temporal response in pyramidal and fast spiking cortical neurons. *J. Neurophysiol.* **96**, 3448–3464 (2006).
21. Lundstrom, B.N., Higgs, M.H., Spain, W.J. & Fairhall, A.L. Fractional differentiation by neocortical pyramidal neurons. *Nat. Neurosci.* **11**, 1335–1342 (2008).
22. Drew, P.J. & Abbott, L.F. Models and properties of power-law adaptation in neural systems. *J. Neurophysiol.* **96**, 826–833 (2006).
23. Fleidervish, I.A., Friedman, A. & Gutnick, M.J. Slow inactivation of Na⁺ current and slow cumulative spike adaptation in mouse and guinea-pig neocortical neurones in slices. *J. Physiol. (Lond.)* **493**, 83–97 (1996).
24. Mickus, T., Jung, H.y. & Spruston, N. Properties of slow, cumulative sodium channel inactivation in rat hippocampal CA1 pyramidal neurons. *Biophys. J.* **76**, 846–860 (1999).
25. Melnick, I.V., Santos, S.F.A. & Saffronov, B.V. Mechanism of spike frequency adaptation in substantia gelatinosa neurones of rat. *J. Physiol. (Lond.)* **559**, 383–395 (2004).
26. Madison, D.V. & Nicoll, R.A. Control of the repetitive discharge of rat CA 1 pyramidal neurones *in vitro*. *J. Physiol. (Lond.)* **354**, 319–331 (1984).
27. Schwindt, P.C., Spain, W.J. & Crill, W.E. Long-lasting reduction of excitability by a sodium-dependent potassium current in cat neocortical neurons. *J. Neurophysiol.* **61**, 233–244 (1989).
28. Sanchez-Vives, M.V., Nowak, L.G. & McCormick, D.A. Cellular mechanisms of long-lasting adaptation in visual cortical neurons *in vitro*. *J. Neurosci.* **20**, 4286–4299 (2000).
29. Wang, X.J., Liu, Y., Sanchez-Vives, M.V. & McCormick, D.A. Adaptation and temporal decorrelation by single neurons in the primary visual cortex. *J. Neurophysiol.* **89**, 3279–3293 (2003).
30. Rieke, F., Warland, D., de Ruyter van Steveninck, R. & Bialek, W. *Spikes: Exploring the neural code* (MIT Press, Cambridge, Massachusetts, 1999).
31. Crochet, S., Poulet, J.F.A., Kremer, Y. & Petersen, C.C.H. Synaptic mechanisms underlying sparse coding of active touch. *Neuron* **69**, 1160–1175 (2011).
32. Gerstner, W. & Kistler, W. *Spiking Neuron Models: Single Neurons, Populations, Plasticity* (Cambridge University Press, New York, 2002).
33. Mensi, S. *et al.* Parameter extraction and classification of three cortical neuron types reveals two distinct adaptation mechanisms. *J. Neurophysiol.* **107**, 1756–1775 (2012).
34. Jolivet, R. *et al.* The quantitative single-neuron modeling competition. *Biol. Cybern.* **99**, 417–426 (2008).
35. Naud, R., Gerhard, F., Mensi, S. & Gerstner, W. Improved similarity measures for small sets of spike trains. *Neural Comput.* **23**, 3016–3069 (2011).
36. Atick, J. Could information theory provide an ecological theory of sensory processing? *Network* **22**, 4–44 (2011).
37. Benda, J. & Herz, A.V.M. A universal model for spike-frequency adaptation. *Neural Comput.* **15**, 2523–2564 (2003).
38. Köndgen, H. *et al.* The dynamical response of neocortical neurons to temporally modulated noisy inputs *in vitro*. *Cereb. Cortex* **18**, 2086–2097 (2008).
39. Toib, A., Lyakhov, V. & Marom, S. Interaction between duration of activity and time course of recovery from slow inactivation in mammalian brain Na⁺ channels. *J. Neurosci.* **18**, 1893–1903 (1998).
40. Lowen, S.B., Liebovitch, L.S. & White, J.A. Fractal ion-channel behavior generates fractal firing patterns in neuronal models. *Phys. Rev. E Stat. Phys. Plasmas Fluids Relat. Interdiscip. Topics* **59**, 5970–5980 (1999).
41. Bohte, S.M. & Rombouts, J.O. Fractionally predictive spiking neurons. *Neural Inf. Process. Syst.* **23**, 253–261 (2010).
42. Mar, D.J., Chow, C.C., Gerstner, W., Adams, R.W. & Collins, J.J. Noise shaping in populations of coupled model neurons. *Proc. Natl. Acad. Sci. USA* **96**, 10450–10455 (1999).
43. Shin, J. Adaptation in spiking neurons based on the noise shaping neural coding hypothesis. *Neural Netw.* **14**, 907–919 (2001).
44. Chacron, M.J., Lindner, B. & Longtin, A. Noise shaping by interval correlations increases information transfer. *Phys. Rev. Lett.* **92**, 080601 (2004).
45. Avila-Akerberg, O. & Chacron, M.J. Nonrenewal spike train statistics: causes and functional consequences on neural coding. *Exp. Brain Res.* **210**, 353–371 (2011).

ONLINE METHODS

In vitro electrophysiological recordings. All animal experiments were performed using published procedures^{46,47} in accordance with the rules of the Swiss Federal Veterinary Office. Briefly, somatosensory brain slices were obtained from postnatal day 14–18 wild-type mice (C57BL6/J) and whole-cell patch-clamp recordings were performed at 35 °C from L5 pyramidal neurons. The pipette solution consisted of 135 mM potassium gluconate, 4 mM KCl, 4 mM Mg-ATP, 10 mM sodium phosphocreatine, 0.3 mM Na₃-GTP and 10 mM HEPES (pH 7.3, 290 mOsm). During the experiments, we blocked all excitatory synaptic transmission by adding CNQX (20 μM) and D-AP5 (50 μM) to the bath solution. All electrophysiological data were low-pass Bessel filtered at 10 kHz and digitized at 20 kHz. Measurements were not corrected for the liquid junction potential. Recordings characterized by instabilities in the action potential shape or large drifts in the baseline firing rate r_0 were excluded from the data set upon visual inspection.

Current injections. To characterize single neurons with the standard tools of linear system analysis, we performed 64-s-long experiments in which noisy currents modulated by sinusoidal means were delivered in current-clamp mode. The injected current, denoted $I_{\text{ext}}(t)$, was generated according to the following equation

$$I_{\text{ext}}(t) = I_0 + \Delta I_{\text{mean}} \cdot \sin\left(\frac{2\pi}{T}t\right) + \Delta I_{\text{noise}} \cdot N(t) \quad (4)$$

where I_0 is a constant offset, ΔI_{mean} controls the amplitude of the sinusoidal mean and ΔI_{noise} defines the s.d. of the noise. The noise $N(t)$ was generated with an Ornstein-Uhlenbeck process with zero mean, unitary variance and a temporal correlation of 3 ms.

Each experiment consisted of many injections of currents generated according to equation (4). In the first half of the experiment (training set), we performed six injections using different periods of modulation $T \in \{0.5, 1, 2, 4, 8, 16\}$ in seconds. Stimuli were delivered in random order and, for each of the six injections, a new realization of the noise was used. In the second part of the experiment (test set), one of the two slowest modulations ($T = 8$ or 16 s) was chosen and more injections were performed. To assess the reliability of single neurons, the same realization of noise was used (frozen noise). Injections were performed with an interstimulus interval of 1 min.

Before and after each injection, we stimulated the neuron with two additional inputs. The first input was a 2.5-s-long current composed of a hyperpolarizing step followed, after 500 ms, by a suprathreshold step. We used the response to this stimulus to identify the neuronal type (L5 burst-generating cells were not included in the data set). The second input was a 4-s-long subthreshold noisy current generated with an Ornstein-Uhlenbeck process with zero mean and temporal correlation of 3 ms. We used this second injection to characterize the electrode response and perform active electrode compensation (see below and **Supplementary Data Preprocessing**).

At the beginning of each experiment, we tuned the input parameters I_0 , ΔI_{mean} and ΔI_{noise} to obtain a firing rate r_{mean} that oscillated periodically between 2 and 6 Hz. Typical values obtained after calibration were comprised in the range 100–450 pA for I_0 , 15–30 pA for ΔI_{mean} and 50–150 pA for ΔI_{noise} .

Linear analysis. For each neuron, we estimated the transfer function $\hat{H}(\omega)$ (Fig. 4h–j) using previously described methods^{21,38}. Briefly, the experimental spike train $\{\hat{t}_j\}$ was built by selecting the times at which the membrane potential $V(t)$ crossed 0 mV from below. We then obtained the firing rate $r(t)$ by building a histogram of the spike times. The bin size was such that each period of modulation T was divided into 30 bins. For each input frequency $\omega = 2\pi/T$, we finally obtained the transfer function by minimizing the sum of squared errors between the sinusoidal function $r_{\text{linear}}(t) = C_0 + C_1 \cdot \sin(\omega t + \phi)$ and the experimental firing rate $r(t)$, with $\{C_0, C_1, \phi\}$ being the only free parameters. The transfer functions of GLIF- ξ models (Fig. 4h–j) were obtained with the same method.

GLIF- ξ model. The spiking neuron models discussed are GLIF models equipped with a spike-triggered mechanism for SFA and with escape-rate noise for stochastic spike emission (Fig. 1). Spikes are produced according to a point process with conditional firing intensity $\lambda(t)$, which exponentially

depends on the momentary distance between the membrane potential $V(t)$ and the effective firing threshold $V_T(t)$ (ref. 48)

$$\lambda(t) = \lambda_0 \exp\left(\frac{V(t) - V_T(t)}{\Delta V}\right) \quad (5)$$

where λ_0 has units of s⁻¹ so that $\lambda(t)$ is in Hz and ΔV defines the sharpness of the threshold. Consequently, the probability of a spike occurring at a time $\hat{t} \in [t, t + \Delta t]$ is

$$P(\hat{t} \in [t, t + \Delta t]) = 1 - \exp\left(-\int_t^{t+\Delta t} \lambda(s) ds\right) \approx \lambda(t) \Delta t \quad (6)$$

In the limit of $\Delta V \rightarrow 0$, the model becomes deterministic and action potentials are emitted at the moment when the membrane potential crosses the firing threshold. For finite ΔV and a membrane potential at threshold (that is, when $V = V_T$), λ_0^{-1} defines the mean latency until a spike is emitted.

The subthreshold dynamics is modeled as a standard leaky integrator defined by the following ordinary differential equation for the membrane potential V

$$C\dot{V} = -g_L(V - E_L) + I_{\text{ext}} \quad (7)$$

where the three parameters C , g_L and E_L determine the passive properties of the membrane, the dot denotes the temporal derivative and I_{ext} is the injected current.

The dynamics of the effective firing threshold $V_T(t)$ in equation (5) is given by

$$V_T(t) = V_T^* + \sum_{\hat{t}_j < t} \xi(t - \hat{t}_j - T_{\text{ref}}) \quad (8)$$

where V_T^* is a constant, $\{\hat{t}_1, \hat{t}_2, \hat{t}_3, \dots\}$ are the times at which action potentials have been fired and $\xi(s)$ is an effective adaptation filter that accounts for all the biophysical events triggered by the emission of an action potential. According to equation (8), each time a spike is emitted, a threshold movement with stereotypical shape $\xi(s)$ is triggered, after a delay of absolute refractoriness T_{ref} . Threshold movements induced by different spikes accumulate and therefore produce SFA, if $\xi > 0$. For $s < 0$, we fixed $\xi(s) = 0$ so that only spikes in the past can affect the momentary value of the firing threshold. The adaptation filter $\xi(s)$ also accounts for adaptation processes mediated by spike-triggered currents. Consequently, $V_T(t)$ does not describe the physiological threshold (that is, the membrane potential at which action potentials are initiated *in vitro*), but has to be interpreted as a phenomenological model of spike-triggered adaptation. Finally, the functional shape of $\xi(s)$ was not defined a priori, but was obtained by combining the effects of both spike-triggered currents and movements of the physiological threshold, which were in turn extracted from the experimental data.

In principle, an absolute refractory period can be included in the adaptation kernel $\xi(s)$. Instead, we preferred to work with an explicit reset after a dead time. Each time a spike is emitted the membrane potential is reset to V_i and the numerical integration is restarted after a short period of absolute refractoriness T_{ref} . The GLIF- ξ model only differs from a generalized linear model^{49,50} as a result of this explicit reset.

The three GLIF- ξ models discussed differ in the duration and shape of the adaptation filter $\xi(s)$. In GLIF- ξ_L and GLIF- ξ_S , the functional shape of $\xi(s)$ is directly extracted from intracellular recordings and the duration of the adaptation filters are 22 s and 1 s, respectively. In GLIF- ξ_{PL} , the adaptation filter $\xi(s)$ is modeled as a truncated power law and lasts for 22 s.

Data preprocessing. *In vitro* recordings were preprocessed to remove the bias resulting from the voltage drop across the recording electrode. For that, we performed active electrode compensation⁵¹ following the procedure described in ref. 52. The electrode response was estimated before, during and after each 64-s-long injection. Consequently, we were able to remove experimental drifts resulting from slow changes in the electrode properties (**Supplementary Data Preprocessing and Supplementary Figs. 6–8**).

Fitting the GLIF- ξ model on *in vitro* recordings. To fit GLIF- ξ models, we extended the method introduced in ref. 33 by adding a hidden variable, $I_{\text{drift}}(t)$, able to absorb small drifts that are likely to occur in long recordings.

To get an accurate estimation of the effective adaptation filter $\xi(s)$, we first fitted a two-process GLIF model (**Supplementary Fig. 1**) that explicitly features both a spike-triggered current $\eta(s)$ and a spike-triggered movement of the firing threshold $\gamma(s)$ (**Fig. 2**). The effective adaptation filter $\xi(s)$, was then obtained by combining $\eta(s)$ and $\gamma(s)$ according to the following formula

$$\xi(t) = \int_0^\infty K_m(t-s)\eta(s)ds + \gamma(t) \quad (9)$$

where $K_m(s) = \Theta(s) \cdot R\tau_m^{-1}e^{-s/\tau_m}$ is the membrane filter, $\Theta(s)$ is the Heaviside step function, $R = g_L^{-1}$ and $\tau_m = RC$. The functional shapes of $\eta(s)$ and $\gamma(s)$ were not assumed a priori, but were directly extracted from the experimental data by the following two-step procedure.

In the first step, we extracted the functional shape of $\eta(s)$, together with all the parameters that determine the subthreshold dynamics, by fitting $\dot{V}(t)$ to the experimental voltage derivative $\dot{V}_{\text{data}} = [V_{\text{data}}(t + \Delta T) - V_{\text{data}}(t)] / \Delta T$, where $\Delta T = 0.05$ ms was given by the experimental sampling frequency. Given that adaptation currents directly affect the membrane potential dynamics, we fitted \dot{V}_{data} with the following model

$$C\dot{V} = -g_L(V - E_L) + I_{\text{ext}} - \sum_{\hat{t}_j < t} \eta(t - \hat{t}_j - T_{\text{ref}}) + I_{\text{drift}}(t) \quad (10)$$

where equation (7) was extended with a spike-triggered current $\eta(s)$ and the additional term $I_{\text{drift}}(t)$ is an unknown current that averages out to zero over time and captures experimental drifts during individual injections. To avoid any a priori assumption on the functional shape of the spike-triggered current, we defined $\eta(s)$ as linear combination of basis functions

$$\eta(s) = \sum_{k=1}^K \alpha_k f_k(s) \quad (11)$$

where the coefficients α_k control the shape of $\eta(s)$ and $f_k(s) = \text{rect}\left(\frac{s - T_k}{\Delta_k}\right)$ are rectangular functions of width Δ_k and centered at T_k . For GLIF- ξ_L , we used $K = 45$ log-spaced non-overlapping bins with Δ_k ranging from 0.5 ms to 4 s. For GLIF- ξ_S , we set $K = 30$ and $\Delta_k \in [0.5, 200]$ ms. Similarly, we defined $I_{\text{drift}}(t)$ as a piecewise constant function

$$I_{\text{drift}}(t) = \sum_{l=1}^L \beta_l \cdot \text{rect}\left(\frac{t - \Delta(l - 0.5)}{\Delta}\right) \quad (12)$$

For both GLIF- ξ_L and GLIF- ξ_S , we constrained $I_{\text{drift}}(t)$ to vary slowly in time by choosing a small number $L = 5$ of regularly spaced bins of size $\Delta = 12.8$ s.

As in refs. 33 and 53, given the injected current I_{ext} and the estimate of the membrane potential obtained after electrode compensation V_{data} , optimal parameters (minimizing the sum of squared errors between \dot{V}_{data} and \dot{V} of equation (10)) were obtained by solving a multilinear regression problem in discrete time. As GLIF models do not account for the action potential waveform, all of the data points $\{t | t \in [\hat{t}_j - 5 \text{ ms}; \hat{t}_j + T_{\text{ref}}]\}$ were excluded from the fit. Finally, we fixed the absolute refractory period at $T_{\text{ref}} = 2$ ms and obtained the voltage reset V_T^* by averaging the membrane potential measured T_{ref} milliseconds after the spikes.

Performing parameter extraction in presence of the term $I_{\text{drift}}(t)$ slightly improved the predictive power of the model (**Supplementary Fig. 3**). Note, however, that the term $I_{\text{drift}}(t)$ was not part of the model, but was only used in the fitting procedure to absorb slow changes in the subthreshold potential that could not be explained by spike-triggered processes.

Given the subthreshold dynamics, the second step consisted of estimating the parameters of the firing threshold. Given that spike-triggered currents were already captured by the filter $\eta(s)$, the effective threshold defined in equation (8) was replaced by

$$V_T^{\text{bio}}(t) = V_T^* + \sum_{\hat{t}_j < t} \gamma(t - \hat{t}_j - T_{\text{ref}}) \quad (13)$$

where $V_T^{\text{bio}}(t)$ describes the physiological threshold at which action potentials were initiated *in vitro*. In contrast to $\xi(s)$, $\gamma(s)$ is not a phenomenological model,

but describes physiological changes of the firing threshold triggered by the emission of previous spikes. Similarly to $\eta(s)$, we defined the moving threshold $\gamma(s)$ as a linear combination of rectangular basis function

$$\gamma(s) = \sum_{k=1}^K \delta_k f_k(s) \quad (14)$$

with $f_k(s)$ as in equation (11). Finally, the functional shape of $\gamma(s)$, along with the parameters V_T^* and ΔV , were extracted from experimental data by maximizing the log-likelihood of the observed spike-train⁵⁴

$$\log L(\theta) = \log p(\{\hat{t}_j\} | V; \theta) = \sum_{\hat{t}_j} \log \lambda_\theta(\hat{t}_j) - \int_{\Omega} \lambda_\theta(s) ds \quad (15)$$

where $\theta = \{\delta_1, \dots, \delta_K, V_T^*, \Delta V\}$ are the threshold parameters, $\Omega = \{t | t \notin [\hat{t}_j, \hat{t}_j + T_{\text{ref}}]\}$ is a set that excludes periods of absolute refractoriness and the conditional firing intensity $\lambda_\theta(s)$ is given by

$$\lambda_\theta(t) = \lambda_0 \exp\left(\frac{V(t) - V_T^* - \sum_{\hat{t}_j < t} \gamma(t - \hat{t}_j - T_{\text{ref}})}{\Delta V}\right) \quad (16)$$

where $V(t)$ was obtained by integrating equation (10) and, without loss of generality, we set $\lambda_0 = \Delta T^{-1}$. With the exponential function in equation (16), the log-likelihood to maximize is a concave function of θ (ref. 55). Consequently, the fit could be performed in discrete time using standard gradient ascent methods^{33,49,50}.

With this fitting procedure, an inaccurate estimation of the spike-triggered current $\eta(s)$ would affect the measure of the moving threshold $\gamma(s)$. To ensure that the estimation of $\gamma(s)$ that we obtained could indeed be attributed to a movement of the physiological threshold, we also extracted the threshold parameters using the experimental membrane potential V_{data} , rather than V (**Fig. 2a**).

Power-law fit of the effective adaptation filter $\xi(s)$. For GLIF- ξ_{PL} , the effective adaptation filter $\xi_L(s)$ extracted from the intracellular recordings was fitted with a truncated power-law $\xi_{PL}(s)$ (equation (2)). The fit was performed in two steps. First, we estimated the magnitude α_ξ and the scaling exponent β_ξ using a least-square linear regression performed in log-log space. For that, data points were logarithmically resampled and excluded from the fit if $\xi_L(s) < 5 \times 10^{-3}$ mV or $s < 5$ ms. Second, we obtained the cutoff T_ξ by calculating the intercept between the power-law fitted in the first step and the average value of the extracted kernel $\xi_L(s)$ computed on the first 5 ms. A similar procedure (that is, least-square linear regression in log-log space with logarithmically resampled points) was used for the power-law fit of the spike-triggered current $\eta(s)$ and the spike-triggered movement of the firing threshold $\gamma(s)$ shown in **Figure 2a**.

Performance evaluation. All of the performances reported in this study were evaluated on data sets that have not been used for parameter extraction. For the predictions reported in **Figures 3 and 4a-g**, the model fitted on the first half of the experiment (training set) was used to predict the responses observed in the second half (test set). Given that, in certain experiments, the average firing rates r_0 observed in the test set were slightly different than the ones of the training set, the parameter V_T^* was readjusted using the first 16 s of all the test set injections and models were validated on the responses recorded in the remaining 48 s. According to this procedure, models that do not capture SFA on slow timescales were expected to overestimate the average firing rate r_0 . For the predictions reported in **Figure 4h-j**, a leave-one-out strategy was used. In this case, models fitted on the responses to five different periods of modulation were used to predict the sixth one.

To evaluate spike-timing prediction, we used the similarity measure M_d^* introduced in ref. 35. M_d^* quantifies the similarity between two groups of spike trains generated by two stochastic processes and corrects the bias caused by the small number of available repetitions. M_d^* takes values between 0 and 1, where $M_d^* = 0$ indicates that the model is unable to predict any of the observed spikes and $M_d^* = 1$ means that the two groups of spike trains have the same instantaneous firing rate and are statistically indistinguishable. M_d^* can also be interpreted as the number of spikes correctly predicted (here with a precision of ± 4 ms) divided by an estimate of the number of reliable spikes.

Estimating the statistical properties of the input current received *in vivo* by neocortical pyramidal neurons. To test the hypothesis that power-law adaptation contributes to efficient coding by whitening the single neuron output, we estimated the power spectrum of the currents $\Delta I(f)$ received as input at the somata of neocortical pyramidal neurons *in vivo*. According to equation (10), in the absence of spikes, the membrane potential $\Delta V(t)$ is a low pass-filtered version of the input current, where the cutoff frequency $f = \tau_m^{-1}$ is defined by the membrane timescale. Consequently, at all frequencies $f \ll f_c$, we have $\Delta I(f) = \Delta V(f)/R^2$, with $\Delta V(f)$ being the power spectrum of the subthreshold membrane potential fluctuations and R being the cell resistance.

We estimated $\Delta V(f)$ using 20-s-long whole-cell recordings ($n = 57$) of the synaptically driven membrane potential dynamics obtained from seven different L2/3 pyramidal neurons of behaving mice (data from ref. 31). All the *in vivo* recordings were performed in primary somatosensory barrel cortex during active whisker sensation (see ref. 31 for more details). Spike-triggered currents last for more than 20 s and can in principle affect $\Delta V(f)$ even at very low frequencies. For this reason, only trials with low firing rates $r_0 < 0.5$ Hz were used. However, including recordings with $r_0 > 0.5$ Hz did not affect the results.

Simulating the population response to *in vivo*-like inputs. To obtain the results reported in Figure 5, we simulated a population of $N = 100$ unconnected GLIF- ξ_{PI} neurons in response to a 4,000-s-long current $I(t)$ characterized by a power spectrum $\Delta I(f) \propto f^{-\beta_I}$, with $\beta_I = 0.67$. Model parameters are given in Supplementary Table 1 and input currents were generated by numerically solving the following inverse Fourier transform²⁹

$$I(t) = I_0 + \Lambda \cdot \int_{-\infty}^{+\infty} \sqrt{\Delta I(f)} N(f) e^{i(2\pi ft + \phi(f))} df \quad (17)$$

where $N(f)$ is a Gaussian white-noise process, the phases $\phi(f)$ were independently sampled from a uniform distribution and the scaling factor Λ was adjusted to fit the power spectrum of the subthreshold membrane potential fluctuations observed *in vivo* (Fig. 5a). To avoid unrealistically large power at low frequencies, we introduced a cutoff $\Delta I(f) = 0$, for $f < 0.025$ Hz. The highest frequency in the

signal was determined by the time step $\Delta T = 0.5$ ms used for numerical simulations. The mean input I_0 was adjusted to obtain a plausible average activity of $A_0 = 4$ Hz, which was consistent with the firing rates obtained *in vitro*. Finally, the population activity $A(t)$ was constructed by counting the number of spikes falling in bins of 50 ms and its power spectrum $A(f)$ was computed using time series of 40 s.

Statistics. The number of cells used for the analysis ($n = 12$ or $n = 14$) was limited by experimental constraints. Data analysis only started after complete data collection and no data were excluded. Two-sided Student *t* test was used as a standard. Normality was verified using the Anderson-Darling test. Multiple comparison correction was not appropriate and was therefore not used.

46. Lefort, S., Tomm, C., Sarria, J.C.F. & Petersen, C.C.H. The excitatory neuronal network of the C2 barrel column in mouse primary somatosensory cortex. *Neuron* **61**, 301–316 (2009).
47. Avermann, M., Tomm, C., Mateo, C., Gerstner, W. & Petersen, C.C.H. Microcircuits of excitatory and inhibitory neurons in layer 2/3 of mouse barrel cortex. *J. Neurophysiol.* **107**, 3116–3134 (2012).
48. Jolivet, R., Rauch, A., Lüscher, H. & Gerstner, W. Predicting spike timing of neocortical pyramidal neurons by simple threshold models. *J. Comput. Neurosci.* **21**, 35–49 (2006).
49. Truccolo, W., Eden, U.T., Fellows, M.R., Donoghue, J.P. & Brown, E.N. A point process framework for relating neural spiking activity to spiking history, neural ensemble and extrinsic covariate effects. *J. Neurophysiol.* **93**, 1074–1089 (2005).
50. Pillow, J.W. *et al.* Spatio-temporal correlations and visual signaling in a complete neuronal population. *Nature* **454**, 995–999 (2008).
51. Brette, R. *et al.* High-resolution intracellular recordings using a real-time computational model of the electrode. *Neuron* **59**, 379–391 (2008).
52. Badel, L. *et al.* Dynamic I-V curves are reliable predictors of naturalistic pyramidal-neuron voltage traces. *J. Neurophysiol.* **99**, 656–666 (2008).
53. Paninski, L., Pillow, J. & Simoncelli, E. Comparing integrate-and-fire models estimated using intracellular and extracellular data. *Neurocomputing* **65**, 379–385 (2005).
54. Brillinger, D.R. Maximum likelihood analysis of spike trains of interacting nerve cells. *Biol. Cybern.* **59**, 189–200 (1988).
55. Paninski, L. Maximum likelihood estimation of cascade point-process neural encoding models. *Network* **15**, 243–262 (2004).

Temporal whitening by power-law adaptation in neocortical neurons

Christian Pozzorini, Richard Naud, Skander Mensi and Wulfram Gerstner

School of Computer and Communication Sciences and School of Life Science,
Brain Mind Institute,
Ecole Polytechnique Federale de Lausanne
1015 Lausanne EPFL, Switzerland

E-mail: christian.pozzorini@epfl.ch

Supplementary material

The Supplementary Material is organized as follows. In the section *Supplementary Data Pre-processing*, we show how known methods of Active Electrode Compensation (AEC) have been applied to preprocess *in vitro* recordings. The section *Supplementary Modeling* contains background information on the difference between temporal whitening and noise-shaping. This document also includes Supplementary Figures S1-S8 and Supplementary Table S1.

Figure S1: Two-process GLIF model fitted on intracellular recordings

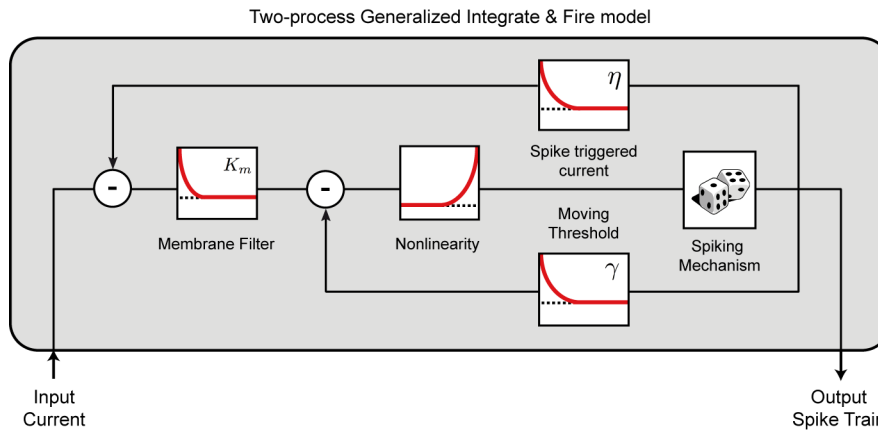


Figure S1: Schematic representation of the two-process GLIF model used to extract the spike-triggered current $\eta(s)$ and the moving threshold $\gamma(s)$ from intracellular recordings. As in GLIF- ξ , the input current I_{ext} is first low-pass filtered by the membrane kernel $K_m(s)$ and then transformed into a firing intensity by an exponential nonlinearity. Spikes are emitted stochastically and, in contrast to GLIF- ξ , trigger both a stereotypical current $\eta(s)$ and a movement of the firing threshold $\gamma(s)$.

Figure S2: Effective adaptation filters of individual L5 pyramidal neurons

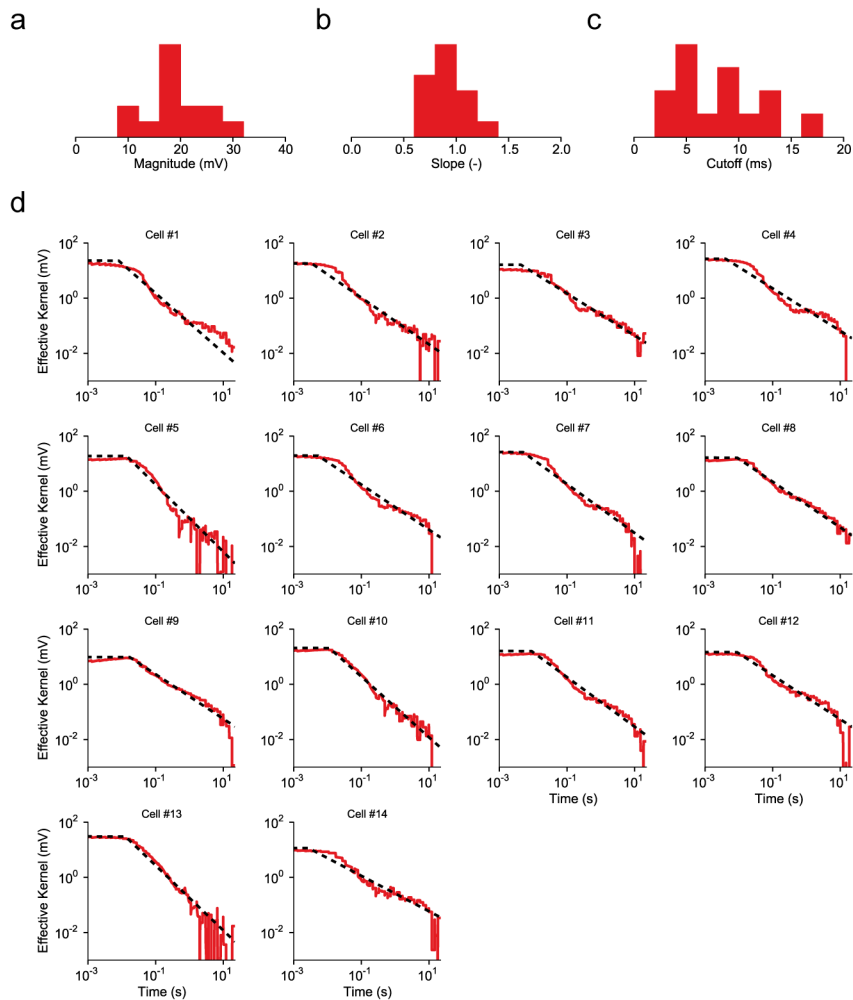


Figure S2: The 14 effective filters $\xi(t)$ measured in individual cells were fitted with the truncated power-law function $\xi_{PL}(t)$ (Eq. 2). **a:** Distribution of magnitudes α_ξ measured in different cells. **b:** Distribution of scaling exponents β_ξ measured in different cells. **c:** Distribution of cutoff values T_ξ measured in different cells. **d:** Effective filters $\xi(t)$ extracted from different cells (red) with optimal truncated power-law fit (dashed black). Each sub-panel corresponds to a different cell.

Figure S3: Influence of extending the fitting procedure with the hidden variable I_{drift}

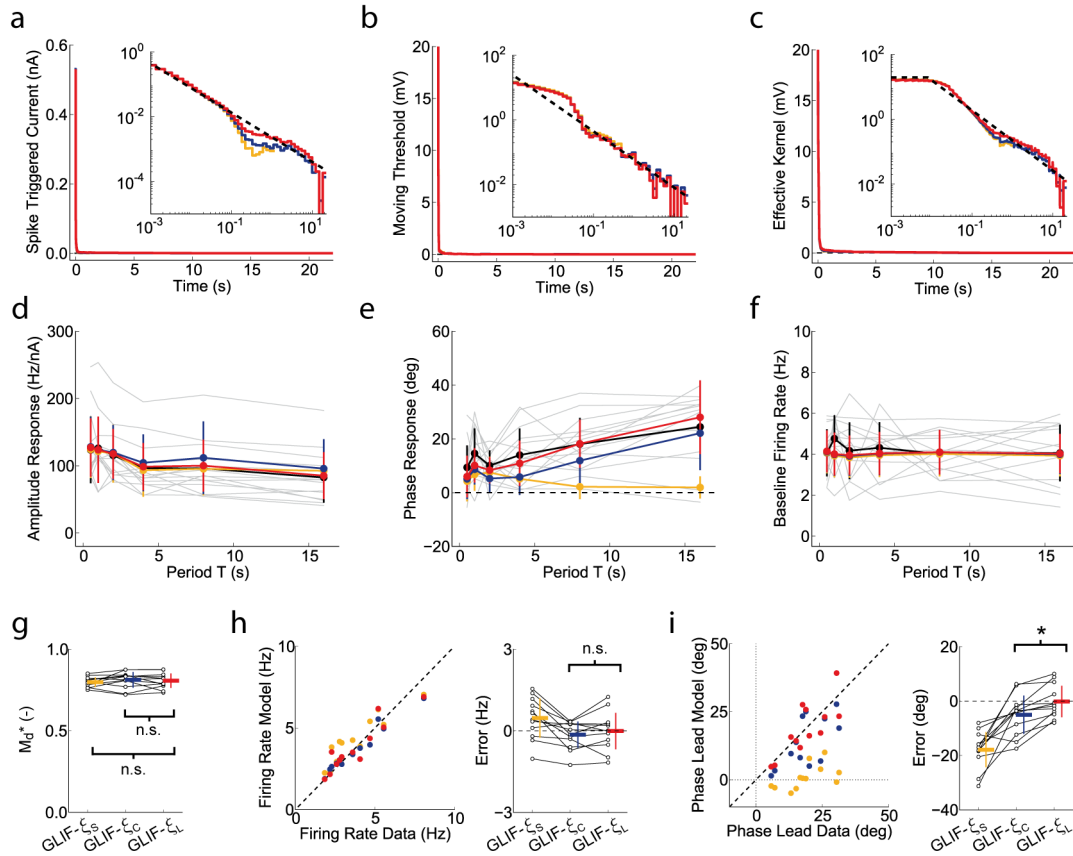


Figure S3: Extending the fitting procedure with an additional hidden variable I_{drift} improves the accuracy of parameter extraction. To assess the influence of the additional variable I_{drift} , we quantified the predictive power of a GLIF- ξ model fitted under the assumption that E_L does not change over time (i.e. the fit is performed without I_{drift}). In the following we will refer to this model as GLIF- ξ_C (where “C” stands for *control*). **a:** Spike-triggered current $\eta(t)$ of the model GLIF- ξ_C (blue). To allow for a comparison, the adaptation filters of the two models discussed in the main text (GLIF- ξ_S and GLIF- ξ_L) are plotted in orange and red, respectively. The inset of panel *a* shows that performing the fit with the auxiliary variable I_{drift} qualitatively changes the functional shape of the resulting spike-triggered current $\eta(t)$. **b:** The moving threshold $\gamma(t)$ is not affected. Colors as in *a*. **c:** Effective adaptation filter $\xi(t)$ obtained by combining the spike-triggered current and the moving threshold. Colors are as in *a*. **d-e:** Figures 4h and j are completed with the predictions of the GLIF- ξ_C model (blue). **f:** Control showing that the average firing rate r_0 does not depend on the period of modulation T (c.f. Eq. 1). Colors are as in panel *d* and *e*. **g:** The performance of different models in predicting the occurrence of individual spikes with a precision of ± 4 ms was quantified using the similarity measure M_d^* . Performance of GLIF- ξ_L were not significantly different from GLIF- ξ_C ($n = 12$ cells, paired t -test, $t_{11} = 0.58$, $p = 0.59$) and GLIF- ξ_S ($n = 12$ cells, paired t -test, $t_{11} = -0.82$, $p = 0.43$). **h:** Figure 4f is completed with the performance of GLIF- ξ_C ($n = 12$ cells, paired t -test, $t_{11} = -1.10$, $p = 0.29$). **i:** Figure 4g is completed with the performance of GLIF- ξ_C ($n = 12$ cells, paired t -test, $t_{11} = -4.17$, $p = 0.002$). In all panels, error bars indicate one standard deviation.

Figure S4: GLIF- ξ_L and GLIF- ξ_{PL} have a similar predictive power

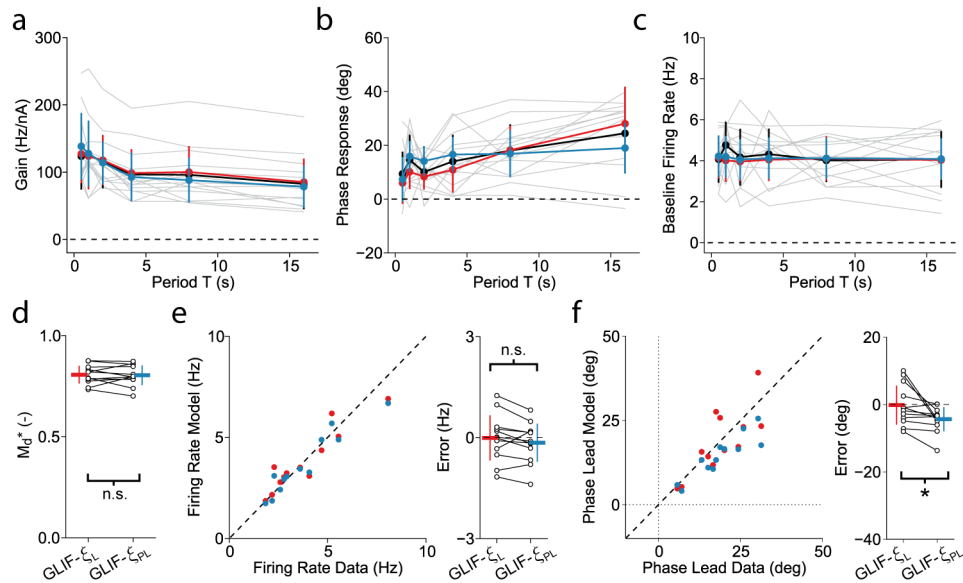


Figure S4: **a:** The amplitude response \hat{H}_A predicted by GLIF- ξ_L (red) is compared with that of GLIF- ξ_{PL} (light blue). **b:** The phase response \hat{H}_Φ predicted by GLIF- ξ_L is compared with that of GLIF- ξ_{PL} . In panels *a* and *b*, experimental data (black and gray) and GLIF- ξ_L predictions (red) are as in Figure 4h and 4j, respectively. **c:** Control showing that the average firing rate r_0 does not depend on the period of modulation T (c.f. Eq. 1). Colors are as in panel *a* and *b*. **d:** The performance of GLIF- ξ_L (red) and GLIF- ξ_{PL} (light blue) in predicting the occurrence of individual spikes with a precision of ± 4 ms was quantified using the similarity measure M_d^* . Each couple of open circles shows the performance of the two models on one cell. Performance of GLIF- ξ_{PL} ($M_d^* = 0.804$, s.d. = 0.05) were not significantly different ($n = 12$, paired t -test, $t_{11} = 0.30$, $p = 0.77$) from the ones obtained with GLIF- ξ_L ($M_d^* = 0.807$, s.d. = 0.04). **e:** Figure 4f is completed with the performance of GLIF- ξ_{PL} (average error $\Delta r_0 = -0.15$ Hz, s.d. = 0.57). Predictions of GLIF- ξ_L and GLIF- ξ_{PL} were not significantly different ($n = 12$ cells, paired t -test, $t_{11} = 1.80$, $p = 0.10$). **f:** Figure 4g is completed with the performance of GLIF- ξ_{PL} (average error $\Delta \hat{H}_\Phi = -4.4$ deg, s.d.=3.57). Predictions of GLIF- ξ_L and GLIF- ξ_{PL} were significantly different ($n = 12$ cells, paired t -test, $t_{11} = 2.73$, $p = 0.02$), however the difference was small. In all panels, the error bars indicate one standard deviation.

Figure S5: The functional shape of the effective adaptation filter does not shape the power spectrum of the effective noise

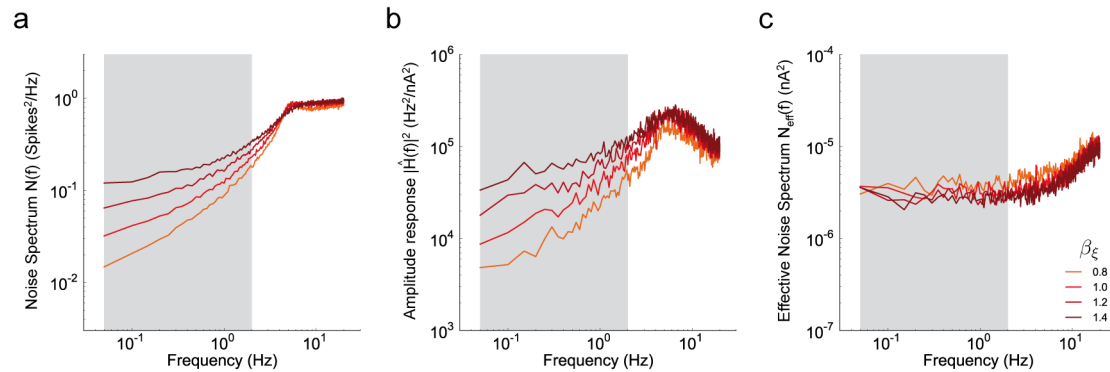


Figure S5: The precise shape of spike-triggered adaptation does not affect the power spectrum of the effective noise $N_{\text{eff}}(f)$. The squared amplitude response $|\hat{H}(f)|^2$ and the noise spectrum $N(f)$ of a small population of $M = 10$ uncoupled GLIF- ξ_{PL} neurons were numerically calculated for 4 different scaling exponents $\beta_{\xi} = \{0.8, 1.0, 1.2, 1.4\}$. The results were then used to compute the power spectrum of the effective noise defined as $N_{\text{eff}}(f) = N(f)/|\hat{H}(f)|^2$. To do so, the population activities $A_i(t)$ in response to repetitive injections of the same white-noise stimulus were simulated. Each neuron in the population received the same current. Furthermore, in all the simulations, the input current was tuned to evoke small fluctuations around a mean activity of $A_0 = 5$ Hz. Firing rates were computed by counting spikes in bins of 25 ms. **a:** The noise spectra $N(f)$ obtained for 4 different scaling exponents β_{ξ} are shown with 4 different colors. In the temporal domain, the noise is defined as $n_i(t) = A_i(t) - \langle A_i(t) \rangle_i$, where $A_i(t)$ denotes the population response to a single stimulation and $\langle \cdot \rangle_i$ denotes an average across repetitions i of the same injection. **b:** Squared amplitude response $|\hat{H}(f)|^2$ computed by dividing the power spectrum of the average population response $\langle A_i(t) \rangle_i$ by the constant defining the power spectrum of the white-noise input. **c:** The power spectrum of the effective noise $N_{\text{eff}}(f)$ is not affected by the scaling exponent β_{ξ} of the adaptation filter $\xi_{\text{PL}}(t)$. In the frequency band $0.05 < f \leq 2$ Hz (gray area) the effective noise spectrum is approximately flat, regardless of the value of the scaling exponent β_{ξ} . Indeed, changes in the noise spectra $N(f)$ are counterbalanced by similar changes in the amplitude response. Colors are the same in all panels and correspond to different scaling exponents β_{ξ} (see legend in panel c).

Figure S6: AEC removes artifacts on short timescales

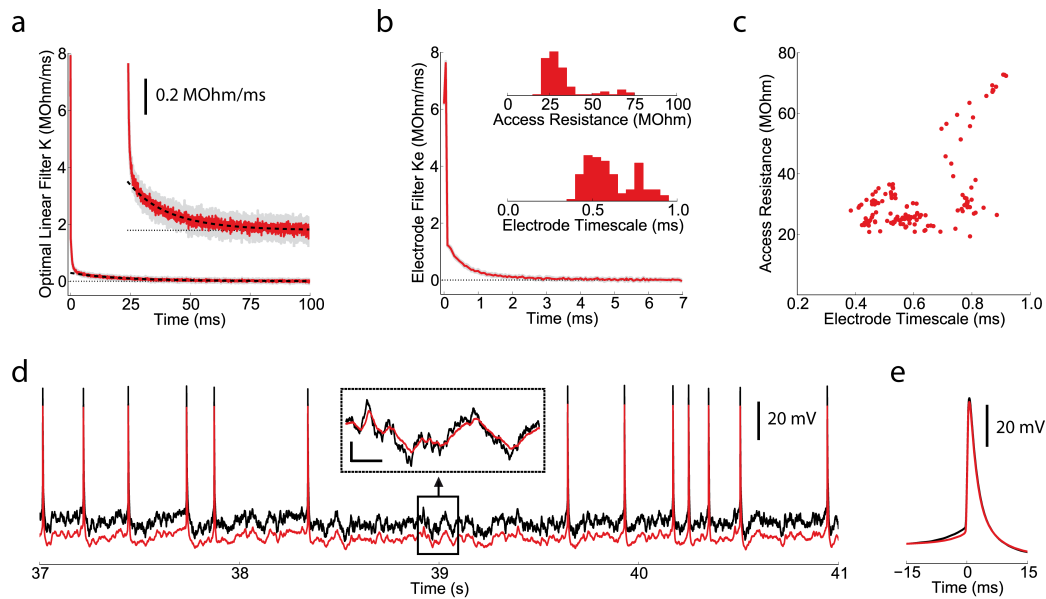


Figure S6: Active Electrode Compensation (AEC) removes artifacts on short timescales. **a:** Typical optimal filter $K^{(i)}(s)$ (red) obtained by averaging across bootstrap repetitions. The gray area represents one standard deviation. The tail of $K^{(i)}(s)$ is well fitted by an exponential function (dashed black line). The inset shows a zoom on the y-axis. **b:** For each bootstrap-repetition, the exponential fit from **a** (i.e. the putative membrane filter $K_m^{(i)}$) is subtracted from the optimal filter to obtain an estimate of the electrode filter. The electrode filters obtained in the 15 bootstrap-repetitions are then averaged to obtain the electrode filter $K_e^{(i)}(s)$ (red) used for AEC. The gray area represents one standard deviation. Each electrode filter $K_e^{(i)}(s)$ is characterized by its timescale (estimated by fitting an exponential function) and by the access resistance $R_e^{(i)} = \int_0^\infty K_e^{(i)}(s)ds$. The two distributions plotted in the inset show the electrode properties measured in all the recordings included in this paper. **c:** The access resistance is plotted as a function of the electrode timescale. This plot indicates that high access resistances are often associated with longer electrode timescales. **d:** For each injection i , the membrane potential V (red) is estimated by subtracting from the recorded signal V_{rec} (black) the potential drop across the electrode. Since the injected current has a baseline $I_0 > 0$, the membrane potential is, on average, lower than the recorded potential. This difference is given by $I_0 R_e^{(i)}$. Inset: zoom illustrating the fact that AEC acts as a low-pass filter to remove artifacts on the short timescales. The signals shown in the inset have been shifted to have the same mean. Scale bars: 5 mV and 30 ms. **e:** Average shapes of the action potentials obtained from V_{rec} (black) and V (red). The two traces have been shifted to have the same mean.

Figure S7: AEC removes artifacts on slow timescales

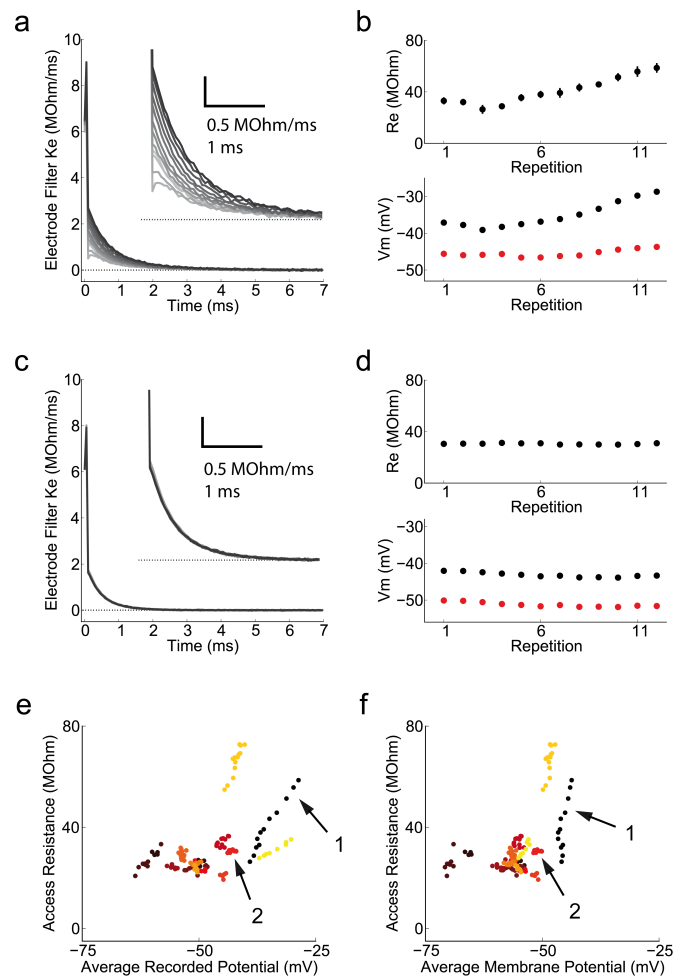


Figure S7: Active Electrode Compensation (AEC) removes artifacts caused by slow changes in the access resistance. According to our protocol, each experiment is divided in several 64-second long injections. AEC was performed using electrode filters $K_e^{(i)}$ estimated independently at each repetition (labeled i). **a:** Electrode filters K_e estimated in 12 consecutive stimulations (gray level increases from $K_e^{(1)}$ to $K_e^{(12)}$). In this specific experiment, the properties of the electrode filter clearly change over time. **b:** *Top:* the access resistance $R_e^{(i)}$ (computed by integrating the electrode filters in panel *a*) increases with time. This produces a drift in the recorded potential that we were able correct with AEC. *Bottom:* black and red dots show the average subthreshold potential computed using the recorded signal V_{rec} and the membrane potential V estimated with AEC, respectively. **c, d:** Same plots as in *a* and *b* showing the data of a typical experiment in which the electrode properties are stable. **e:** For each experiment included in this study, the access resistance is plotted as a function of the average subthreshold recorded potential. Groups of dots having the same color represent injections into the same neuron. Different colors represent different neurons. Most of the recordings are stable (in these cases the data points form a small cloud). Slow drifts in the recorded potential are always associated with changes in the access resistance. **f:** The access resistance is plotted as a function of the average membrane potential estimated with AEC. These results demonstrate the ability of AEC to compensate drifts due to changes in the access resistance. Arrows (1) and (2) indicate the two experiments shown in panels (a-b) and (c-d), respectively.

Figure S8: Electrode properties are stable during single injections

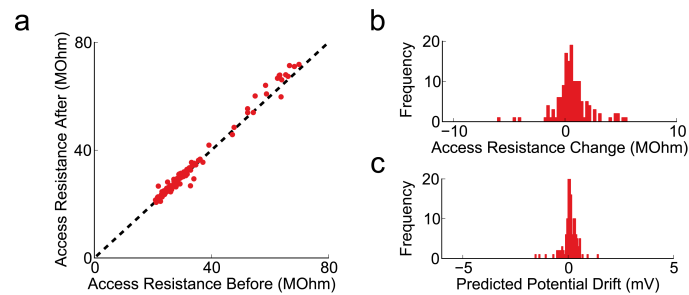


Figure S8: According to our protocol, each experiment consists of many injections of 64 seconds, labeled i . To remove artifacts due to changes in the electrode properties, a new electrode filter $K_e^{(i)}$ is estimated at each injection i . The assumption was made that the electrode filter does not change during single injections. To validate this hypothesis, the electrode filter was estimated before and after each injection using the response to 4-second long subthreshold noise (see Methods). **a:** Each dot shows the access resistance measured before and after each injection. **b:** Distribution of the changes $\Delta R_e^{(i)}$ in the access resistance observed in the experiments. $\Delta R_e^{(i)}$ is defined as the difference between the access resistance measured *after* and *before* the injection i . On average, the change in the access resistance was of $\Delta R_e = 0.62 \text{ M}\Omega$ ($s.d. = 1.54$). **c:** Histogram of membrane potential drifts expected to occur during single injections. The expected drift $\Delta V^{(i)}$ is computed by multiplying $\Delta R_e^{(i)}$ with the baseline current I_0 . On average, the expected drift was of $\Delta V = 0.09 \text{ mV}$ ($s.d. = 0.34$) confirming the hypothesis that electrode properties are sufficiently stable during individual injections.

Table S1: GLIF- ξ_{PL} model parameters extracted from experimental recordings of (n=14) L5 pyramidal neurons

Term	Description	Mean	S.d.	Units
R	Cell resistance	93.2	39.2	$\text{M}\Omega$
τ_m	Membrane timescale	15.3	7.1	ms
E_L	Reversal potential	-69.4	6.2	mV
V_T^*	Firing threshold baseline	-51.9	5.4	mV
ΔV	Firing threshold sharpness	0.75	0.15	mV
α_ξ	Magnitude of the effective adaptation filter ξ_{PL} [†]	19.42	5.72	mV
β_ξ	Scaling exponent of the effective adaptation filter ξ_{PL} [†]	0.90	0.17	-
T_ξ	Cutoff of the effective adaptation filter ξ_{PL} [†]	8.05	4.12	ms
T_{ref}	Absolute refractory period	2.0	-	ms
V_r	Reset potential	-38.8	9.0	mV

[†] The parameters obtained by fitting the average kernel shown in Fig. 2b are: $\alpha_\xi=19.3$ mV, $\beta_\xi=0.93$ and $T_\xi=8.3$.

1 Supplementary Data Preprocessing

To obtain an accurate estimate of the membrane potential, the recorded voltage was preprocessed with Active Electrode Compensation (AEC) [51] following the procedure described in ref. [52].

Background information on AEC

Due to the voltage drop across the electrode resistance (V_e), the potential recorded by a stimulating electrode (V_{rec}) is a biased version of the real membrane potential (V):

$$V(t) = V_{\text{rec}}(t) - V_e(t). \quad (\text{S1})$$

In AEC, the electrode is assumed to be an arbitrary linear system operating on the timescale of a few milliseconds so that $V_e(t)$ can be modeled as a filtered version of the injected current. If the input current I_{ext} has both a DC component I_0 and a time varying component $\delta I(t)$, then the electrode potential is described by the following equation

$$V_e(t) = I_0 R_e(t) + \int_0^t K_e(s, t) \delta I(t - s) ds, \quad (\text{S2})$$

where $K_e(s, t)$ is the electrode kernel at time t and $R_e(t) = \int_0^\infty K_e(s, t) ds$ is the electrode resistance (i.e. the *access resistance*). The argument t enables us to incorporate a potential slow drift of electrode parameters. The timescales on which electrodes operate are much faster than the timescales on which the electrode properties change. Consequently, the two terms on the right hand side of Equation S2 are responsible for slow-frequency and high-frequency artifacts, respectively. As it has already been shown, AEC removes high-frequency artifacts (Supplementary Fig. S6). Here, we show that this technique also enables to compensate slow-frequency artifacts due to changes in the access resistance (Supplementary Fig. S7).

Implementation

In practice, it is impossible to accurately estimate the electrode filter K_e at each moment in time. However, since changes in the electrode properties usually occur slowly, it was sufficient to estimate K_e within separable experimental blocks. In our protocol, each experiment consists of many 64-second long injections (i.e. experimental blocks). The assumption that the electrode properties are stable throughout each injection was verified by a systematic estimation of the electrode filter *before*, *during* and *after* each injection (Supplementary Fig. S8).

For each experimental block i , the electrode kernel $K_e^{(i)}$ was estimated following the procedure already used in ref. [52]. Briefly, far from spikes (i.e. in the subthreshold regime) we assume the neuron to act as a linear system described by the membrane filter $K_m^{(i)}$. Consequently, the

recorded potential can be modeled as a filtered version of the input current

$$V_{\text{rec}}(t) = V_0 + \int_0^t K^{(i)}(s) I_{\text{ext}}(t-s) ds, \quad (\text{S3})$$

where V_0 is the resting potential and $K^{(i)} = K_e^{(i)} + K_m^{(i)}$ accounts for both the electrode and the passive membrane. The filter $K^{(i)}$ was extracted from segments of subthreshold data by calculating the Wiener-Hopf optimal filter that provides the best estimate of the derivative of the recorded potential:

$$\dot{V}_{\text{rec}}(t) = \int_0^\infty K^{(i)}(s) \dot{I}_{\text{ext}}(t-s) ds. \quad (\text{S4})$$

The electrode kernel $K_e^{(i)}$ was then obtained by subtracting from the optimal filter an exponential function fitted on the tail of $K^{(i)}$. As in [51, 52], this tail is interpreted as the membrane filter $K_m^{(i)}$. To improve the accuracy, this procedure was repeated 15 times by resampling experimental data from the available subthreshold segments. The final estimate of $K_e^{(i)}(t)$ was obtained by averaging across bootstrap-repetitions. The maximal length of the two kernels $K^{(i)}(t)$ and $K_e^{(i)}(t)$ was set to 100 ms and 7 ms, respectively.

Results

Supplementary Figures S6a,b show a typical Wiener-Hopf filter $K^{(i)}(s)$ and a typical electrode filter $K_e^{(i)}(s)$, respectively. Each electrode filter was characterized by its timescale (estimated by fitting $K_e^{(i)}(s)$ with an exponential function) and by its access resistance $R_e^{(i)} = \int_0^\infty K_e^{(i)}(s) ds$. A distribution of the electrode properties measured in all the recordings included in this study are shown in Supplementary Figure S6b (inset) and Supplementary Figure S6c. As shown by a comparison between the recorded potential $V_{\text{rec}}(t)$ and the membrane potential obtained after Active Electrode Compensation $V(t)$ (Supplementary Fig. S6d,e), AEC acts as a low-pass filter by removing artifacts on short timescales.

As already mentioned, since the electrode properties were estimated *during* each individual experimental block i , we were also able to compensate artifacts that are due to changes in the *access resistance*. Supplementary Figures S7a,b illustrate a typical example in which AEC successfully removed a drift of the recorded potential induced by slow changes in the electrode filter. Note however that, in most cases, the electrode properties were stable throughout the whole experiment (a typical example of a stable recording is shown in Supplementary Figures S7c,d and summary data are presented in Supplementary Figures S7e,f).

Finally, the stability within individual injections was verified by estimating the electrode filter $K_e^{(i)}$ *before* and *after* each experiment block i . The results shown in Supplementary Figure S8 confirm our assumption that the electrode properties were stable within individual injections.

2 Supplementary modeling

In the Fourier domain, the linearized population model defined in Equation 3 can be written as

$$\Delta A(f) = \Delta I(f) \cdot |\hat{H}(f)|^2 + N(f), \quad (\text{S5})$$

with $\Delta A(f)$ being the power spectrum of the population activity fluctuations, $\Delta I(f)$ the power spectrum of the input fluctuations, $N(f)$ the noise spectrum and $|\hat{H}(f)|^2$ the squared amplitude response of the population. For populations of unconnected neurons, the noise produced by individual cells is independent. Consequently, the noise power is inversely proportional to the population size M (i.e. $N(f) \propto M^{-1}$).

In the following, we briefly review some known results useful for an understanding of how a population operating in a linear regime could optimally encode information. In particular we will discuss the case of small populations (where noise plays an important role) and large populations (where noise can be neglected).

Small populations: *Noise-Shaping*

Under the assumption that the input current $I(t)$ and the intrinsic noise $n(t)$ are both stationary Gaussian processes, the linearized population dynamics defined by Equation 3 constitutes a Gaussian channel (*Shannon, C. E. Communication in the Presence of Noise. Proc. IRE 37, pp. 10-21, 1949*). Consequently, the rate at which information is transmitted can be obtained by summing the signal-to-noise ratio (SNR) of different Fourier components according to the following formula [30]:

$$R_{\text{info}} = \frac{1}{2} \int \log_2(1 + \text{SNR}(f)) df, \quad (\text{S6})$$

where the signal-to-noise ratio is defined as $\text{SNR}(f) = |\hat{H}(f)|^2 \cdot \Delta I(f)/N(f)$. Alternatively, the same ratio can be written as $\text{SNR}(f) = \Delta I(f)/N_{\text{eff}}(f)$, where the power spectrum of the *effective noise* $N_{\text{eff}}(f)$ is obtained by putting the noise spectrum $N(f)$ back into the same units as the input signal:

$$N_{\text{eff}}(f) = \frac{N(f)}{|\hat{H}(f)|^2}. \quad (\text{S7})$$

The technique of *noise-shaping* consists of increasing the SNR at important input frequencies by transferring the effective noise $N_{\text{eff}}(f)$ to other bands. By doing so, systems can achieve higher information rates R_{info} . In particular, it is possible to prove that the *channel capacity* is optimally exploited when $\Delta I(f) + N_{\text{eff}}(f) = \text{Const.}$ [30]. This result is known as the *water-filling theorem* and tells that information transmission is maximized when the input and effective noise spectra complement each other.

As mentioned in the main text, our results indicate that changing the shape of the adaptation

filter $\xi_{\text{PL}}(t)$ does not affect $N_{\text{eff}}(f)$ (see Supplementary Fig. S5), suggesting that *noise-shaping* is not the functional role of power-law adaptation.

Large populations: *Temporal Whitening*

For large populations, the noise term $N(f)$ becomes very small and, in the limit of $M \rightarrow \infty$, encoding becomes deterministic. In such a situation the *noise entropy* vanishes and one can show that the information transmission is optimal when the entropy of the encoded signal is maximized. For this reason, at low noise, an optimal encoder should remove temporal correlations from the input signal [30, 36].

The idea of considering redundancy reduction as an optimization principle for neural coding is quite old and has been put forward by Barlow and others [3]. Experimental evidence was then provided by studies of the early visual system [4, 6]. In this paper, we extended these ideas to cortical neurons known to be embedded in a highly recurrent network [47]. In particular, we have shown that redundancy reduction provides an interpretation of power-law adaptation. Finally, it is worth remembering that redundancy reduction is only useful at low noise. Indeed, in highly noisy channels, redundancy can be exploited to provide robustness against noise [30].